
VOLTTRON Documentation

Release 6.0

The VOLTTRON Community

Aug 26, 2020

Contents

1	Features	3
2	Background	5
3	License	7
3.1	Overview	7
3.2	Join the Community	12
3.3	Installing VOLTTRON	17
3.4	Developing VOLTTRON	69
3.5	Core Services	292
3.6	Platform Specifications	456
3.7	Applications	518
3.8	Indices and tables	539



VOLTRON™ is an open-source platform for distributed sensing and control. The platform provides services for collecting and storing data from buildings and devices and provides an environment for developing applications that interact with that data.

Out of the box VOLTTRON provides:

- a secure *message bus* allowing agents to subscribe to data sources and publish results and messages.
- secure connectivity between multiple instances.
- BACnet, ModBus and other device/system protocol connectivity through our *driver framework* for collecting data from and sending control actions to buildings and devices.
- automatic data capture and retrieval through our *historian framework*.
- platform based *agent lifecycle management*.
- a *web based management* tool for managing several instances from a central instance.
- the ability to easily extend the functionality of existing agents or create new ones for your specific purposes.

CHAPTER 2

Background

VOLTTRON™ is written in Python 2.7 and runs on Linux Operating Systems. For users unfamiliar with those technologies, the following resources are recommended:

- <https://docs.python.org/2.7/tutorial/>
- <http://ryanstutorials.net/linuxtutorial/>

The project is *licensed* under Apache 2 license.

Contents:

3.1 Overview

VOLTTRON™ is an open-source distributed control and sensing platform for integrating buildings and the power grid. VOLTTRON connects devices, agents in the platform, agents in the Cloud, and signals from the power grid. The platform also supports use cases such as demand response and integration of distributed renewable energy sources.

VOLTTRON provides an environment for agent execution and serves as a single point of contact for interfacing with devices (rooftop units, building systems, meters, etc.), external resources, and platform services such as data archival and retrieval. VOLTTRON applications are referred to as agents since VOLTTRON provides an agent-based programming paradigm to ease application development and minimize the lines of code that need to be written by domain experts such as buildings engineers. VOLTTRON provides a collection of utility and helper classes that simplifies agent development.

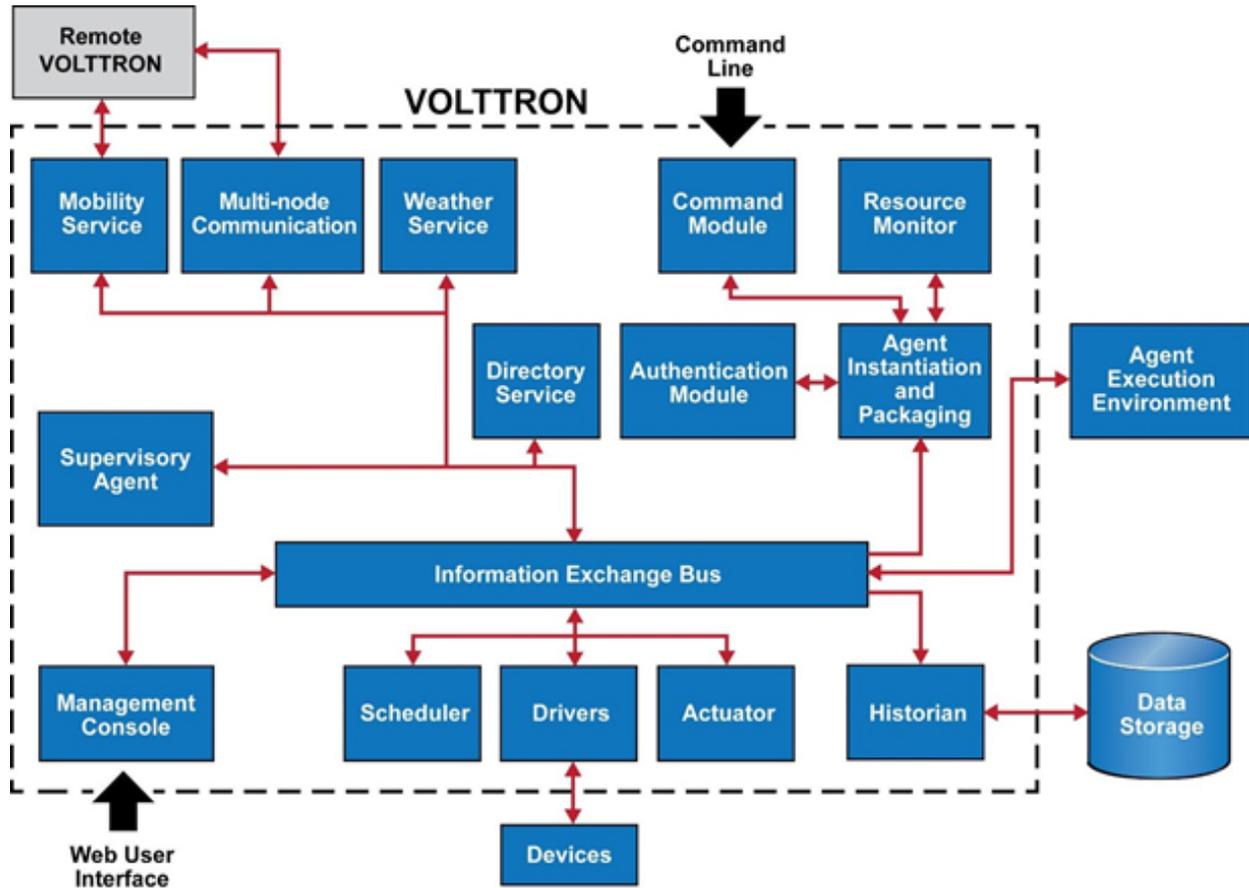
The VOLTTRON white paper provides an overview of the capabilities of the platform: https://volttron.org/sites/default/files/publications/PNNL-25499_VOLTTRON_2016.pdf

3.1.1 Components

An overview of the VOLTTRON platform components is illustrated in the figure below. The platform comprises several components and agents that provide services to other agents. Of these components, the Information Exchange Bus (IEB), or Message Bus is central to the platform. All other VOLTTRON components communicate through it using the publish/subscribe paradigm over a variety of topics.

Drivers communicate with devices allowing their data to be published on the IEB. Agents can control devices by interacting with the *Actuator Agent* to schedule and send commands. The *Historian* framework takes data published on the messages bus and stores it to a database, file, or sends it to another location.

The agent lifecycle is controlled by the Agent Instantiation and Packaging (AIP) component which launches agents in an Agent Execution Environment. This isolates agents from the platform while allowing them to interact with the IEB.



3.1.2 Agents in the Platform

Agents deployed on VOLTTRON can perform one or more roles which can be broadly classified into the following groups:

- **Platform Agents:** Agents which are part of the platform and provide a service to other agents. Examples are agents which interface with devices to publish readings and handle control signals from other agents.
- **Cloud Agents:** These agents represent a remote application which needs access to the messages and data on the platform. This agent would subscribe to topics of interest to the remote application and would also allow it publish data to the platform.
- **Control Agents:** These agents control the devices of interest and interact with other resources to achieve some goal.

Platform Services:

- **Message Bus:** All agents and services publish and subscribe to topics on the message bus. This provides a single interface that abstracts the details of devices and agents from each other. Components in the platform basically produce and consume events.
- **Weather Information:** This agent periodically retrieves data from the Weather Underground site. It then reformats it and publishes it out to the platform on a weather topic.

- **Modbus-based device interface:** The Modbus driver publishes device data onto the message bus. It also handles the locking of devices to prevent multiple conflicting directives.
- **Application Scheduling:** This service allows the scheduling of agents' access to devices in order to prevent conflicts.
- **Logging service:** Agents can publish arbitrary strings to a logging topic and this service will push them to a historian for later analysis.

3.1.3 Definition of Terms

This page lays out a common terminology for discussing the components and underlying technologies used by the platform. The first section discusses capabilities and industry standards that volttron conforms to while the latter is specific to the VOLTRON domain.

Industry Terms

- **BACNet:** Building Automation and Control network, that leverages ASHRAE, ANSI, and IOS 16484-5 standard protocols.
- **JSON-RPC:** JSON-encoded remote procedure call
- **JSON:** JavaScript object notation is a text-based, human-readable, open data interchange format, similar to XML, but less verbose
- **Publish/subscribe:** A message delivery pattern where senders (publishers) and receivers (subscribers) do not communicate directly nor necessarily have knowledge of each other, but instead exchange messages through an intermediary based on a mutual class or topic
- **ZeroMQ or ØMQ:** A library used for inter-process and inter-computer communication
- **Modbus:** Communications protocol for talking with industrial electronic devices
- **SSH:** Secure shell is a network protocol providing encryption and authentication of data using public-key cryptography
- **SSL:** Secure sockets layer is a technology for encryption and authentication of network traffic based on a chain of trust
- **TLS:** Transport layer security is the successor to SSL

VOLTRON Terms

Activated Environment An activated environment is the environment a VOLTRON instance is run in. The bootstrap process creates the environment from the shell and to activate it the following command is executed.

```
user@computer> source env/bin/activate

# Note once the above command has been run the prompt will have changed
(volttron)user@computer>
```

Bootstrap Environment The process by which an operating environment (activated environment) is produced. From the VOLTRON_ROOT directory executing `python bootstrap.py` will start the bootstrap process.

VOLTTRON_HOME The location for a specific VOLTTRON_INSTANCE to store its specific information. There can be many VOLTTRON_HOMEs on a single computing resource(VM, machine, etc.)

VOLTTRON_INSTANCE A single volttron process executing instructions on a computing resource. For each VOLTTRON_INSTANCE there WILL BE only one VOLTTRON_HOME associated with it. In order for a VOLTTRON_INSTANCE to be able to participate outside its computing resource it must be bound to an external ip address.

VOLTTRON_ROOT The cloned directory from github. When executing the command

```
git clone http://github.com/VOLTTRON/volttron
```

the top volttron folder is the VOLTTRON_ROOT

VIP VOLTTRON Interconnect Protocol is a secure routing protocol that facilitates communications between agents, controllers, services and the supervisory VOLTTRON_INSTANCE.

3.1.4 Version History

VOLTTRON 1.0 – 1.2

- Agent execution platform
- Message bus
- Modbus and BACnet drivers
- Historian
- Data logger
- Device scheduling
- Device actuation
- Multi-node communication
- Weather service

VOLTTRON 2.0

- Advanced Security Features
- Guaranteed resource allocation to agents using execution contracts
- Signing and verification of agent packaging
- Agent mobility
- Admin can send agents to another platform
- Agent can request to move
- Enhanced command framework

VOLTTRON 3.0

- Modularize Data Historian
- Modularize Device Drivers
- Secure and accountable communication using the VIP
- Web Console for Monitoring and Administering VOLTTRON Deployments

VOLTRON 4.0

- Documentation moved to ReadTheDocs
- VOLTRON Configuration Wizard
- Configuration store to dynamically configure agents
- Aggregator agent for aggregating topics
- More reliable remote install mechanism
- UI for device configuration
- Automatic registration of VOLTRON instances with management agent

VOLTRON 5.0

- Tagging service for attaching metadata to topics for simpler retrieval
- Message bus performance improvement
- Multi-platform publish/subscribe for simpler coordination across platforms
- Drivers contributed back for SEP 2.0 and ChargePoint EV

3.1.5 License

Copyright 2017, Battelle Memorial Institute.

Licensed under the Apache License, Version 2.0 (the “License”); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

The patent license grant shall only be applicable to the following patent and patent application (Battelle IPID 17008-E), as assigned to the Battelle Memorial Institute, as used in conjunction with this Work: • US Patent No. 9,094,385, issued 7/28/15 • USPTO Patent App. No. 14/746,577, filed 6/22/15, published as US 2016-0006569.

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an “AS IS” BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

3.1.6 Terms

This material was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the United States Department of Energy, nor Battelle, nor any of their employees, nor any jurisdiction or organization that has cooperated in the development of these materials, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness or any information, apparatus, product, software, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or Battelle Memorial Institute. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

PACIFIC NORTHWEST NATIONAL LABORATORY operated by BATTELLE for the UNITED STATES DEPARTMENT OF ENERGY under Contract DE-AC05-76RL01830

3.2 Join the Community

The VOLTTRON team aims to work with users and contributors to continuously improve the platform with features requested by the community as well as architectural features that improve robustness, security, and scalability. Contributing back to the project, which is encouraged but not required, enhances its capabilities for the whole community. To learn more, check out *Contributing* and *Documentation*.

3.2.1 Slack Channel

voltron-community.slack.com is where the VOLTTRON™ community at large can ask questions and meet with others using VOLTTRON™ Signup via <https://voltron-community.signup.team/>

3.2.2 Mailing List

Join the mailing list by emailing voltron@pnnl.gov.

3.2.3 Stack Overflow

The VOLTTRON community supports questions being asked and answered through stack overflow. The questions tagged with the voltron tag can be found at <http://stackoverflow.com/questions/tagged/voltron>.

3.2.4 Office Hours

PNNL hosts office hours every other week on Fridays at 11 AM (PST). These meetings are designed to be very informal where VOLTTRON developers can answer specific questions about the inner workings of VOLTTRON. These meetings are also available for topical discussions of different aspects of the VOLTTRON platform. Currently the office hours are available through a Lync meeting.

- [Join Lync Meeting](#)
- Dial in Number: 1-866-528-1882
- Conference Number: 886916

Meetings are recorded and can be reviewed [here](#).

3.2.5 Contributing Back

Contributing to VOLTTRON

As an open source project VOLTTRON requires input from the community to keep development focused on new and useful features. To that end we are revising our commit process to hopefully allow more committers to be part of the community. The following document outlines the process for source code and documentation to be submitted. There are gui tools that may make this process easier, however I am going to focus on what is required from the command line.

The only requirements for this is the program git and your favorite web browser.

Getting Started

Forking the main VOLTRON repository

The first step to editing the repository is to fork it into your own user space. This is done by pointing your favorite web browser to <http://github.com/VOLTRON/volttron> and the clicking Fork on the upper right of the screen. (Note you must have a github account to fork the repository, if you don't have one then click Sign Up in the upper right corner and create one).

Cloning 'YOUR' VOLTRON forked repository

The next step in the process is to get your forked repository down to your computer to work on. This will create an identical copy of the github on your local machine. To do this you need to know the address of your repository. Luckily the github has a convention so your repository address will be <https://github.com/<YOUR USERNAME>/volttron.git>. From a terminal execute the following commands which will create a directory in your home directory and then change to that directory, clone from your repository, and finally change into the cloned repository.

Note: VOLTRON uses develop as its main development branch rather than the standard master branch (the default).

```
mkdir -p ~/git
cd ~/git
git clone -b develop https://github.com/<YOUR USERNAME>/volttron.git
cd volttron
```

Adding and Committing files

Now that you have your repository cloned it's time to start doing some modifications. Using a simple text editor you can create or modify any file under the volttron directory. After making a modification or an creating a file it is time to move it to the stage for review before committing to the local repository. For this example let's assume we have made a change to README.md in the root of the volttron directory and add a new file called foo.py. To get those files on the stage (preparing for committing to the local repository) we would execute the following commands

```
git add foo.py
git add README.md

# Alternatively in one command
git add foo.py README.md
```

After adding the files to the stage you can review the staged files by executing

```
git status
```

Finally in order to commit to the local repository we need to think of what change we actually did and be able to document it. We do that with a commit message (the -m parameter) such as the following.

```
git commit -m "Added new foo.py and updated copyright of README.md"
```

Pushing to the remote repository

Next we want to be able to share our changes with the world through github. We can do this by pushing the commits from your local repository out to your github repository. This is done by the following command.

```
git push
# alternative where origin is the name of the remote you are pushing to
# more on that later.
git push origin
```

Getting modifications to the main VOLTTRON repository

Now we want our changes to get put into the main VOLTTRON repository. After all our foo.py can cure a lot of the world's problems and of course it is always good to have a copyright the correct year. Open your browser to <https://github.com/VOLTTRON/volttron/compare/develop...<YOUR USERNAME>:develop>.

On that page the base fork should always be VOLTTRON/volttron with the base develop whilst the head fork should be <YOUR USERNAME>/volttron and the compare should be the branch in your repository to pull from. Once you have verified that you have got the right changes made then you can enter a title and description that represent your changes.

What happens next?

Once creating a pull request one or more VOLTTRON team members will review your changes and either accept them as is or ask for modifications in order to have your commits accepted. You will be automatically emailed through the github notification system when this occurs.

Next Steps

Merging changes from the main VOLTTRON repository

As time goes on the VOLTTRON code base will continually be modified so the next time you want to work on a change to your files the odds are your local and remote repository will be out of date. In order to get your remote VOLTTRON repository up to date with the main VOLTTRON repository you could simply do a pull request to your remote repository from the main repository. That would involve pointing your browser at <https://github.com/<YOUR USERNAME>/volttron/compare/develop... VOLTTRON:develop>.

Click the 'Create Pull Request' button. On the following page click the 'Create Pull Request' button. On the next page click 'Merge Pull Request' button.

Once your remote is updated you can now pull from your remote repository into your local repository through the following command:

```
git pull
```

The other way to get the changes into your remote repository is to first update your local repository with the changes from the main VOLTTRON repository and then pushing those changes up to your remote repository. To do that you need to first create a second remote entry to go along with the origin. A remote is simply a pointer to the url of a different repository than the current one. Type the following command to create a new remote called 'upstream'

```
git remote add upstream https://github.com/VOLTTRON/volttron
```

To update your local repository from the main VOLTRON repository then execute the following command where upstream is the remote and develop is the branch to pull from.

```
git pull upstream develop
```

Finally to get the changes into your remote repository you can execute

```
git push origin
```

Other commands to know

At this point in time you should have enough information to be able to update both your local and remote repository and create pull requests in order to get your changes into the main VOLTRON repository. The following commands are other commands to give you more information that the preceding tutorial went through

Viewing what the remotes are in our local repository

```
git remote -v
```

Stashing changed files so that you can do a merge/pull from a remote

```
git stash save 'A comment to be listed'
```

Applying the last stashed files to the current repository

```
git stash pop
```

Finding help about any git command

```
git help  
git help branch  
git help stash  
git help push  
git help merge
```

Creating a branch from the branch and checking it out

```
git checkout -b newbranchname
```

Checking out a branch (if not local already will look to the remote to checkout)

```
git checkout branchname
```

Removing a local branch (cannot be current branch)

```
git branch -D branchname
```

Determine the current and show all local branches

```
git branch
```

Hooking into other services

The main VOLTTRON repository is hooked into an automated build tool called travis-ci. Your remote repository can be automatically built with the same tool by hooking your account into travis-ci's environment. To do this go to <https://travis-ci.org> and create an account. You can use your github login directly to this service. Then you will need to enable the syncing of your repository through the travis-ci service. Finally you need to push a new change to the repository. If the build fails you will receive an email notifying you of that fact and allowing you to modify the source code and then pushing new changes out.

Contributing Documentation

The Community is encouraged to contribute documentation back to the project as they work through use cases the developers may not have considered or documented. By contributing documentation back, the community can learn from each other and build up a much more extensive knowledge base.

VOLTTRON™ documentation utilizes ReadTheDocs: <http://volttron.readthedocs.io/en/develop/> and is built using the Sphinx Python library with static content in Restructured Text.

Building the Documentation

Static documentation can be found in the *docs/source* directory. Edit or create new .rst files to add new content using the Restructured Text format. To see the results of your changes, the documentation can be built locally through the command line using the following instructions.

If you've already *bootstrapped* VOLTTRON™, do the following while activated. If not, this will also pull down the necessary VOLTTRON™ libraries.

```
python bootstrap.py --documentation
cd docs
make html
```

Then, open your browser to the created local files:

```
file:///home/<USER>/git/volttron/docs/build/html/overview/index.html
```

When complete, changes can be contributed back using the same process as code *contributions* by creating a pull request. When the changes are accepted and merged, they will be reflected in the ReadTheDocs site.

3.3 Installing VOLTRON

3.3.1 Install Required Software

Ensure that all the *required packages* are installed.

3.3.2 Clone VOLTRON source code

From version 6.0 VOLTRON supports two message bus - ZMQ and RabbitMQ. For the latest build use the develop branch. For a more conservative branch please use the master branch.

```
git clone https://github.com/VOLTRON/volttron --branch <branch name>
```

For other options see: *Getting VOLTRON*

3.3.3 Setup virtual environment

The VOLTRON project includes a bootstrap script which automatically downloads dependencies and builds VOLTRON. The script also creates a Python virtual environment for use by the project which can be activated after bootstrapping with `. env/bin/activate`. This activated Python virtual environment should be used for subsequent bootstraps whenever there are significant changes. The system's Python need only be used on the initial bootstrap.

Steps for ZMQ

```
cd <volttron clone directory>
python bootstrap.py
source env/bin/activate
```

Proceed to *Testing the Installation*.

Steps for RabbitMQ

1. Install Erlang version \geq 21

For RabbitMQ based VOLTRON, some of the RabbitMQ specific software packages have to be installed. If you are running an **Debian or CentOS system**, you can install the RabbitMQ dependencies by running the rabbit dependencies script, passing in the os name and appropriate distribution as a parameter. The following are supported

- debian bionic (for Ubuntu 18.04)
- debian xenial (for Ubuntu 16.04)
- debian xenial (for Linux Mint 18.04)
- debian stretch (for Debian Stretch)
- centos 7 (for CentOS 7)
- centos 6 (for CentOS 6)

Example command

```
./scripts/rabbit_dependencies.sh debian xenial
```

Alternatively

You can download and install Erlang from [Erlang Solution](#) Please include OTP/components - ssl, public_key, asn1, and crypto. Also lock version of Erlang using the [yum-plugin-versionlock](#)

2. Configure hostname

Rabbitmq requires a valid hostname to start. Use the command `hostname` on your linux machine to verify if a valid hostname is set. If not add a valid hostname to the file `/etc/hostname`. You would need `sudo` access to edit this file If you want your rabbitmq instance to be reachable externally, then a hostname should be resolvable to a valid ip. In order to do this you need to have a entry in `/etc/hosts` file. For example, the below shows a valid `/etc/hosts` file

```
127.0.0.1 localhost
127.0.0.1 myhost

192.34.44.101 externally_visible_hostname
```

After the edit, logout and log back in for the changes to take effect.

If you are testing with VMs make please make sure to provide unique host names for each of the VM you are using.

Note: If you change `/etc/hostname` after setting up rabbitmq (<refer to the step that does `vcfg --rabbitmq single`>), you will have to regenerate certificates and restart RabbitMQ.

Note: RabbitMQ startup error would show up in system log (`/var/log/messages`) file and not in RabbitMQ logs (`$RABBITMQ_HOME/var/log/rabbitmq/rabbitmq@hostname.log` where `$RABBITMQ_HOME` is `<install dir>/rabbitmq_server-3.7.7`)

3. Bootstrap

Install the required software by running the bootstrap script with `--rabbitmq` option

```
cd volttron

# bootstrap.py --help will show you all of the "package options" such as
# installing required packages for volttron central or the platform agent.

python bootstrap.py --rabbitmq [optional install directory defaults to
<user_home>/rabbitmq_server]
```

This will build the platform and create a virtual Python environment and dependencies for RabbitMQ. It also installs RabbitMQ server as the current user. If an install path is provided, path should exists and be writeable. RabbitMQ will be installed under `<install dir>/rabbitmq_server-3.7.7` Rest of the documentation refers to the directory `<install dir>/rabbitmq_server-3.7.7` as `$RABBITMQ_HOME`

You can check if RabbitMQ server is installed by checking it's status.

```
$RABBITMQ_HOME/sbin/rabbitmqctl status
```

Please note, RABBITMQ_HOME environment variable can be set in ~/.bashrc. If doing so, it needs to be set to RabbitMQ installation directory (default path is <user_home>/rabbitmq_server/rabbitmq_server-3.7.7)

```
echo 'export RABBITMQ_HOME=$HOME/rabbitmq_server/rabbitmq_server-3.7.7'|tee -
→-append ~/.bashrc | source ~/.bashrc
# Reload the environment variables in the current shell
source ~/.bashrc
```

4. Activate the environment

```
source env/bin/activate
```

5. Create RabbitMQ setup for VOLTRON

```
vcfg --rabbitmq single [optional path to rabbitmq_config.yml]
```

Refer to examples/configurations/rabbitmq/rabbitmq_config.yml for a sample configuration file. At a minimum you would need to provide the host name and a unique common-name (under certificate-data) in the configuration file. Note. common-name must be unique and the general conventions is to use -root-ca.

Running the above command without the optional configuration file parameter will prompt user for all the needed data at the command prompt and use that to generate a rabbitmq_config.yml file in VOLTTRON_HOME directory.

This scripts creates a new virtual host and creates SSL certificates needed for this VOLTRON instance. These certificates get created under the sub directory “certificates” in your VOLTRON home (typically in ~/.volttron). It then creates the main VIP exchange named “volttron” to route message between platform and agents and alternate exchange to capture unroutable messages.

NOTE: We configure RabbitMQ instance for a single volttron_home and volttron_instance. This script will confirm with the user the volttron_home to be configured. volttron instance name will be read from volttron_home/config if available, if not user will be prompted for volttron instance name. To run the scripts without any prompts, save the volttron instance name in volttron_home/config file and pass the volttron home directory as command line argument For example: “vcfg -vhome /home/vdev/new_vhome -rabbitmq single”

Following is the example inputs for “vcfg -rabbitmq single” command. Since no config file is passed the script prompts for necessary details.

```
Your VOLTRON_HOME currently set to: /home/vdev/new_vhome2

Is this the volttron you are attempting to setup? [Y]:
Creating rmq config yml
RabbitMQ server home: [/home/vdev/rabbitmq_server/rabbitmq_server-3.7.7]:
Fully qualified domain name of the system: [cs_cbox.pnl.gov]:

Enable SSL Authentication: [Y]:

Please enter the following details for root CA certificates
```

(continues on next page)

(continued from previous page)

```

Country: [US]:
State: Washington
Location: Richland
Organization: PNNL
Organization Unit: Volttron-Team
Common Name: [volttron1-root-ca]:
Do you want to use default values for RabbitMQ home, ports, and virtual_
↳host: [Y]: N
Name of the virtual host under which RabbitMQ VOLTTRON will be running:
↳[volttron]:
AMQP port for RabbitMQ: [5672]:
http port for the RabbitMQ management plugin: [15672]:
AMQPS (SSL) port RabbitMQ address: [5671]:
https port for the RabbitMQ management plugin: [15671]:
INFO:rmq_setup.pyc:Starting rabbitmq server
Warning: PID file not written; -detached was passed.
INFO:rmq_setup.pyc:**Started rmq server at /home/vdev/rabbitmq_server/
↳rabbitmq_server-3.7.7
INFO:requests.packages.urllib3.connectionpool:Starting new HTTP connection_
↳(1): localhost
INFO:requests.packages.urllib3.connectionpool:Starting new HTTP connection_
↳(1): localhost
INFO:requests.packages.urllib3.connectionpool:Starting new HTTP connection_
↳(1): localhost
INFO:rmq_setup.pyc:
Checking for CA certificate

INFO:rmq_setup.pyc:
Root CA (/home/vdev/new_vhome2/certificates/certs/volttron1-root-ca.crt) NOT_
↳Found. Creating root ca for volttron instance
Created CA cert
INFO:requests.packages.urllib3.connectionpool:Starting new HTTP connection_
↳(1): localhost
INFO:requests.packages.urllib3.connectionpool:Starting new HTTP connection_
↳(1): localhost
INFO:rmq_setup.pyc:**Stopped rmq server
Warning: PID file not written; -detached was passed.
INFO:rmq_setup.pyc:**Started rmq server at /home/vdev/rabbitmq_server/
↳rabbitmq_server-3.7.7
INFO:rmq_setup.pyc:

#####

Setup complete for volttron home /home/vdev/new_vhome2 with instance_
↳name=volttron1
Notes:
- Please set environment variable VOLTTRON_HOME to /home/vdev/new_vhome2_
↳before starting volttron
- On production environments, restrict write access to
/home/vdev/new_vhome2/certificates/certs/volttron1-root-ca.crt to only admin_
↳user. For example: sudo chown root /home/vdev/new_vhome2/certificates/
↳certs/volttron1-root-ca.crt
- A new admin user was created with user name: volttron1-admin and_
↳password=default_passwd.
You could change this user's password by logging into https://cs_cbox.pnl.
↳gov:15671/ Please update /home/vdev/new_vhome2/rabbitmq_config.yml if you_
↳change password

```

(continues on next page)

(continued from previous page)

#####

3.3.4 Testing the Installation

We are now ready to start VOLTRON instance. If configured with RabbitMQ message bus a config file would have been generated in \$VOLTRON_HOME/config with the entry message-bus=rmq. If you need to revert back to ZeroMQ based VOLTRON, you will have to either remove “message-bus” parameter or set it to default “zmq” in \$VOLTRON_HOME/config. The following command starts volttron process in the background

```
volttron -vv -l volttron.log&
```

This enters the virtual Python environment and then starts the platform in debug (vv) mode with a log file named volttron.log. Alternatively you can use the utility script start-volttron script that does the same. To stop stop volttron you can use the stop-volttron script.

```
./start-volttron
```

Warning: If you plan on running VOLTRON in the background and detaching it from the terminal with the disown command be sure to redirect stderr and stdout to /dev/null. Some libraries which VOLTRON relies on output directly to stdout and stderr. This will cause problems if those file descriptors are not redirected to /dev/null

```
#To start the platform in the background and redirect stderr and stdout
#to /dev/null
volttron -vv -l volttron.log > /dev/null 2>&1&
```

3.3.5 Installing and Running Agents

VOLTRON platform comes with several built in services and example agents out of the box. To install a agent use the script install-agent.py

```
python scripts/install-agent.py -s <top most folder of the agent> [-c <config file.>
↳ Might be optional for some agents>]
```

For example, we can use the command to install and start the Listener Agent - a simple agent that periodically publishes heartbeat message and listens to everything on the message bus. Install and start the Listener agent using the following command.

```
python scripts/install-agent.py -s examples/ListenerAgent --start
```

Check volttron.log to ensure that the listener agent is publishing heartbeat messages.

```
tail volttron.log
```

```
2016-10-17 18:17:52,245 (listeneragent-3.2 11367) listener.agent INFO: Peer: 'pubsub',
↳ Sender: 'listeneragent-3.2_1':, Bus: u'', Topic: 'heartbeat/listeneragent-3.2_1',
↳ Headers: {'Date': '2016-10-18T01:17:52.239724+00:00', 'max_compatible_version': u'',
↳ 'min_compatible_version': '3.0'}, Message: {'status': 'GOOD', 'last_updated':
↳ '2016-10-18T01:17:47.232972+00:00', 'context': 'hello'}
```

You can also use the `vctl` or `volttron-ctl` command to start, stop or check the status of an agent

```
(volttron)volttron@volttron1:~/git/rmq_volttron$ vctl status
AGENT          IDENTITY          TAG          STATUS          HEALTH
6 listeneragent-3.2 listeneragent-3.2_1 running [13125] GOOD
f master_driveragent-3.2 platform.driver master_driver
```

```
vctl stop <agent id>
```

To stop the platform:

```
volttron-ctl shutdown --platform
```

or

```
./stop-volttron
```

Note: The default working directory is `~/volttron`. The default directory for creation of agent packages is `~/volttron/packaged`

3.3.6 Next Steps

Now that the project is configured correctly:

See the following links for core services and volttron features:

- [Core Services](#)
- [Platform Specifications](#)

See the following links for agent development:

- [Agent Development](#)
- [VOLTTRON Development in Eclipse](#)
- [VOLTTRON Development in PyCharm](#)

Please refer to related topics to for advanced setup instructions

3.3.7 Related Topics

Using RabbitMQ message bus

Backward Compatibility With ZeroMQ Message Based VOLTTRON

RabbitMQ VOLTTRON supports backward compatibility with ZeroMQ based VOLTTRON. RabbitMQ VOLTTRON has a ZeroMQ router running internally to accept incoming ZeroMQ connections and to route ZeroMQ messages coming in/going out of it's instance. There are multiple ways for an instance with a RabbitMQ message bus, and an instance with ZeroMQ message bus to connect with each other. For example, an agent from one instance can directly connect to the remote instance to publish or pull data from it. Another way is through multi-platform communication, where the VOLTTRON platform is responsible for connecting to the remote instance. For more information on multi-platform communication, see https://volttron.readthedocs.io/en/develop/core_services/multiplatform/Multiplatform-Communication.html.

Agent Connecting Directly to Remote Instance

The following steps are to demonstrate how RabbitMQ VOLTRON is backward compatible with ZeroMQ VOLTRON, using the Forward Historian as an example. This example shows how to forward messages from local ZeroMQ based VOLTRON to remote RabbitMQ based VOLTRON instance. Similar steps can be followed if you needed to move messages from local RabbitMQ based VOLTRON to ZeroMQ based VOLTRON.

1. In order for RabbitMQ and ZeroMQ VOLTRONs to communicate with each other, one needs two instances of VOLTRON_HOME on the same VM. To create a new instance of VOLTRON_HOME use the command.

```
export VOLTRON_HOME=~/.new_volttron_home
```

It is recommended that one uses multiple terminals to keep track of both instances.

2. Start VOLTRON on both instances. Note: since the start-volttron script uses the volttron.log by default, the second instance will need be started manually in the background, using a separate log. For example:

```
volttron -vv -l volttron-two.log&
```

3. Modify the configuration file for both instances. The config file is located at \$VOLTRON_HOME/config

For RabbitMQ VOLTRON, the config file should look similar to:

```
[volttron]
message-bus = rmq
vip-address = tcp://127.0.0.1:22916
instance-name = volttron_rmq
```

The ZeroMQ config file should look similar, with all references to RMQ being replaced with ZMQ, and a different vip-address (e.g. tcp://127.0.0.2:22916).

4. On the instance running ZeroMQ:
 - a. Install the Forward Historian agent using an upgrade script similar to:

```
#!/bin/bash
export CONFIG=$(mktemp /tmp/abc-script.XXXXXX)
cat > $CONFIG <<EOL
{
    "destination-vip": "tcp://127.0.0.1:22916",
    "destination-serverkey": "key"
}
EOL
python /home/username/volttron/scripts/install-agent.py -c $CONFIG -s_
↪services/core/ForwardHistorian --start --force -i forward.historian
# Finally remove the temporary config file
rm $CONFIG
```

Since we are attempting to push data from the local (ZeroMQ in this example) to the remote (RabbitMQ) instance, the we would need the RabbitMQ serverkey, which can be obtained by running `vctl auth serverkey` on the RabbitMQ instance.

Start the Forward Historian.

- b. Install master driver, configure fake device on upstream server and start volttron and master driver. `vcfg --agent master_driver` command can install master driver and setup a fake device.

```
./stop-volttron
vcfg --agent master_driver
./start-volttron
vctl start --tag master_driver
```

5. On the instance running RabbitMQ:
 - a. Run a listener agent which subscribes to messages from all platforms

- Open the file `examples/ListenerAgent/listener/agent.py`. Search for `@PubSub.subscribe('pubsub', '')` and replace that line with `@PubSub.subscribe('pubsub', 'devices', all_platforms=True)`
- upgrade the listener

```
scripts/core/upgrade-listener
```

- b. Provide the RabbitMQ instance with the public key of the Forward Historian running on ZeroMQ instance.

Run `vctl auth public key` on the ZeroMQ instance. Copy the output and add the public key to the RabbitMQ instance's `auth.config` file, using the defaults except for the `user_id` and credentials.

```
vctl auth add
domain []:
address []:
user_id []: forward
capabilities (delimit multiple entries with comma) []:
roles (delimit multiple entries with comma) []:
groups (delimit multiple entries with comma) []:
mechanism [CURVE]:
credentials []: key
comments []:
enabled [True]:
```

Once that is completed you should be able to see data similar to below in `the log of the voltrron instance running RabbitMQ:`

```
2018-12-31 14:48:10,043 (listeneragent-3.2 7175) listener.agent INFO:
↳Peer: pubsub, Sender: forward.historian:, Bus: , Topic: devices/fake-
↳campus/fake-building/fake-device/all, Headers: {'X-Forwarded': True,
↳'SynchronizedTimeStamp': '2018-12-31T19:48:10.000000+00:00', 'TimeStamp
↳': '2018-12-31T19:48:10.001966+00:00', 'max_compatible_version': 'u',
↳'min_compatible_version': '3.0', 'Date': '2018-12-31T19:48:10.
↳001966+00:00'}, Message:
[{'Heartbeat': True, 'PowerState': 0, 'ValveState': 0, 'temperature': 50.
↳0},
  {'Heartbeat': {'type': 'integer', 'tz': 'US/Pacific', 'units': 'On/Off'},
   'PowerState': {'type': 'integer', 'tz': 'US/Pacific', 'units': '1/0'},
   'ValveState': {'type': 'integer', 'tz': 'US/Pacific', 'units': '1/0'},
   'temperature': {'type': 'integer',
                   'tz': 'US/Pacific',
                   'units': 'Fahrenheit'}}]
```

Multi-Platform Connection

The below example demonstrates backward compatibility using multi-platform connection method.

1. Refer to steps 1 -3 in the previous section for configuring two VOLTRON instances (one with RabbitMQ and one with ZeroMQ). For step 3, the VOLTRON config files need to account for a web-server, which is necessary for multi-platform communication. As such, the config files should look similar to the following:

```
[volttron]
message-bus = rmq
vip-address = tcp://127.0.0.1:22916
instance-name = volttron_rmq
bind-web-address = http://127.0.0.1:8080
```

2. Create an external_address.json file in the VOLTRON_HOME directory for both instances, with the IP address and port of the remote instance(s) it will need to connect to. In this example, the RabbitMQ would have the address of the ZeroMQ instance, and vice versa. Below is an example for one instance:

```
[
  "http://127.0.0.2:8080"
]
```

3. On the instance running ZeroMQ:

- a. Install the DataMover agent using an upgrade script similar to:

```
#!/bin/bash
export CONFIG=$(mktemp /tmp/abc-script.XXXXXX)
cat > $CONFIG <<EOL
{
  "destination-vip": "tcp://127.0.0.1:22916",
  "destination-serverkey": "rmq server key",
  "destination-instance-name": "volttron_rmq",
  "destination-message-bus": "zmq"
}
EOL
python /home/osboxes/volttron/scripts/install-agent.py -c $CONFIG -s_
↪services/core/DataMover --start --force -i data.mover
# Finally remove the temporary config file
rm $CONFIG
```

Replace “rmq server key” with the RabbitMQ server key.

In this example the DataMover will be running on the ZeroMQ instance, which means that the destination vip, serverkey, and instance name are configured to the RabbitMQ instance. However, destination-message-bus should be set to zmq. Start DataMover agent.

- b. Install master driver, configure fake device on upstream server and start volttron and master driver. ‘vcfg --agent master_driver’ command can install master driver and setup a fake device.

```
./stop-volttron
vcfg --agent master_driver
./start-volttron
vctl start --tag master_driver
```

4. On the instance running RabbitMQ:

- a. Start SQLHistorian. Easiest way to accomplish this is to stop VOLTRON, reconfigure to have RabbitMQ message bus and install platform.historian already installed, and start VOLTRON again.

```
./stop-volttron
vcfg --agent platform_historian
./start-volttron
vctl start --tag platform_historian
```

- b. Run a listener agent which subscribes to messages from all platforms

- Open the file `examples/ListenerAgent/listener/agent.py`. Search for `@PubSub.subscribe('pubsub', '')` and replace that line with `@PubSub.subscribe('pubsub', 'devices', all_platforms=True)`
- upgrade the listener

```
scripts/core/upgrade-listener
```

- Provide the RabbitMQ instance with the public key of the DataMover running on ZeroMQ instance.

Run `vctl auth public key` on the ZeroMQ instance. Copy the output and add the public key to the RabbitMQ instance's `auth.config` file, using the defaults except for the `user_id` and credentials.

```
vctl auth add
domain []:
address []:
user_id []: forward
capabilities (delimit multiple entries with comma) []:
roles (delimit multiple entries with comma) []:
groups (delimit multiple entries with comma) []:
mechanism [CURVE]:
credentials []: key
comments []:
enabled [True]:
```

- Stop VOLTRON on both instances, and restart using setup mode.

```
volttron -vv -l volttron.log --setup-mode&
```

If you tail the logs of both instances, there should be a message indicating that starting with setup mode was complete upon success.

After a successful start, a new file called `external_platform_discovery.json` should be located in the `$VOLTRON_HOME` directory of both instances. The file will contain the platform discovery information (`()`), of all external platforms the respective VOLTRON instance is aware of. The file will look something like:

```
{"<platform1 name>": {"vip-address": "tcp://<ip1>:<vip port1>",
  "instance-name": "<platform1 name>",
  "serverkey": "<serverkey1>"
},
"<platform2 name>": {"vip-address": "tcp://<ip2>:<vip port2>",
  "instance-name": "<platform2 name>",
  "serverkey": "<serverkey2>"
},
"<platform3 name>": {"vip-address": "tcp://<ip3>:<vip port3>",
  "instance-name": "<platform3 name>",
  "serverkey": "<serverkey3>"
},
.....
}
```

Additionally for different combinations of multi-bus, multi-platform setup, please refer to [Multi-Platform Multi-Bus Walk-through](#)

Monitoring and Controlling RabbitMQ

Some of the important native RabbitMQ control and management commands are now integrated with “volttron-ctl” utility. Using `volttron-ctl rabbitmq` management utility, we can control and monitor the status of RabbitMQ message bus.

```

volttron-ctl rabbitmq --help
usage: volttron-ctl command [OPTIONS] ... rabbitmq [-h] [-c FILE] [--debug]
                                                    [-t SECS]
                                                    [--msgdebug MSGDEBUG]
                                                    [--vip-address ZMQADDR]
                                                    ...

subcommands:

  add-vhost          add a new virtual host
  add-user           Add a new user. User will have admin privileges
                    i.e,configure, read and write
  add-exchange       add a new exchange
  add-queue          add a new queue
  list-vhosts        List virtual hosts
  list-users         List users
  list-user-properties
                    List users
  list-exchanges     add a new user
  list-exchange-properties
                    list exchanges with properties
  list-queues        list all queues
  list-queue-properties
                    list queues with properties
  list-bindings      list all bindings with exchange
  list-federation-parameters
                    list all federation parameters
  list-shovel-parameters
                    list all shovel parameters
  list-policies      list all policies
  remove-vhosts      Remove virtual host/s
  remove-users       Remove virtual user/s
  remove-exchanges   Remove exchange/s
  remove-queues      Remove queue/s
  remove-federation-parameters
                    Remove federation parameter
  remove-shovel-parameters
                    Remove shovel parameter
  remove-policies    Remove policy

```

Multi-Platform Multi-Bus Walk-through

This guide describes the setup process for a multi-platform connection that has a combination of ZeroMQ and RabbitMQ instances. For this example, we want to use the Forwarder to pass device data from two VOLTRON instance to a single “central” instance for storage. It will also have a Volttron Central agent running on the “central” instance and Volttron Central Platform agents on all 3 instances and connected to “central” instance to provide operational status of it’s instance to the “central” instance. For this document “node” will be used interchangeably with VOLTRON instance.

Node Setup

For this example we will have two types of nodes; a data collector and a central node. Each of the data collectors will have different message buses (VOLTTRON supports both RabbitMQ and ZeroMQ). The nodes will be configured as in the following table.

Table 1: Node Configuration

	Central	Node-ZMQ	Node-RMQ
Node Type	Central	Data Collector	Data Collector
Master Driver		yes	yes
Forwarder		yes	yes
SQL Historian	yes		
Volttron Central	yes		
Volttron Central Platform	yes	yes	yes
Exposes RMQ Port	yes		
Exposes ZMQ Port	yes		
Exposes HTTPS Port	yes		

The goal of this is to be able to see the data from Node-ZMQ and Node-RMQ in the Central SQL Historian and on the trending charts of Volttron Central.

Virtual Machine Setup

The first step in creating a VOLTTRON instance is to make sure the machine is ready for VOLTTRON. Each machine should have its hostname setup. For this walkthrough, the hostnames “central”, “node-zmq” and “node-rmq” will be used.

For Central and Node-RMQ follow the instructions Building-VOLTTRON#steps-for-rabbitmq. For Node-ZMQ use Building-VOLTTRON#steps-for-zmq.

Instance Setup

The following conventions/assumptions are made for the rest of this document:

- Commands should be run from the VOLTTRON root
- Default values are used for VOLTTRON_HOME(\$HOME/.volttron), vip port (22916), HTTPS port (8443), rabbitmq ports(5671 for AMQPs and 15671 for RabbitMQ management interface). If using different VOLTTRON_HOME or ports, please replace accordingly.
- Replace central, node-zmq and node-rmq with your own hostnames.

The following will use vcfg (volttron-cfg) to configure the individual platforms.

Central Instance Setup

Note: This instance must have been bootstrapped using `--rabbitmq` see Building-VOLTTRON#steps-for-rabbitmq.

Next step would be to configure the instance to have a web interface to accept/deny incoming certificate signing requests from other instances. Additionally, we will need to install a Volttron Central agent, Volttron Central Platform agent, SQL historian agent and a Listener agent. The following shows an example command output for this setup.

```
(volttron) d3x140@central:~/volttron$ vcfg

Your VOLTRON_HOME currently set to: /home/d3x140/.volttron

Is this the volttron you are attempting to setup? [Y]:
What type of message bus (rmq/zmq)? [zmq]: rmq
Name of this volttron instance: [volttron1]: central
RabbitMQ server home: [/home/d3x140/rabbitmq_server/rabbitmq_server-3.7.7]:
Fully qualified domain name of the system: [central]:
Would you like to create a new self signed root CA certificate for this instance: [Y]:

Please enter the following details for root CA certificate
Country: [US]:
State: WA
Location: Richland
Organization: PNNL
Organization Unit: volttron
Do you want to use default values for RabbitMQ home, ports, and virtual host: [Y]:
2019-06-20 16:28:11,341 volttron.utils.rmq_mgmt DEBUG: Creating new VIRTUAL HOST:
↳ volttron
2019-06-20 16:28:11,422 volttron.utils.rmq_mgmt DEBUG: Create READ, WRITE and
↳ CONFIGURE permissions for the user: central-admin
Create new exchange: volttron, {'durable': True, 'type': 'topic', 'arguments': {
↳ 'alternate-exchange': 'undeliverable'}}
Create new exchange: undeliverable, {'durable': True, 'type': 'fanout'}
Created CA cert

The rmq message bus has a backward compatibility
layer with current zmq instances. What is the
zmq bus's vip address? [tcp://127.0.0.1]: tcp://172.20.214.72
What is the port for the vip address? [22916]:
Is this instance web enabled? [N]: y
Web address set to: https://central
What is the port for this instance? [8443]:
Is this an instance of volttron central? [N]: y
Configuring /home/d3x140/volttron/services/core/VolttronCentral.
Enter volttron central admin user name: admin
Enter volttron central admin password:
Retype password:
Installing volttron central.
Should the agent autostart? [N]: y
Will this instance be controlled by volttron central? [Y]:
Configuring /home/d3x140/volttron/services/core/VolttronCentralPlatform.
What is the name of this instance? [central]:
What is the hostname for volttron central? [https://central]:
What is the port for volttron central? [8443]:
Should the agent autostart? [N]:
Would you like to install a platform historian? [N]: y
Configuring /home/d3x140/volttron/services/core/SQLHistorian.
Should the agent autostart? [N]: y
Would you like to install a master driver? [N]:
Would you like to install a listener agent? [N]: y
Configuring examples/ListenerAgent.
Should the agent autostart? [N]: y
Finished configuration!

You can now start the volttron instance.
```

(continues on next page)

(continued from previous page)

```
If you need to change the instance configuration you can edit
the config file is at /home/d3x140/.volttron/config
```

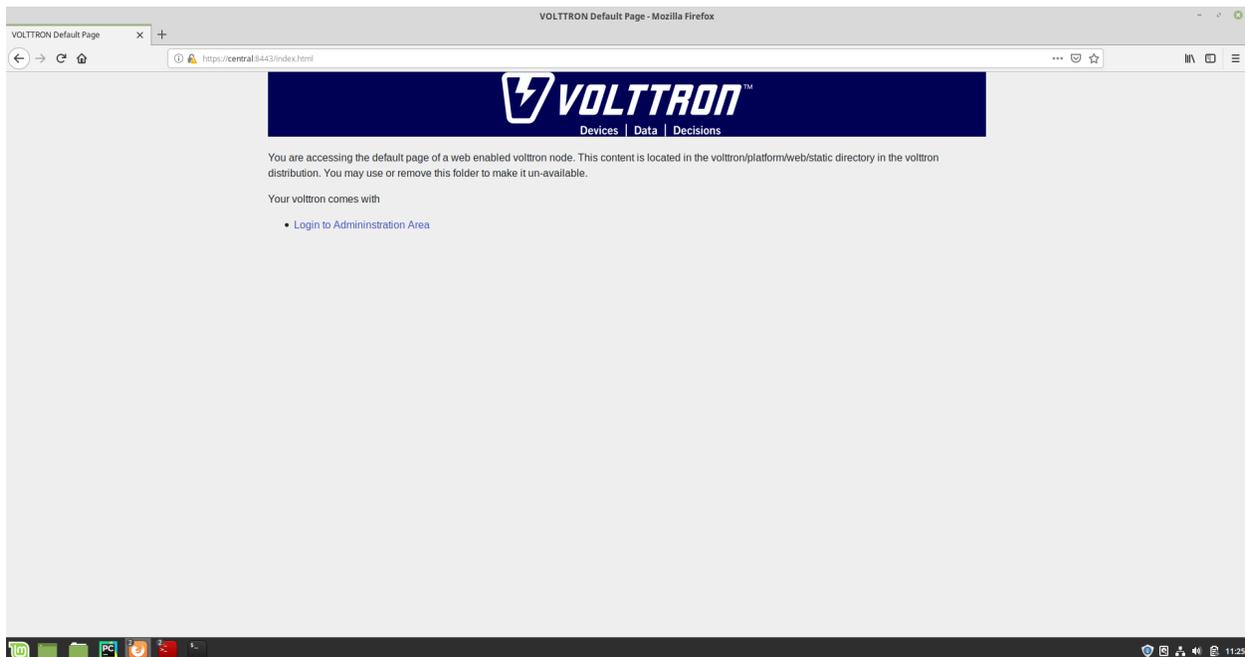
Start VOLTTRON instance and check if the agents are installed.

```
./start-volttron
vctl status
```

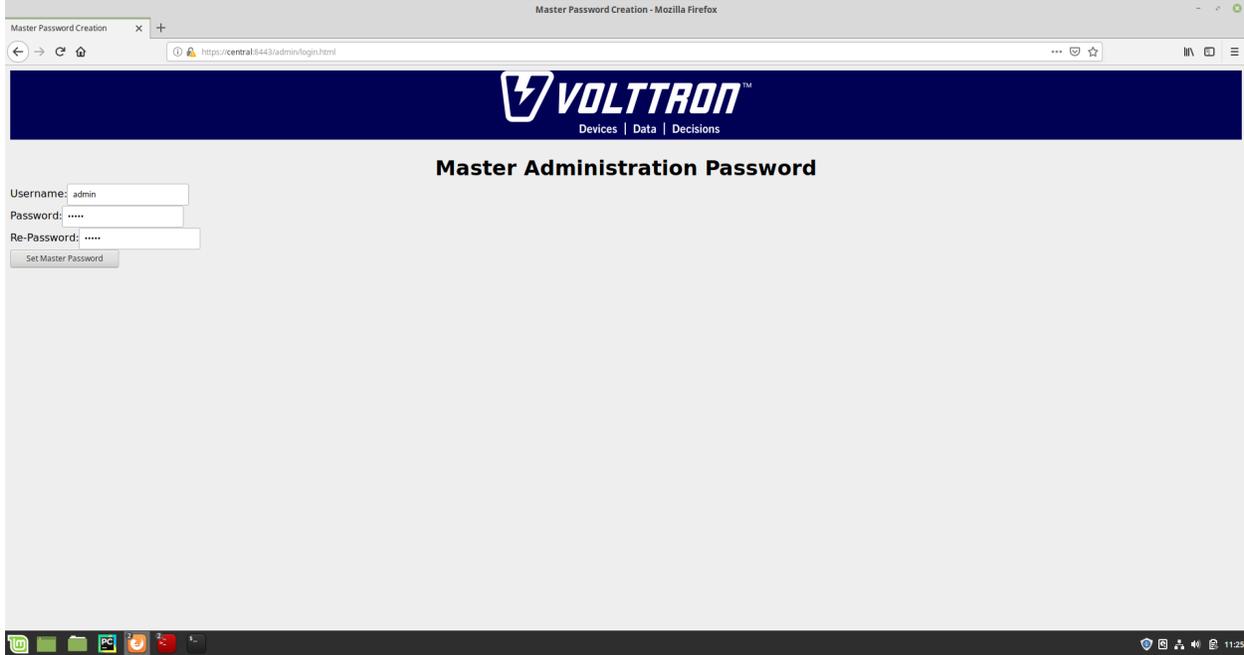
Open browser and go to master admin authentication page <https://central:8443/index.html> to accept/reject incoming certificate signing request (CSR) from other platforms.

Note: Replace “central” with the proper hostname of VC instance in the admin page URL. If opening the admin page from a different system, then please make that the hostname is resolvable in that machine.

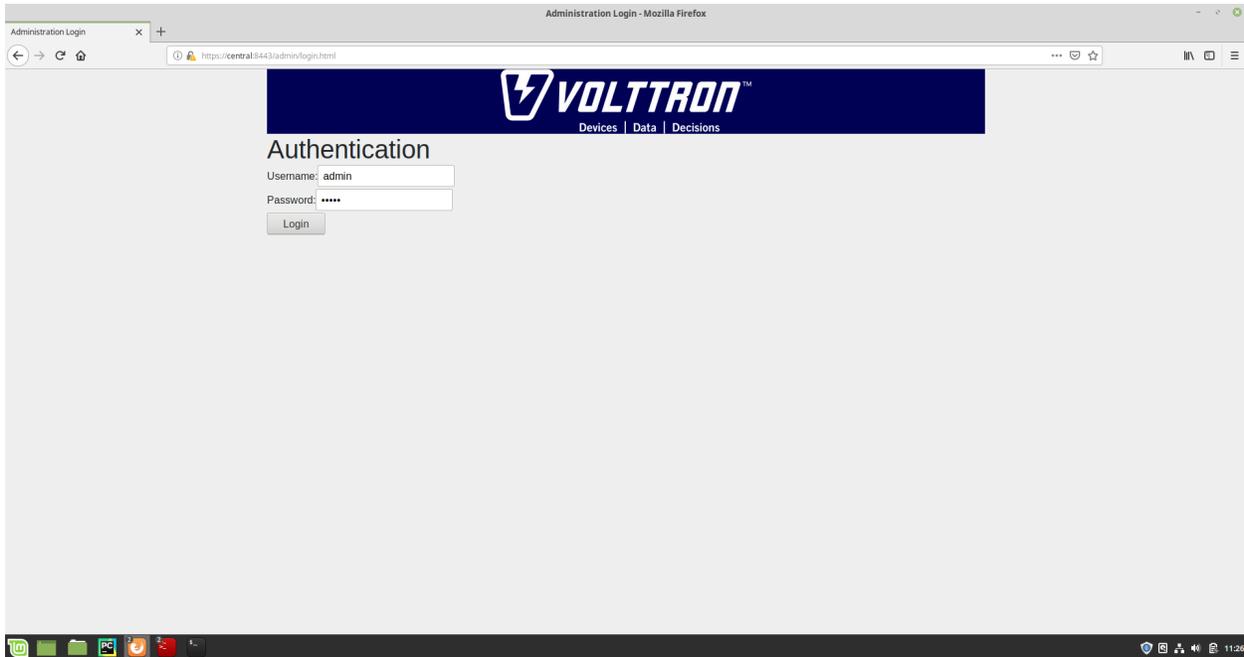
Click on “Login To Administration Area”.



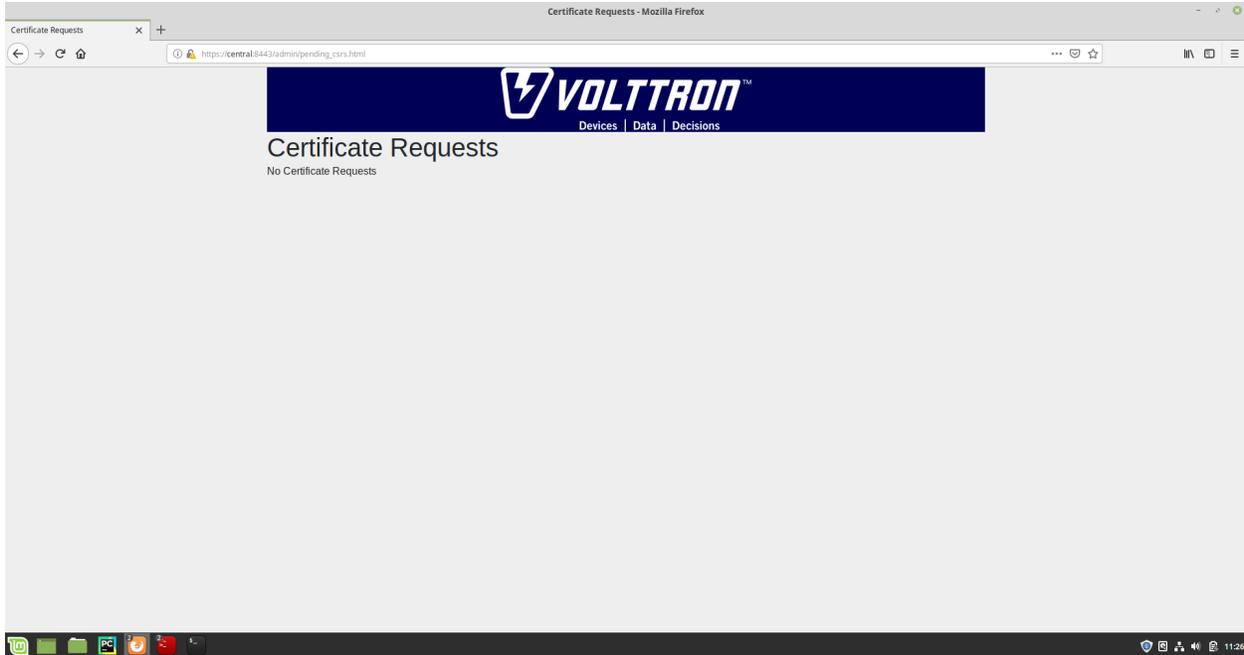
Set the master admin username and password. This can be later used to login into master admin authentication page.



Login into the Master Admin page.



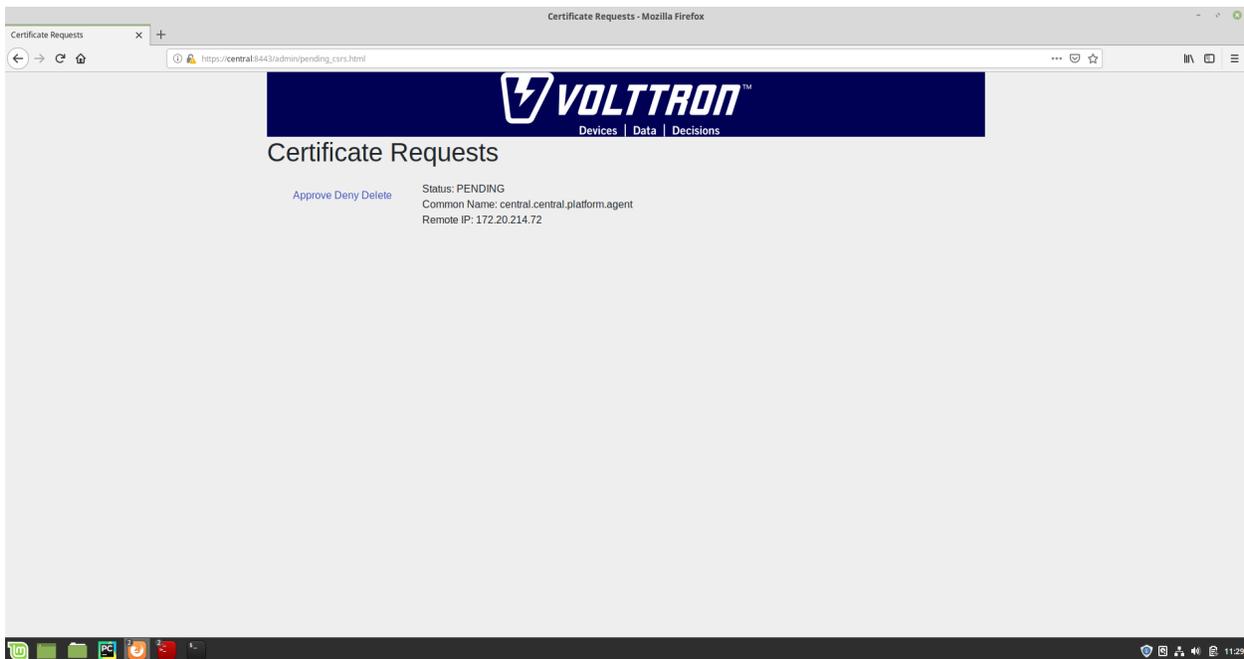
After logging in, you will see no CSR requests initially.



Go back to the terminal and start Voltron Central Platform agent on the “central” instance. The agent will send a CSR request to the web interface.

```
vctl start --tag vcp
```

Now go to master admin page to check if there is a new pending CSR request. You will see a “PENDING” request from “central.central.platform.agent”



Approve the CSR request to allow authenticated SSL based connection to the “central” instance.

Go back to the terminal and check the status of Voltron Central Platform agent. It should be set to “GOOD”.

Node-ZMQ Instance Setup

On the “node-zmq” VM, setup a ZeroMQ based VOLTRON instance. Using “vcfg” command, install Volttron Central Platform agent, a master driver agent with a fake driver.

Note: This instance will use old ZeroMQ based authentication mechanism using CURVE keys.

```
(volttron)d3x140@node-zmq:~/volttron$ vcfg

Your VOLTRON_HOME currently set to: /home/d3x140/.volttron

Is this the volttron you are attempting to setup? [Y]:
Message bus set to zmq
What is the vip address? [tcp://127.0.0.1]:
What is the port for the vip address? [22916]:
Is this instance web enabled? [N]:
Will this instance be controlled by volttron central? [Y]:
Configuring /home/d3x140/volttron/services/core/VolttronCentralPlatform.
What is the name of this instance? [volttron1]: collector1
What is the hostname for volttron central? [http://node-zmq]: https://central
What is the port for volttron central? [8080]:
Should the agent autostart? [N]:
Would you like to install a platform historian? [N]:
Would you like to install a master driver? [N]: y
Configuring /home/d3x140/volttron/services/core/MasterDriverAgent.
Would you like to install a fake device on the master driver? [N]: y
Should the agent autostart? [N]: y
Would you like to install a listener agent? [N]:
Finished configuration!

You can now start the volttron instance.

If you need to change the instance configuration you can edit
the config file is at /home/d3x140/.volttron/config
```

Please note the Volttron Central web-address should point to that of the “central” instance.

Start VOLTRON instance and check if the agents are installed.

```
./start-volttron
vctl status
```

Start Volttron Central Platform on this platform manually.

```
vctl start --tag vcp
```

Check the VOLTRON log in the “central” instance, you will see “authentication failure” entry from the incoming connection. You will need to add the public key of VCP agent on the “central” instance.

```
2019-06-13 11:37:52,023 () volttron.platform.auth INFO: authentication failure: domain='vip', address
='172.20.214.69', mechanism='CURVE', credentials=['f5XjmsCN_n4MLV5uKpZmWn49ZzleRYinGrMpgeYmQwg']
```

On the “node-zmq” box execute this command and grab the public key of the VCP agent.

```
vctl auth publickey
```

Add auth entry corresponding to VCP agent on “central” instance using the below command. Replace the user id value and credentials value appropriately before running

```
vctl auth add --user_id <any unique user id. for example zmq_node_vcp> --credentials  
↪<public key of vcp on zmq node>
```

Complete similar steps to start a forwarder agent that connects to “central” instance. Modify the configuration in *services/core/ForwardHistorian/rmq_config.yml* to have a destination VIP address pointing to VIP address of the “central” instance and server key of the “central” instance.

```
----  
destination-vip: tcp://<ip>:22916  
destination-serverkey: <serverkey>
```

Note: Replace <ip> with public facing IP-address of “central” instance and <serverkey> with serverkey of “central” instance. Use the command **vctl auth serverkey** on the “central” instance to get the server key of the instance

Install and start forwarder agent.

```
python scripts/install-agent.py -s services/core/ForwardHistorian -c services/core/  
↪ForwardHistorian/rmq_config.yml --start
```

Grab the public key of the forwarder agent.

```
vctl auth publickey
```

Add auth entry corresponding to VCP agent on **central** instance.

```
vctl auth add --user_id <any unique user id. for example zmq_node_forwarder> --  
↪credentials <public key of forwarder on zmq node>
```

You should start seeing messages from “collector1” instance on the “central” instance’s VOLTTRON log now.

```
2019-06-13 12:02:45.099 (listeneragent-3.2 10405) listener_agent INFO: Peer: pubsub, Sender: proxy_router, Bus: , Topic: devices/fake-campus/fake-building/fake-device/  
all, Headers: {'X-Forwarded': True, 'SynchronizedTimeStamp': '2019-06-13T19:02:45.000000+00:00', 'TimeStamp': '2019-06-13T19:02:45.001920+00:00', 'X-Forwarded-From': 'c  
ollector1', 'Date': '2019-06-13T19:02:45.001920+00:00', 'min_compatible_version': '5.0', 'max_compatible_version': 'u'}, Message:  
[{'Heartbeat': True, 'PowerState': 0, 'ValveState': 0, 'temperature': 50.0},  
{ 'Heartbeat': {'type': 'integer', 'tz': 'US/Pacific', 'units': 'On/Off'},  
  'PowerState': {'type': 'integer', 'tz': 'US/Pacific', 'units': '1/0'},  
  'ValveState': {'type': 'integer', 'tz': 'US/Pacific', 'units': '1/0'},  
  'temperature': {'type': 'integer',  
                  'tz': 'US/Pacific',  
                  'units': 'Fahrenheit'}}]  
2019-06-13 12:02:45.097 (listeneragent-3.2 10403) listener_agent INFO: Peer: pubsub, Sender: proxy_router, Bus: , Topic: devices/fake-campus/fake-building/fake-device/  
all, Headers: {'X-Forwarded': True, 'SynchronizedTimeStamp': '2019-06-13T19:02:45.000000+00:00', 'TimeStamp': '2019-06-13T19:02:45.001920+00:00', 'X-Forwarded-From': 'c  
ollector1', 'Date': '2019-06-13T19:02:45.001920+00:00', 'min_compatible_version': '5.0', 'max_compatible_version': 'u'}, Message:  
[{'Heartbeat': True, 'PowerState': 0, 'ValveState': 0, 'temperature': 50.0},  
{ 'Heartbeat': {'type': 'integer', 'tz': 'US/Pacific', 'units': 'On/Off'},  
  'PowerState': {'type': 'integer', 'tz': 'US/Pacific', 'units': '1/0'},  
  'ValveState': {'type': 'integer', 'tz': 'US/Pacific', 'units': '1/0'},  
  'temperature': {'type': 'integer',  
                  'tz': 'US/Pacific',  
                  'units': 'Fahrenheit'}}]
```

Node-RMQ Instance Setup

Note: This instance must have been bootstrapped using `--rabbitmq` see `Building-VOLTTRON#steps-for-rabbitmq`.

Using “vcfg” command, install Volttron Central Platform agent, a master driver agent with fake driver. The instance name is set to “collector2”.

```

(volttron)nidd494@node-rmq:~/volttron$ vcfg

Your VOLTRON_HOME currently set to: /home/nidd494/.volttron

Is this the volttron you are attempting to setup? [Y]:
What type of message bus (rmq/zmq)? [zmq]: rmq

The rmq message bus has a backward compatibility
layer with current zmq instances. What is the
zmq bus's vip address? [tcp://127.0.0.1]:
What is the port for the vip address? [22916]:
Is this instance web enabled? [N]:
Will this instance be controlled by volttron central? [Y]:
Configuring /home/nidd494/volttron/services/core/VolttronCentralPlatform.
What is the name of this instance? [collector2]:
What is the hostname for volttron central? [http://node-rmq]: https://central
What is the port for volttron central? [8080]: 8443
Should the agent autostart? [N]:
Would you like to install a platform historian? [N]:
Would you like to install a master driver? [N]: y
Configuring /home/nidd494/volttron/services/core/MasterDriverAgent.
Would you like to install a fake device on the master driver? [N]: y
Should the agent autostart? [N]: y
Would you like to install a listener agent? [N]:
Finished configuration!

You can now start the volttron instance.

If you need to change the instance configuration you can edit
the config file is at /home/nidd494/.volttron/config

(volttron)nidd494@node-rmq:~/volttron$ sudo vi /etc/hosts
[sudo] password for nidd494:
(volttron)nidd494@node-rmq:~/volttron$ █

```

Note: The Volttron Central web-address should point to that of the “central” instance.

Start VOLTRON instance and check if the agents are installed.

```

./start-volttron
vctl status

```

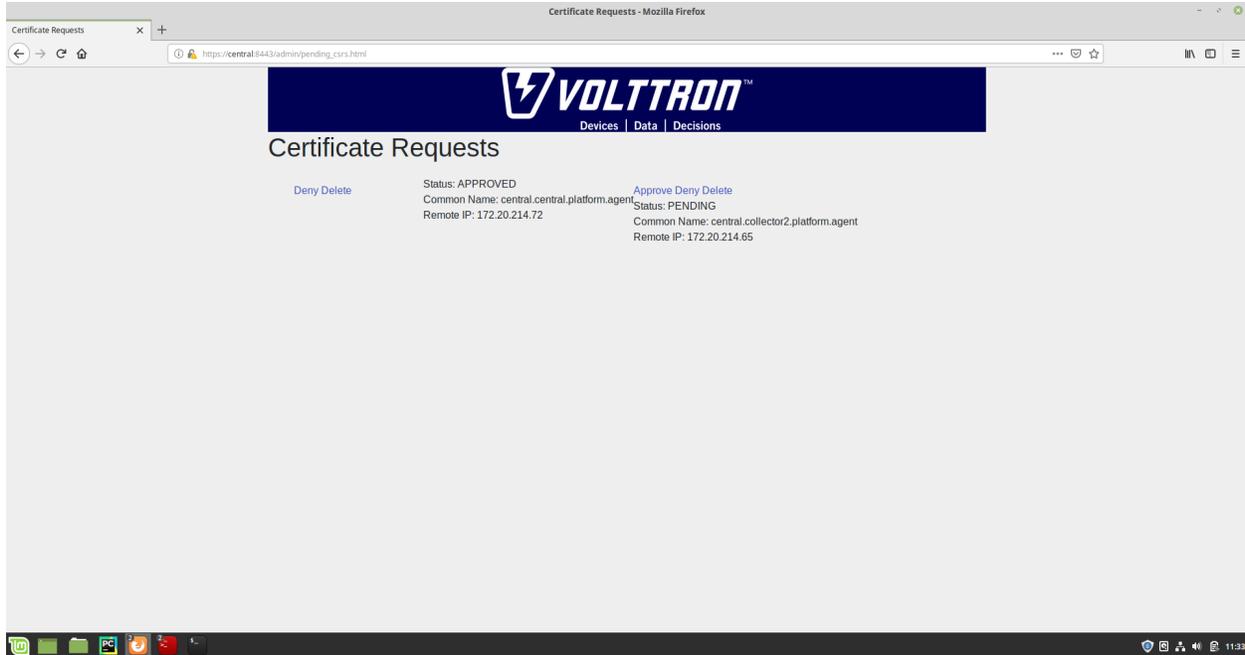
Start Volttron Central Platform on this platform manually.

```

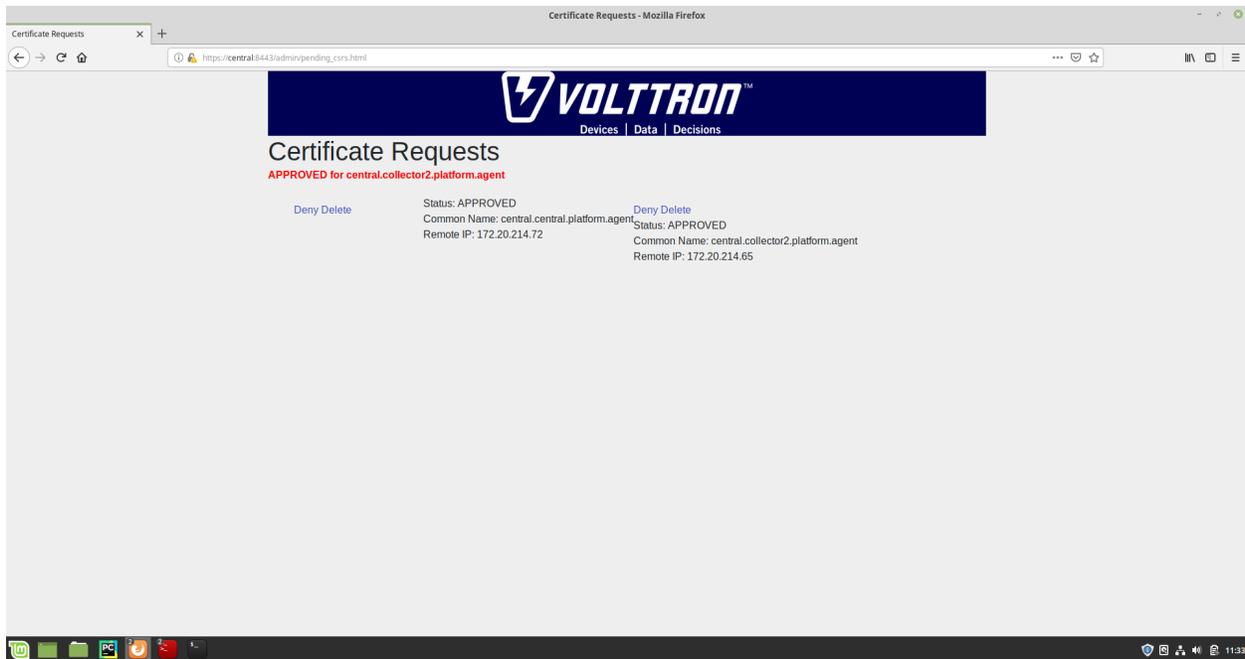
vctl start --tag vcp

```

Go the master admin authentication page and check if there is a new pending CSR request from VCP agent of “collector2” instance.



Approve the CSR request to allow authenticated SSL based connection to the “central” instance.



Now go back to the terminal and check the status of Volttron Central Platform agent. It should be set to “GOOD”.

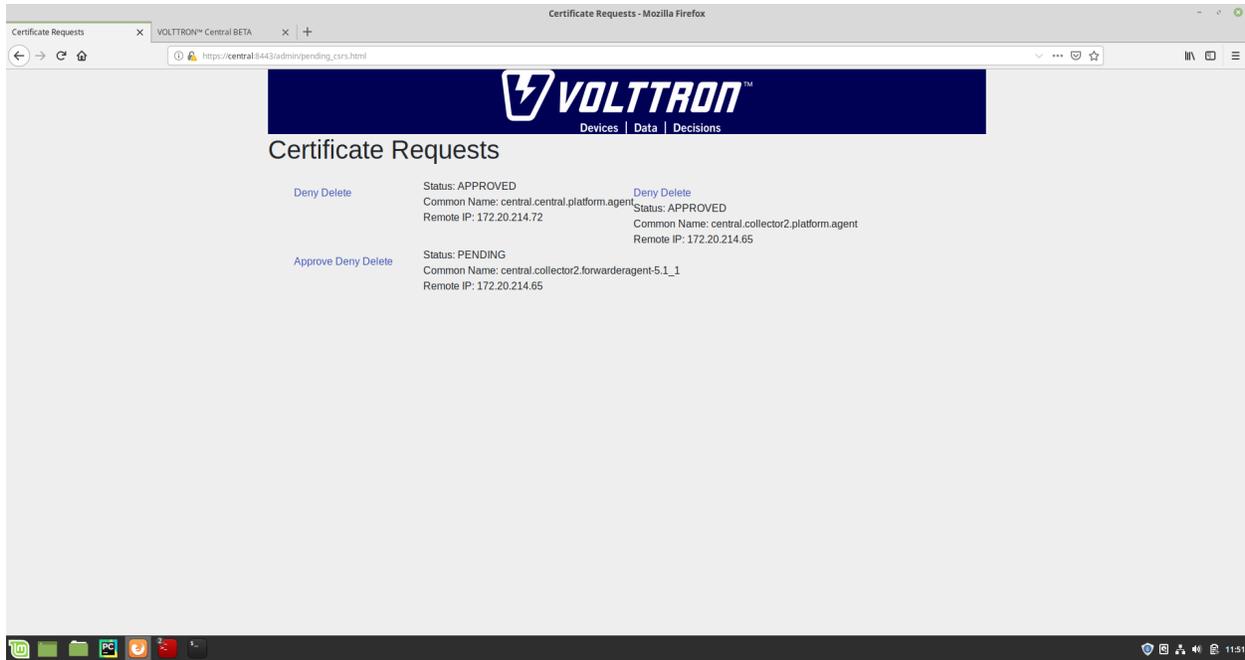
Let’s now install a forwarder agent on this instance to forward local messages matching “devices” topic to external “central” instance. Modify the configuration in `services/core/ForwardHistorian/rmq_config.yml` to have a destination address pointing to web address of the “central” instance.

```
---
destination-address: https://central:8443
```

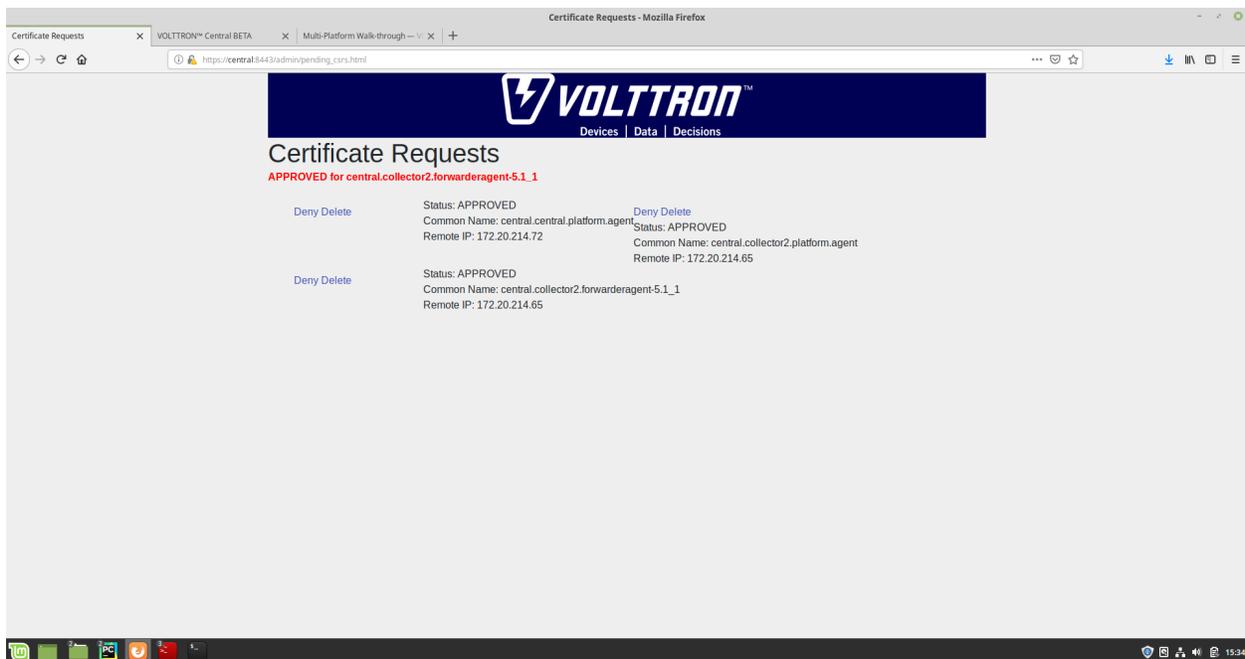
Start forwarder agent.

```
python scripts/install-agent.py -s services/core/ForwardHistorian -c services/core/
↵ForwardHistorian/rmq_config.yml --start
```

Go to the master admin authentication page and check if there is a new pending CSR request from forwarder agent of “collector2” instance.



Approve the CSR request to allow authenticated SSL based connection to the “central” instance.



Now go back to the terminal and check the status of forwarder agent. It should be set to “GOOD”.

Check the VOLTRON log of “central” instance. You should see messages with “devices” topic coming from “collector2” instance.

```
2019-06-13 12:00:00.095 (listeneragent-3.2 10403) listener.agent INFO: Peer: pubsub, Sender: collector2.forwarderagent-5.1.1, Bus: , Topic: devices/fake-campus/fake-building/fake-device/all, Headers: {'X-Forwarded': True, 'SynchronizedTimestamp': '2019-06-13T19:00:00.000000+00:00', 'Timestamp': '2019-06-13T19:00:00.004159+00:00', 'X-Forwarded-From': 'collector2', 'Date': '2019-06-13T19:00:00.004159+00:00', 'min-compatible-version': '5.0', 'max-compatible-version': '0'}, Message:
[{'Heartbeat': True, 'PowerState': 0, 'ValveState': 0, 'Temperature': 50.0},
 {'Heartbeat': {'type': 'integer', 'tz': 'US/Pacific', 'units': 'On/Off'},
  'PowerState': {'type': 'integer', 'tz': 'US/Pacific', 'units': '1/0'},
  'ValveState': {'type': 'integer', 'tz': 'US/Pacific', 'units': '1/0'},
  'Temperature': {'type': 'integer',
                  'tz': 'US/Pacific',
                  'units': 'Fahrenheit'}}]
```

To confirm that VolttronCentral is monitoring the status of all the 3 platforms, open a browser and type this URL <https://central:8443/vc/index.html>. Login using credentials (username and password) earlier set during the VC configuration step (using `vcfg` command in “central” instance). Click on “platforms” tab in the far right corner. You should see all three platforms listed in that page. Click on each of the platforms and check the status of the agents.



Multi-Platform Deployment With RabbitMQ Message bus

In ZeroMQ based VOLTTRON, if multiple instances needed to be connected together and be able to send or receive messages to/from remote instances we would do it in few different ways.

1. Write an agent that would connect to remote instance directly and publish/subscribe to messages or perform RPC communication directly. This is described in [Agent connection to remote volttron instance](#)
2. Use special agents such as forwarder/data puller agents to forward/receive messages to/from remote instances. This can be achieved using RabbitMQ’s shovel plugin and is described at [Using Shovel Plug-in](#)
3. Configure vip address of all remote instances that an instance has to connect to in its `$VOLTTRON_HOME/external_discovery.json` and let the router module in each instance manage the connection and take care of the message routing for us. This is the most seamless way to do multi-platform communication. This can be achieved using RabbitMQ’s federation plugin. Setup for this is described at [Using Federation Plug-in](#)

RabbitMQ Trouble Shooting

Check the status of the federation connection

```
`${RABBITMQ_HOME}/sbin/rabbitmqctl eval 'rabbit_federation_status:status().'
```

If everything is properly configured, then the status is set to “running”. If not look for the error status. Some of the typical errors are,

- a. “failed_to_connect_using_provided_uris” - Check if RabbitMQ user is created in downstream server node. Refer to step 3b of federation setup
- b. “unknown ca” - Check if the root CAs are copied to all the nodes correctly. Refer to step 2 of federation setup
- c. “no_suitable_auth_mechanism” - Check if the AMPQ/S ports are correctly configured.

Check the status of the shovel connection

```
RABBITMQ_HOME/sbin/rabbitmqctl eval 'rabbit_shovel_status:status().'
```

If everything is properly configured, then the status is set to “running”. If not look for the error status. Some of the typical errors are,

- a. “failed_to_connect_using_provided_uris” - Check if RabbitMQ user is created in subscriber node. Refer to step 3 b of shovel setup
- b. “unknown ca” - Check if the root CAs are copied to remote servers correctly. Refer to step 2 of shovel setup
- c. “no_suitable_auth_mechanism” - Check if the AMPQ/S ports are correctly configured.

Check the RabbitMQ logs for any errors

```
tail -f <voltrtron source dir>/rabbitmq.log
```

Rabbitmq startup hangs

- a. Check for errors in rabbitmq log. There is a rabbitmq.log file in your voltrtron source directory that is a symbolic link to the rabbitmq server logs.
- b. Check for errors in syslog (/var/log/syslog or /var/log/messages)
- c. If there are no errors in either of the logs, stop rabbitmq and starting rabbitmq server in foreground and see if there are any errors written on the console. Once you find the error you can kill the process by entering Ctl+C, fix the error and start rabbitmq again using ./start-rabbitmq from voltrtron source directory.

```
./stop-voltrtron
./stop-rabbitmq
@RABBITMQ_HOME/sbin/rabbitmq-server
```

SSL trouble shooting

There are few things that are essential for SSL certificates to work right.

- a. Please use a unique common-name for CA certificate for each voltrtroninstance. This is configured under certificate-data in the rabbitmq_config.yml or if no yml file is used while configuring a voltrtron single instance (using vcfg -rabbitmq single). Certificate generated for agent will automatically get agent’s vip identity as the certificate’s common-name
- b. host name in ssl certificate should match hostname used to access the server. For example, if the fully qualified domain name was configured in the certificate-data, you should use the fully qualified domain name to access rabbitmq’s management url.
- c. Check if your system time is correct especially if you are running virtual machines. If the system clock is not right, it could lead to ssl certificate errors

DataMover troubleshooting

If output from voltrtron.log is not as expected check for {'alert_key': 'historian_not_publishing'} in the callee node’s voltrtron.log. Most likely cause is the historian is not running properly or credentials between caller and callee nodes was not set properly.

RabbitMQ Multi-Platform Deployment Using Federation Plugin

Federation plugin allows us to send and receive messages to/from remote instances with few simple connection settings. Once a federation link is established to remote instance, the messages published on the remote instance become available to local instance as if it were published on the local instance. Before, we illustrate the steps to setup a federation link, let us start by defining the concept of upstream and downstream server.

Upstream Server - The node that is publishing some message of interest

DownStream Server - The node that wants to receive messages from the upstream server

A federation link needs to be established from downstream server to the upstream server. The data flows in single direction from upstream server to downstream server. For bi-directional data flow we would need to create federation links on both the nodes.

1. Setup two VOLTRON instances using the instructions at RMQ Setup. **Please note that each instance should have a unique instance name and should be running on machine/VM that has a unique host name.**
2. In a multi platform setup that need to communicate with each other with RabbitMQ over SSL, each VOLTRON instance should should trust the ROOT CA of the other instance(RabbitMQ root ca)
 - a. Transfer (scp/sftp/similar) `volttron_home/certificates/certs/<instance_name>-root-ca.crt` to a temporary location on the other volttron instance machine. For example, if you have two instance v1 and v2, scp v1's v1-root-ca.crt to v2 and v2-root-ca.crt to v1.

Note: If using VMs, in order to scp files between VM openssh should be installed and running.

- b. Append the contents of the transferred root ca to the instance's `trusted-cas.crt` file. Do this on both the instances. Now both the instances `<instance_name>-trusted-cas.crt` will have two certificates.

For example:

On v1: `cat /tmp/v2-root-ca.crt >> VOLTRON_HOME/certificates/certs/v1-trusted-cas.crt`

On v2: `cat /tmp/v1-root-ca.crt >> VOLTRON_HOME/certificates/certs/v2-trusted-cas.crt`

3. Stop volttron, stop rabbitmq server and start volttron on both the instances. This is required only when you update the root certificate and not required when you add a new shovel/federation between the same hosts

```
./stop-volttron
./stop-rabbitmq
./start-volttron
```

4. Identify upstream servers (publisher nodes) and downstream servers (collector nodes). To create a RabbitMQ federation, we have to configure upstream servers on the downstream server and make the VOLTRON exchange "federated".

- a. On the downstream server (collector node)

```
vcfg --rabbitmq federation [optional path to rabbitmq_federation_config.yml
containing the details of the upstream hostname, port and vhost.
```

Example configuration for federation is available in `examples/configurations/rabbitmq/rabbitmq_federation_config.yml`

If no config file is provided, the script will prompt for hostname (or IP address), port, and vhost of each upstream node you would like to add. Hostname provided should match the hostname in the SSL certificate of the upstream server. For bi-directional data flow, we will have to run the same script on both the nodes.

- b. Create a user in the upstream server(publisher) with `username=<downstream admin user name>` (i.e. (instance-name)-admin) and provide it access to the virtual host of the upstream RabbitMQ server. Run the below command in the upstream server

```

volttron-ctl rabbitmq add-user <username> <password>
Do you want to set READ permission [Y/n]
Do you want to set WRITE permission [Y/n]
Do you want to set CONFIGURE permission [Y/n]

```

5. Test the federation setup.

- a. On the downstream server run a listener agent which subscribes to messages from all platforms
 - Open the file `examples/ListenerAgent/listener/agent.py`. Search for `@PubSub.subscribe('pubsub', '')` and replace that line with `@PubSub.subscribe('pubsub', 'devices', all_platforms=True)`
 - upgrade the listener

```
scripts/core/upgrade-listener
```

- b. Install master driver, configure fake device on upstream server and start volttron and master driver. `vcfg --agent master_driver` command can install master driver and setup a fake device.

```

./stop-volttron
vcfg --agent master_driver
./start-volttron
vctl start --tag master_driver

```

- c. Verify listener agent in downstream VOLTRON instance is able to receive the messages. downstream volttron instance's `volttron.log` should display device data scrapped by master driver agent in upstream volttron instance
6. Open ports and https service if needed On Redhat based systems ports used by RabbitMQ (defaults to 5671, 15671 for SSL, 5672 and 15672 otherwise) might not be open by default. Please contact system administrator to get ports opened on the downstream server.

Following are commands used on centos 7.

```

sudo firewall-cmd --zone=public --add-port=15671/tcp --permanent
sudo firewall-cmd --zone=public --add-port=5671/tcp --permanent
sudo firewall-cmd --reload

```

7. How to remove federation link

- a. Using the management web interface

Log into management web interface using downstream server's admin username. Navigate to admin tab and then to federation management page. The status of the upstream link will be displayed on the page. Click on the upstream link name and delete it.

- b. Using "volttron-ctl" command on the upstream server.

```

vctl rabbitmq list-federation-parameters
NAME                               URI
upstream-volttron2-rabbit-2        amqps://rabbit-2:5671/volttron2?
↪cacertfile=/home/nidd494/.volttron1/certificates/certs/volttron1-
↪root-ca.crt&certfile=/home/nidd494/.volttron1/certificates/certs/
↪volttron1-admin.crt&keyfile=/home/nidd494/.volttron1/certificates/
↪private/volttron1-admin.pem&verify=verify_peer&fail_if_no_peer_
↪cert=true&auth_mechanism=external&server_name_indication=rabbit-2

```

Grab the upstream link name and run the below command to remove it.

```
vctl rabbitmq remove-federation-parameters upstream-volttron2-
↔rabbit-2
```

Agent communication to Remote RabbitMQ instance

Communication between two RabbitMQ based VOLTTRON instances must be done using SSL certificate based authentication. Non SSL based authentication will not be supported for communication to remote RabbitMQ based VOLTTRON instances. An volttron instance that wants to communicate with a remote instance should first request a SSL certificate that is signed by the remote instance. To facilitate this process there will be a web based server api for requesting, listing, approving and denying certificate requests. This api will be exposed via the MasterWebService and will be available to any RabbitMQ based VOLTTRON instance with ssl enabled. This api will be tested and used in the following agents:

- ForwarderAgent
- DataPuller
- VolttronCentralPlatform

For the following document we will assume we have two instances a local-instance and remote-volttron-instance. The remote-volttron-instance will be configured to allow certificate requests to be sent to it from the local-instance. A remote-agent running in local-instance will attempt to establish a connection to the remote-volttron-instance

Configuration

Both volttron-server and volttron-client must be configured for RabbitMQ message bus with SSL using the step described at *Installing Volttron*.

In addition the remote-volttron-instance configuration file must have a https bind-web-address specified in the instance config file. Below is an example config file with bind-web-address. Restart volttron after editing the config file

```
[volttron]
message-bus = rmq
vip-address = tcp://127.0.0.1:22916
bind-web-address = https://volttron1:8443
instance-name = volttron1
```

By default the *bind-web-address* parameter will use the *MasterWebService* agent's certificate and private key. Both private and public key are necessary in order to bind the port to the socket for incoming connections. This key pair is auto generated for RabbitMQ based VOLTTRON at the time of platform startup. Users can provide a different certificate and private key to be used for the bind-web-address by specifying web-ssl-cert and web-ssl-key in the config file. Below is an example config file with the additional entries

```
[volttron]
message-bus = rmq
vip-address = tcp://127.0.0.1:22916
bind-web-address = https://volttron1:8443
instance-name = volttron1
web-ssl-cert = /path/to/cert/cert.pem
web-ssl-key = /path/to/cert/key.pem
```

Note:

- The */etc/hosts* file should be modified in order for the dns name to be used for the bound address.
-

remote-agent on local-instance

The *auth* subsystem of the volttron architecture is how a remote-agent on local instance will connect to the remote volttron instance.

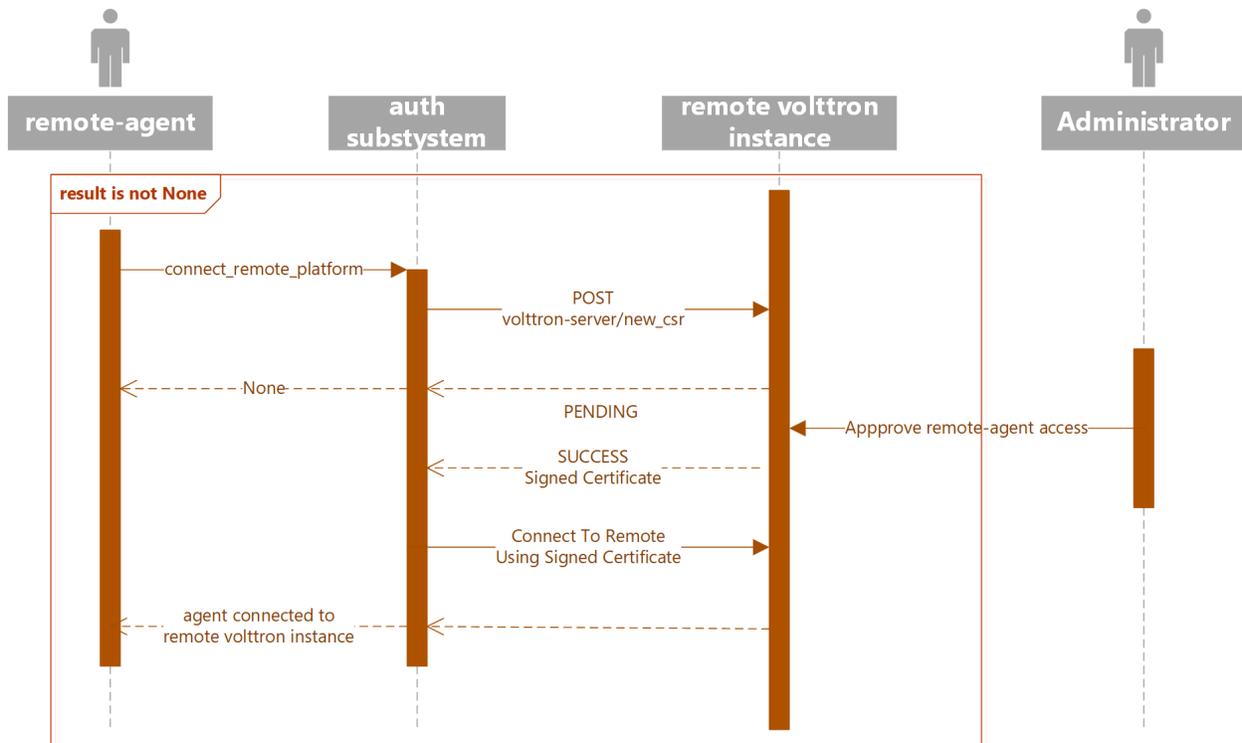
The following is a code snippet from the remote-agent to connect to the remote volttron instance.

```
...
value = self.vip.auth.connect_remote_platform(address)
```

The above function call will return an agent that connects to the remote instance only after the request is approved by an administrator of the remote instance. It is up to the agent to repeat calling *connect_remote_platform* periodically until an agent object is obtained.

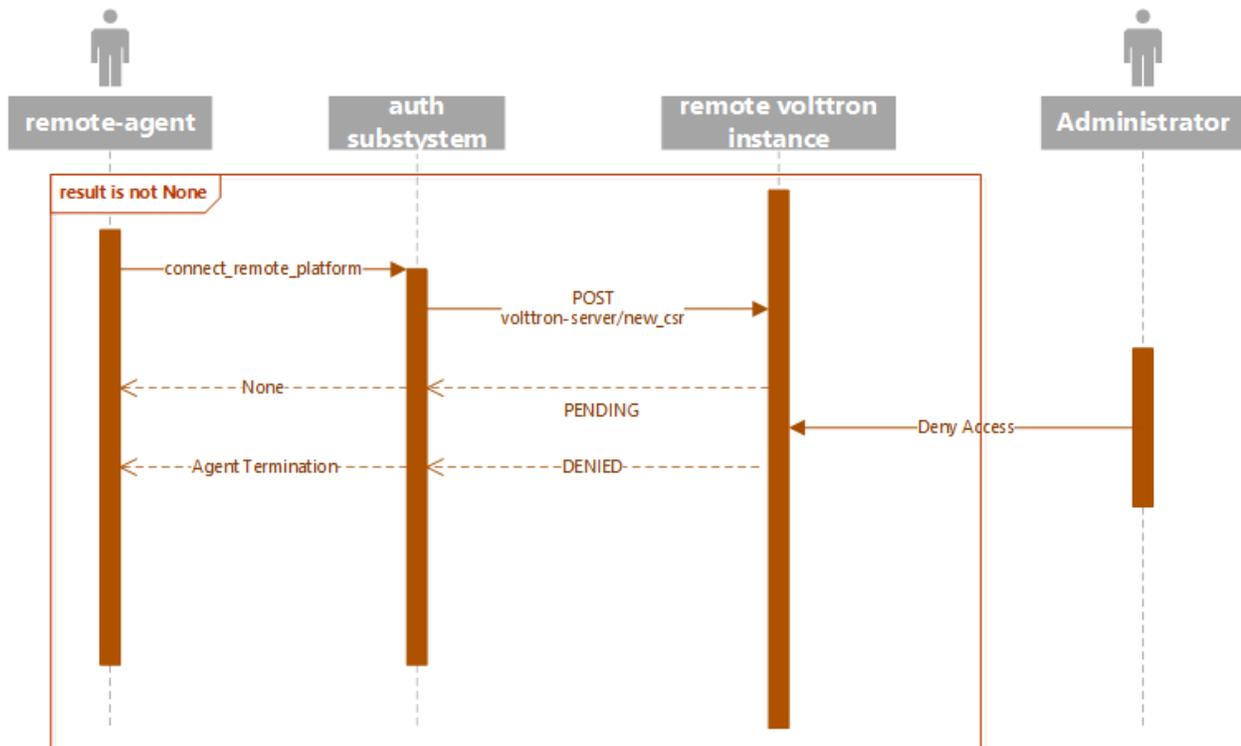
Approving a CSR Request

The following diagram shows the sequence of events when an access request is approved by the administrator of remote volttron instance. In this case, the volttron-client agent will get a *Agent* object that is connected to the remote instance. The diagram shows the client agent repeating the call to *connect_remote_platform* until the return value is not *None*.



Denying a CSR Request

The following diagram shows the sequence of events when an access request is denied by the administrator. The client agent repeats the call to *connect_remote_platform* until the return value is not *None*. When the remote instance's administrator denies a access request, the auth subsystem will raise an alert and shutdown the agent.



Follow walk-through in Multi-Platform Multi-Bus Walk-through for setting up different combinations of multi-bus multi-platform setup using CSR.

RabbitMQ Multi-Platform Deployment Using Shovel Plugin

In RabbitMQ based VOLTTRON, forwarder and data mover agents will be replaced by shovels to send or receive remote pubsub messages. Shovel behaves like a well written client application that connects to its source (can be local or remote) and destination (can be local or remote instance), reads and writes messages, and copes with connection failures. In case of shovel, apart from configuring the hostname, port and virtual host of the remote instance, we will also have to provide list of topics that we want to forward to remote instance. Shovels can also be used for remote RPC communication in which case we would have to create shovel in both the instances, one to send the RPC request and other to send the response back.

Pubsub Communication

1. Setup two VOLTTRON instances using the steps described in installation section. Please note that each instance should have a unique instance name.
2. In a multi platform setup that need to communicate with each other with RabbitMQ over SSL, each VOLTTRON instance should should trust the ROOT CA of the other instance(RabbitMQ root ca)
 - a. Transfer (scp/sftp/similar) volttron_home/certificates/certs/<instance_name>-root-ca.crt to a temporary location on the other volttron instance machine. For example, if you have two instance v1 and v2, scp v1's v1-root-ca.crt to v2 and v2-root-ca.crt to v1.
 - b. Append the contents of the transferred root ca to the instance's root ca.

For example:

On v1

```
cat /tmp/v2-root-ca.crt >> VOLTRON_HOME/certificates/v1-root-ca.crt
```

On v2

```
cat /tmp/v1-root-ca.crt >> VOLTRON_HOME/certificates/v2-root-ca.crt
```

3. Identify the instance that is going to act as the “publisher” instance. Suppose “v1” instance is the “publisher” instance and “v2” instance is the “subscriber” instance. Then we need to create a shovel on “v1” to forward messages matching certain topics to remote instance “v2”.

- a. On the publisher node,

```
vcfg --rabbitmq shovel [optional path to rabbitmq_shovel_config.yml]
```

rabbitmq_shovel_config.yml should contain the details of the remote hostname, port, vhost and list of topics to forward. Example configuration for shovel is available in examples/configurations/rabbitmq/rabbitmq_shovel_config.yml

For this example, let’s set the topic to “devices”

If no config file is provided, the script will prompt for hostname (or IP address), port, vhost and list of topics for each remote instance you would like to add. For bi-directional data flow, we will have to run the same script on both the nodes.

- b. Create a user in the subscriber node with username set to publisher instance’s agent name ((instance-name)-PublisherAgent) and allow the shovel access to the virtual host of the subscriber node.

```
cd $RABBITMQ_HOME
vctl add-user <username> <password>
```

4. Test the shovel setup.

- a. Start VOLTRON on publisher and subscriber nodes.

- b. On the publisher node, start a master driver agent that publishes messages related to a fake device. (Easiest way is to run volttron-cfg command and follow the steps)

- c. On the subscriber node, run a listener agent which subscribes to messages from all platforms (set @PubSub.subscribe(‘pubsub’, ‘devices’, all_platforms=True) instead of @PubSub.subscribe(‘pubsub’, ‘’))

- d. Verify listener agent in subscriber node is able to receive the messages matching “devices” topic.

5. How to remove the shovel setup.

- a. Using the management web interface

Log into management web interface using publisher instance’s admin username. Navigate to admin tab and then to shovel management page. The status of the shovel will be displayed on the page. Click on the shovel name and delete the shovel.

- b. Using “volttron-ctl” command on the publisher node.

```
vctl rabbitmq list-shovel-parameters
NAME                               SOURCE ADDRESS
↪                               DESTINATION ADDRESS
↪                               BINDING KEY
shovel-rabbit-3-devices  amqps://rabbit-1:5671/volttron1?cacertfile=/home/
↪nidd494/.volttron1/certificates/certs/volttron1-root-ca.crt&certfile=/
↪home/nidd494/.volttron1/certificates/certs/volttron1-admin.crt&keyfile=/
↪home/nidd494/.volttron1/certificates/private/volttron1-admin.pem&
↪verify=verify_peer&fail_if_no_peer_cert=true&auth_mechanism=external&
↪server_name_indication=rabbit-1  amqps://rabbit-3:5671/volttron3?
↪cacertfile=/home/nidd494/.volttron1/certificates/certs/volttron1-root-
↪ca.crt&certfile=/home/nidd494/.volttron1/certificates/certs/volttron1-
↪admin.crt&keyfile=/home/nidd494/.volttron1/certificates/private/
↪volttron1-admin.pem&verify=verify_peer&fail_if_no_peer_cert=true&auth_
↪mechanism=external&server_name_indication=rabbit-3  __pubsub__.
↪volttron1.devices.#
```

(continued from previous page)

Grab the shovel name and run the below command to remove it.

```
vctl rabbitmq remove-shovel-parameters shovel-rabbit-3-devices
```

RPC Communication

Following are the steps to create Shovel for multi-platform RPC communication.

1. Setup two VOLTRON instances using the steps described in installation section. Please note that each instance should have a unique instance name.
2. In a multi platform setup that need to communicate with each other with RabbitMQ over SSL, each VOLTRON instance should should trust the ROOT CA of the other instance(RabbitMQ root ca)

- a. Transfer (scp/sftp/similar) `volttron_home/certificates/certs/<instance_name>-root-ca.crt` to a temporary location on the other volttron instance machine. For example, if you have two instance v1 and v2, scp v1's v1-root-ca.crt to v2 and v2-root-ca.crt to v1.

- b. Append the contents of the transferred root ca to the instance's root ca. For example:

On v1

```
cat /tmp/v2-root-ca.crt >> VOLTRON_HOME/certificates/v1-root-ca.crt
```

On v2

```
cat /tmp/v1-root-ca.crt >> VOLTRON_HOME/certificates/v2-root-ca.crt
```

3. Typically RPC communication is 2 way communication so we will to setup shovel in both the VOLTRON instances. In RPC calls there are two instances of shovel. One serving as the caller (makes RPC request) and the other acting as a callee (replies to RPC request). Identify the instance is the "caller" and which is the "callee." Suppose "v1" instance is the "caller" instance and "v2" instance is the "callee" instance.
 - a. On both the caller and callee nodes, shovel instances need to be created. In this example, v1's shovel would forward the RPC call request from an agent on v1 to v2 and similarly v2's shovel will forward the RPC reply from agent on v2 back to v1.

```
vcfg --rabbitmq shovel [optional path to rabbitmq_shovel_config.yml]
```

`rabbitmq_shovel_config.yml` should contain the details of the **remote** hostname, port, vhost, volttron instance name (so in v1's yml file parameters would point to v2 and vice versa), and list of agent pair identities (local caller, remote callee). Example configuration for shovel is available in `examples/configurations/rabbitmq/rabbitmq_shovel_config.yml`.

For this example, let's say that we are using the `schedule-example` and `acutator` agents.

For v1, the agent pair identities would be:

- [Scheduler, platform.actuator]

For v2, they would be:

- [platform.actuator, Scheduler]

Indicating the flow from local agent to remote agent.

- b. On the caller node create a user with username set to callee instance's agent name ((instance-name)-RPCcallee) and allow the shovel access to the virtual host of the callee node. Similarly, on the callee

node, create a user with username set to caller instance's agent name ((instance-name)-RPCCaller) and allow the shovel access to the virtual host of the caller node.

```
cd $RABBITMQ_HOME
vctl add-user <username> <password>
```

4. Test the shovel setup

a. On caller node:

Make necessary changes to RPC methods of caller agent.

For this example, in `volttron/examples/SchedulerExample/schedule_example/agent.py`:

- Search for 'campus/building/unit' in `publish_schedule` method. Replace with 'devices/fake-campus/fake-building/fake-device'
- Search for ['campus/building/unit3',start,end] in the `use_rpc` method, replace with:

```
msg = ['fake-campus/fake-building/fake-device',start,end].
```
- Add: `kwargs = {"external_platform": 'v2'}` on the line below
- On the `result = self.vip.rpc.call` method below, replace "`msg.get(timeout=10)`" with:

```
msg, **kwargs).get(timeout=10),
```

- In the second try clause of the `use_rpc` method:
- Replace `result['result']` with `result[0]['result']`
- Add `kwargs = {"external_platform": 'v2'}` as the first line of the if statement
- Replace 'campus/building/unit3/some_point' with 'fake-campus/fake-building/fake-device/PowerState'
- Below 'fake-campus/fake-building/fake-device/PowerState' add: 0,
- Replace

```
'0.0').get(timeout=10) with **kwargs).get(timeout=10)
```

Next, install an example scheduler agent and start it:

```
#!/bin/bash
python /home/username/volttron/scripts/install-agent.py -c /home/username/
↪ volttron/examples/SchedulerExample/schedule-example.agent -s examples/
↪ SchedulerExample --start --force -i Scheduler
```

b. On the callee node:

- Run upgrade script to install actuator agent.

```
#!/bin/bash
python /home/username/volttron/scripts/install-agent.py -s services/
↪ core/ActuatorAgent --start --force -i platform.actuator
```

- Run the upgrade script to install the listener agent.

```
scripts/core/upgrade-listener
```

- Install master driver, configure fake device on upstream callee and start volttron and master driver. `vcfg -agent master_driver` command can install master driver and setup a fake device.

```
./stop-volttron
vcfg --agent master_driver
./start-volttron
vctl start --tag master_driver
```

- Start actuator agent and listener agents.

The output for the callee node with a successful shovel run should look similar to:

```
2018-12-19 15:38:00,009 (listeneragent-3.2 13039) listener.agent INFO:↵
↵Peer: pubsub, Sender: platform.driver:, Bus: , Topic: devices/fake-
↵campus/fake-building/fake-device/all, Headers: {'Date': '2018-12-
↵19T20:38:00.001684+00:00', 'TimeStamp': '2018-12-19T20:38:00.
↵001684+00:00', 'min_compatible_version': '5.0', 'max_compatible_version
↵': 'u', 'SynchronizedTimeStamp': '2018-12-19T20:38:00.000000+00:00'},↵
↵Message:
  [{'Heartbeat': True, 'PowerState': 0, 'ValveState': 0, 'temperature': 50.
↵0},
  {'Heartbeat': {'type': 'integer', 'tz': 'US/Pacific', 'units': 'On/Off'}
↵},
  {'PowerState': {'type': 'integer', 'tz': 'US/Pacific', 'units': '1/0'},
  'ValveState': {'type': 'integer', 'tz': 'US/Pacific', 'units': '1/0'},
  'temperature': {'type': 'integer',
                  'tz': 'US/Pacific',
                  'units': 'Fahrenheit'}}]
```

DataMover Communication

The DataMover historian running on one instance makes RPC call to platform historian running on remote instance to store data on remote instance. Platform historian agent returns response back to DataMover agent. For such a request-response behavior, shovels need to be created on both instances.

1. Please ensure that preliminary steps for multi-platform communication are completed (namely, steps 1-3 described above).
2. To setup a data mover to send messages from local instance (say v1) to remote instance (say v2) and back, we would need to setup shovels on both instances.

Example of RabbitMQ shovel configuration on v1

```
shovel:
# hostname of remote machine
rabbit-2:
  port: 5671
  rpc:
    # Remote instance name
    v2:
      # List of pair of agent identities (local caller, remote callee)
      - [data.mover, platform.historian]
  virtual-host: v1
```

This says that DataMover agent on v1 wants to make RPC call to platform historian on v2.

```
vcfg --rabbitmq shovel [optional path to rabbitmq_shovel_config.yml]
```

Example of RabbitMQ shovel configuration on v2

```
shovel:
# hostname of remote machine
rabbit-1:
  port: 5671
  rpc:
    # Remote instance name
  v1:
    # List of pair of agent identities (local caller, remote callee)
    - [platform.historian, data.mover]
virtual-host: v2
```

This says that Hplatform historian on v2 wants to make RPC call to DataMover_↵
↵agent on v1.

a. On v1, run below command to setup a shovel from v1 to v2.

```
vcfg --rabbitmq shovel [optional path to rabbitmq_shovel_config.yml]
```

b. Create a user on v2 with username set to remote agent's username
(for example, v1.data.mover i.e., <instance_name>.<agent_identity>) and_↵
↵allow
the shovel access to the virtual host of v2.

```
cd $RABBITMQ_HOME
vctl add-user <username> <password>
```

c. On v2, run below command to setup a shovel from v2 to v1

```
vcfg --rabbitmq shovel [optional path to rabbitmq_shovel_config.yml]
```

d. Create a user on v1 with username set to remote agent's username
(for example, v2.platform.historian i.e., <instance_name>.<agent_identity>
↵) and allow
the shovel access to the virtual host of the v1.

```
cd $RABBITMQ_HOME
vctl add-user <username> <password>
```

3. Start Master driver agent on v1

```
./stop-volttron
vcfg --agent master_driver
./start-volttron
vctl start --tag master_driver
```

4. Install DataMover agent on v1. Contents of the install script can look like below.

```
#!/bin/bash
export CONFIG=$(mktemp /tmp/abc-script.XXXXXX)
cat > $CONFIG <<EOL
{
  "destination-vip": "",
  "destination-serverkey": "",
  "destination-instance-name": "volttron2",
  "destination-message-bus": "rmq"
}
```

(continues on next page)

(continued from previous page)

```
EOL
python scripts/install-agent.py -s services/core/DataMover -c $CONFIG --start -
↪-force -i data.mover
```

Execute the install script.

5. Start platform historian of your choice on v2. Example shows starting SQLiteHistorian

```
./stop-volttron
vcfg --agent platform_historian
./start-volttron
vctl start --tag platform_historian
```

6. Observe data getting stored in sqlite historian on v2.

Speeding Up VOLTRON™ Builds

The **VOLTRON** build process is straightforward enough, but it can be a bit slow. It relies on `pip` to download, build, and install required third-party packages. The problem is that `pip` does very little to cache the results of builds, especially those which require compilation. In fact, the only thing `pip` caches is the downloaded source archives. While this speeds up the download process and lightens the burden of the Python package index (PyPi) server(s), it does little to improve the overall build speed. The majority of this document will focus on three techniques for improving VOLTRON build times, including the `pip` download cache. But before we begin, let's discuss what is involved in building VOLTRON.

Introducing bootstrap.py

VOLTRON can actually be built just like any other Python project. It includes a `setup.py` script in the project root so one can perform the standard `build`, `install`, `sdist_*`, and `bdist_*`, etc. commands. That's great if we have some project which requires VOLTRON, but that is not the typical use case for VOLTRON. Usually, and especially for developers, VOLTRON is run in a virtual environment so its dependencies can be easily met. Enter `bootstrap.py`.

Sitting next to `setup.py` in the project root is `bootstrap.py`, a script designed to *bootstrap* a virtual environment and make dependency installation a repeatable process. Bootstrapping occurs in two stages: download `virtualenv`, using it to create a virtual environment, and download and install dependencies. The first, or bootstrap, stage typically happens once. The second, or update, stage happens many times as dependencies are added or updated. It is the update stage that takes the majority of the time and is the stage we focus on in this document.

To perform the bootstrap stage, `bootstrap.py` must be executed using the system Python interpreter.

```
$ python2.7 bootstrap.py
```

The above command assumes `python2.7` is in the `PATH` and creates the virtual environment in the default `env` directory. After the virtual environment is created, the update stage is automatically started by executing `bootstrap.py` using the `python` interpreter in the newly created virtual environment. Subsequent updates must also use the interpreter in the virtual environment.

```
$ env/bin/python bootstrap.py
```

Multiple options are available to alter the behavior of `bootstrap.py`. Use the `--help` option to list options and show the script's usage message.

```

$ env/bin/python bootstrap.py --help
usage:
  bootstrap: python2.7 bootstrap.py [options]
  update:      $VIRTUAL_ENV/bin/python bootstrap.py [options]

Bootstrap and update a virtual Python environment for VOLTRON development.

optional arguments:
  -h, --help            show this help message and exit
  -q, --quiet           produce less output
  -v, --verbose        produce extra output

bootstrap options:
  --envdir VIRTUAL_ENV alternate location for virtual environment
  --force              force installing in non-empty directory
  -o, --only-virtenv  create virtual environment and exit (skip install)
  --prompt PROMPT     provide alternate prompt in activated environment
                     (default: (volttron))

update options:
  -u, --upgrade        upgrade installed packages
  -w, --wheel          build wheels in the pip wheelhouse

The first invocation of this script, which should be made using the system
Python, will create a virtual Python environment in the 'env' subdirectory in
the same directory as this script or in the directory given by the --envdir
option. Subsequent invocations of this script should use the Python executable
installed in the virtual environment.

```

Enough about `bootstrap.py`. Let's move on to the magic. As we do, please note the following about the output of commands:

- Ellipses (...) are used to denote where excessive drivel was cut out to make this document shorter and more readable
- Lines beginning with a plus (+) in `bootstrap.py` output show the actual calls to `pip` or `easy_install`, including all arguments.

Oh, yeah. That reminds me that two packages require special handling. `BACpypes` must be installed using `easy_install` because it is only offered as an egg and `pip` doesn't install from eggs. It will always be downloaded, if it isn't already installed, and will not benefit from any of the speedups below. And `pzmq` is handled separately to pass options to its `setup.py`.

Note: As of 26 Feb 2015, `BACpypes` provides a wheel. Yay! Next step: Python 3.

Okay. On with the show.

Preparation

Before building, we need to clone the [VOLTRON](#) repository. We make sure to checkout the master branch to get the latest bootstrap script which has the special sauce for the real speed-up.

```

[volttron@inamatus ~]$ git clone -b master https://github.com/VOLTRON/volttron
Cloning into 'volttron'...
remote: Counting objects: 3268, done.
remote: Compressing objects: 100% (122/122), done.
Receiving objects: 100% (3268/3268), 14.25 MiB | 749.00 KiB/s, done.

```

(continues on next page)

(continued from previous page)

```
Resolving deltas: 100% (2070/2070), done.
Checking connectivity... done.
```

Let's move into that directory where the remainder of our time will be spent.

```
[volttron@inamatus ~]$ cd volttron/
```

Now that we have the code, we are ready for testing. We'll start with the slow method and work toward the fastest.

Meh (a.k.a. Slow Method)

Since version 6.0, `pip caches` downloaded source files to speed the download and install process when a package is once again required. The default location for this cache on Linux is in `$HOME/.cache/pip` (or `~/.cache/pip`). As can be seen by the next command, we currently have no cache.

```
[volttron@inamatus volttron]$ find ~/.cache/pip
find: `/home/volttron/.cache/pip': No such file or directory
```

So let's try bootstrapping the environment. We'll use `bash's` built-in `time` command to time the execution of each bootstrap command for comparison.

```
[volttron@inamatus volttron]$ time python2.7 bootstrap.py
Creating virtual Python environment
Downloading virtualenv DOAP record
Downloading virtualenv 12.0.7
New python executable in /home/volttron/volttron/env/bin/python2.7
Also creating executable in /home/volttron/volttron/env/bin/python
Installing setuptools, pip...done.
Installing required packages
+ easy_install BACpypes>=0.10,<0.11
...
+ pip install --global-option --quiet --install-option --zmq=bundled --no-deps pyzmq>
↪=14.3,<15
...
+ pip install --global-option --quiet --editable ./lib/jsonrpc --editable . --
↪requirement ./requirements.txt
...
Successfully installed Smap-2.0.24c780d avro-1.7.7 configobj-5.0.6 ecdsa-0.13.0
↪flexible-jsonrpc
gevent-1.0.1 greenlet-0.4.5 monotonic-0.1 numpy-1.9.1 pandas-0.15.2 paramiko-1.15.2
pycrypto-2.6.1 pymodbus-1.2.0 pyserial-2.7 python-dateutil-2.4.0 pytz-2014.10.0
↪requests-2.5.3
simplejson-3.6.5 six-1.9.0 twisted-15.0.0 volttron-2.0 wheel-0.24.0 zope.interface-4.
↪1.2

real    9m2.299s
user    7m51.790s
sys     0m14.450s
```

Whew! The build took just over nine minutes on my nearly-4-year-old MacBook Pro running Arch Linux. In case you are wondering about my system's name, as seen in the bash prompt, *inamatus* is Latin for *unloved*. I'll leave it as an exercise for the user to determine why my system is unloved (hint: it has to do with a wonderful fruit with a bite missing from the side).

Anyway, let's have another look at the pip download cache.

```
[volttron@inamatus volttron]$ find ~/.cache/pip -type f
/home/volttron/.cache/pip/http/9/a/b/2/1/
↪9ab21efc4225c8eb9aa41dlc76abef2a53babcefa438a79fa4e981ce
/home/volttron/.cache/pip/http/9/2/6/7/2/
↪92672ab99ac77960252018fbc4f40984eef60ba5588229a729f18f5
/home/volttron/.cache/pip/http/9/e/6/1/9/
↪9e61964f51d8a05a20ecf21eef694877f28cb654a123ce1316ff77e5
/home/volttron/.cache/pip/http/9/7/7/1/a/
↪9771a6b64f3294ac335fdb8574cd3564e21c130924697381d72fd04d
/home/volttron/.cache/pip/http/a/a/7/e/8/
↪aa7e8bc2af1068a43747b0f771b426b7dcf7708283ca3ce3d92a2afc
...
/home/volttron/.cache/pip/http/8/f/9/0/d/
↪8f90d7cf09a2b5380a319b0df8eed268be28d590b6b5f71598a3b56f
/home/volttron/.cache/pip/http/8/d/e/d/a/
↪8deda849bcfd627b8587addf049f79bb333dd8fe1eae1d5053881039
/home/volttron/.cache/pip/http/8/8/7/a/6/
↪887a67fb460d57a10a50deef3658834b9ac01722244315227d334628
/home/volttron/.cache/pip/http/5/5/4/e/2/
↪554e2be8d96625aa74a4e0c4ee4a4b1ca10a442c2877bd3fff96e2a6
/home/volttron/.cache/pip/http/1/d/c/8/3/
↪1dc83c11a861a2bc20d9c0407b41089eba236796ba80c213511f1f74
/home/volttron/.cache/pip/log/debug.log
```

The output is truncated because it was long and boring. The important thing is that it now exists. Next let's remove the virtual environment and rebuild to see what effect the download cache has on our build time.

```
[volttron@inamatus volttron]$ rm -rf env
[volttron@inamatus volttron]$ time python2.7 bootstrap.py
...
real    8m35.387s
user    7m50.770s
sys     0m14.170s
```

Notice that our CPU time was nearly the same, about 8 minutes (user + sys). So the remaining time was likely spent on I/O, which was reduced by about 30 seconds. We need something else to reduce CPU time. Enter `ccache`.

Better

What is `ccache`? According to the official [ccache](#) site,

`ccache` is a compiler cache. It speeds up recompilation by caching the result of previous compilations and detecting when the same compilation is being done again.

Sounds like just the thing we need. `ccache` is already properly configured on my system, it just needs to be placed early in the `PATH` to be found before the official `gcc` compilers.

```
[volttron@inamatus volttron]$ which gcc
/usr/bin/gcc
[volttron@inamatus volttron]$ export PATH=/usr/lib/ccache/bin:$PATH
[volttron@inamatus volttron]$ which gcc
/usr/lib/ccache/bin/gcc
```

Now to prove to ourselves that the cache will be filled during the next run, let's have a look at the cache status.

```
[volttron@inamatus volttron]$ ccache -s
cache directory           /home/volttron/.ccache
primary config            /home/volttron/.ccache/ccache.conf
secondary config         (readonly) /etc/ccache.conf
cache hit (direct)       0
cache hit (preprocessed) 0
cache miss                0
files in cache           0
cache size                0.0 kB
max cache size           5.0 GB
```

The cache is indeed empty.

Nothing up my sleeve... Presto!

```
[volttron@inamatus volttron]$ rm -rf env
[volttron@inamatus volttron]$ time python2.7 bootstrap.py
...

real 6m3.496s
user 4m57.960s
sys 0m10.880s
```

One might expect a ccache build to take slightly longer than the baseline on the first build within a single project. This build completed about two minutes faster. Let's look at the ccache status to discover why.

```
[volttron@inamatus volttron]$ ccache -s
cache directory           /home/volttron/.ccache
primary config            /home/volttron/.ccache/ccache.conf
secondary config         (readonly) /etc/ccache.conf
cache hit (direct)       204
cache hit (preprocessed) 23
cache miss                633
called for link           140
called for preprocessing 95
compile failed            1139
preprocessor error        4
bad compiler arguments    5
autoconf compile/link    103
no input file             19
files in cache            1316
cache size                26.1 MB
max cache size           5.0 GB
```

Ah ha. There were a total of 227 cache hits, meaning that some of the files were identical across all the built packages and the cached version could be used rather than recompiling. Let's see how subsequent builds improve with few cache misses.

```
[volttron@inamatus volttron]$ rm -rf env
[volttron@inamatus volttron]$ time python2.7 bootstrap.py
...

real 3m15.811s
user 2m24.890s
sys 0m7.090s
```

Wow! Now we're cooking with gas. Build times have been cut to nearly 1/3 of our baseline. This ccache status shows only 14 cache misses over our previous run:

```
[volttron@inamatus volttron]$ ccache -s
cache directory                /home/volttron/.ccache
primary config                 /home/volttron/.ccache/ccache.conf
secondary config (readonly)   /etc/ccache.conf
cache hit (direct)            1038
cache hit (preprocessed)      35
cache miss                    647
called for link                280
called for preprocessing      190
compile failed                2278
preprocessor error            8
bad compiler arguments        10
autoconf compile/link        206
no input file                  38
files in cache                1365
cache size                     35.0 MB
max cache size                 5.0 GB
```

So using ccache is a big win. Anyone compiling C or C++ on a Linux system should have ccache enabled. Wait, make that *must*. Go, now, and enable it on your Linux boxen. Or maybe finish reading this and then go do it. But do it!

Best

Now you're thinking "how could it get any better," right? Well, it can. What if those compiled packages only needed to be rebuilt when a new version was required instead of every time they are installed.

When pip installs a package, it downloads the source and executes the packages `setup.py` like so: `python setup.py install`. The install command builds the package and installs it directly into the file system. What if we could package up the build results into an archive and just extract them to the file system when the package is installed. Enter **wheel**.

pip supports the latest Python packaging format known as wheel. Typically this just means that it can install packages in the [wheel format](#). However, if the `wheel` package is installed, pip can also [build wheels](#) from source, executing `python setup.py bdist_wheel`. By default, wheels are placed in the `wheelhouse` directory in the current working directory. But we can alter that location by setting an environment variable (read more on configuring pip [here](#)).

```
[volttron@inamatus volttron]$ export PIP_WHEEL_DIR=$HOME/.cache/pip/wheelhouse
```

We also need to tell pip to look for the wheels, again using an environment variable. The directory needs to exist because while the wheel command will create the directory when creating the packages, pip may try to search the directory first.

```
[volttron@inamatus volttron]$ export PIP_FIND_LINKS=file://$PIP_WHEEL_DIR
[volttron@inamatus volttron]$ mkdir $PIP_WHEEL_DIR
```

So to get this all working, bootstrapping now has to occur in three steps: install the virtual environment, build the wheels, and install the requirements. `bootstrap.py` takes options that control its behavior. The first pass requires the `-o` or `--only-virtenv` option to stop bootstrap after installing the virtual environment and prevent the update stage.

```
[volttron@inamatus volttron]$ rm -rf env
[volttron@inamatus volttron]$ time python2.7 bootstrap.py --only-virtenv
Creating virtual Python environment
Downloading virtualenv DOAP record
```

(continues on next page)

(continued from previous page)

```

Downloading virtualenv 12.0.7
New python executable in /home/volttron/volttron/env/bin/python2.7
Also creating executable in /home/volttron/volttron/env/bin/python
Installing setuptools, pip...done.

real    0m3.866s
user    0m1.480s
sys     0m0.230s

```

The second step requires the `-w` or `--wheel` option to build the wheels. Because the virtual environment already exists, `bootstrap.py` must be called with the virtual environment Python, not the system Python.

```

[volttron@inamatus volttron]$ time env/bin/python bootstrap.py --wheel
Building required packages
+ pip install --global-option --quiet wheel
...
+ pip wheel --global-option --quiet --build-option --zmq=bundled --no-deps pyzmq>=14.
↪3,<15
...
+ pip wheel --global-option --quiet --editable ./lib/jsonrpc --editable . --
↪requirement ./requirements.txt
...
Destination directory: /home/volttron/.cache/pip/wheelhouse
Successfully built numpy pandas gevent monotonic pymodbus simplejson Smap greenlet_
↪pycrypto
twisted pyserial configobj avro zope.interface

real    3m15.431s
user    2m17.980s
sys     0m5.630s

```

It took 3.25 minutes to build the wheels (with `ccache` still enabled). Repeating this command results in nothing new being compiled and takes only 4 seconds. Only new versions of packages meeting the requirements will be built.

```

[volttron@inamatus volttron]$ time env/bin/python bootstrap.py --wheel
Building required packages
...
Skipping numpy, due to already being wheel.
Skipping pandas, due to already being wheel.
Skipping python-dateutil, due to already being wheel.
Skipping requests, due to already being wheel.
Skipping flexible-jsonrpc, due to being editable
Skipping pyzmq, due to already being wheel.
Skipping gevent, due to already being wheel.
Skipping monotonic, due to already being wheel.
Skipping paramiko, due to already being wheel.
Skipping pymodbus, due to already being wheel.
Skipping setuptools, due to already being wheel.
Skipping simplejson, due to already being wheel.
Skipping Smap, due to already being wheel.
Skipping wheel, due to already being wheel.
Skipping volttron, due to being editable
Skipping pytz, due to already being wheel.
Skipping six, due to already being wheel.
Skipping greenlet, due to already being wheel.
Skipping ecdsa, due to already being wheel.

```

(continues on next page)

(continued from previous page)

```

Skipping pycrypto, due to already being wheel.
Skipping pyserial, due to already being wheel.
Skipping twisted, due to already being wheel.
Skipping configobj, due to already being wheel.
Skipping avro, due to already being wheel.
Skipping zope.interface, due to already being wheel.

real    0m3.998s
user    0m3.580s
sys     0m0.360s

```

And let's see what is in the wheelhouse.

```

[volttron@inamatus volttron]$ ls ~/.cache/pip/wheelhouse
Smmap-2.0.24c780d-py2-none-any.whl
Twisted-15.0.0-cp27-none-linux_x86_64.whl
avro-1.7.7-py2-none-any.whl
configobj-5.0.6-py2-none-any.whl
ecdsa-0.13-py2.py3-none-any.whl
gevent-1.0.1-cp27-none-linux_x86_64.whl
greenlet-0.4.5-cp27-none-linux_x86_64.whl
monotonic-0.1-py2-none-any.whl
numpy-1.9.1-cp27-none-linux_x86_64.whl
pandas-0.15.2-cp27-none-linux_x86_64.whl
paramiko-1.15.2-py2.py3-none-any.whl
pycrypto-2.6.1-cp27-none-linux_x86_64.whl
pymodbus-1.2.0-py2-none-any.whl
pyserial-2.7-py2-none-any.whl
python_dateutil-2.4.0-py2.py3-none-any.whl
pytz-2014.10-py2.py3-none-any.whl
pymq-14.5.0-cp27-none-linux_x86_64.whl
requests-2.5.3-py2.py3-none-any.whl
setuptools-12.2-py2.py3-none-any.whl
simplejson-3.6.5-cp27-none-linux_x86_64.whl
six-1.9.0-py2.py3-none-any.whl
wheel-0.24.0-py2.py3-none-any.whl
zope.interface-4.1.2-cp27-none-linux_x86_64.whl

```

Now `bootstrap.py` can be run without options to complete the bootstrap process, again using the virtual environment Python.

```

[volttron@inamatus volttron]$ time env/bin/python bootstrap.py
Installing required packages
+ easy_install BACpypes>=0.10,<0.11
...
+ pip install --global-option --quiet --install-option --zmq=bundled --no-deps pymq>
↳=14.3,<15
...
+ pip install --global-option --quiet --editable ./lib/jsonrpc --editable . --
↳requirement ./requirements.txt
...
Successfully installed Smmap-2.0.24c780d avro-1.7.7 configobj-5.0.6 ecdsa-0.13
↳flexible-jsonrpc
gevent-1.0.1 greenlet-0.4.5 monotonic-0.1 numpy-1.9.1 pandas-0.15.2 paramiko-1.15.2
pycrypto-2.6.1 pymodbus-1.2.0 pyserial-2.7 python-dateutil-2.4.0 pytz-2014.10
↳requests-2.5.3

```

(continues on next page)

(continued from previous page)

```
simplejson-3.6.5 six-1.9.0 twisted-15.0.0 volttron-2.0 zope.interface-4.1.2

real    0m11.137s
user    0m8.930s
sys     0m0.950s
```

Installing from wheels completes in only 11 seconds. And if we blow away the environment and bootstrap again, it takes under 15 seconds.

```
[volttron@inamatus volttron]$ rm -rf env
[volttron@inamatus volttron]$ time python2.7 bootstrap.py
...

real    0m14.644s
user    0m10.380s
sys     0m1.240s
```

Building a clean environment now occurs in less than 15 seconds instead of the 9 minute baseline. That, my friends, is fast.

Why care?

The average VOLTTRON developer probably won't care or see much benefit from the wheel optimization. The typical developer workflow does not include regularly removing the virtual environment and rebuilding. This is, however, very important for continuous integration (CI). With CI, a build server should check out a fresh copy of the source code, build it in a clean environment, and perform unit tests, notifying offending users when their changes break things. Ideally, notification of breakage should happen as soon as possible. We just shaved nearly nine minutes off the turnaround time. It also reduces the load on a shared CI build server, which is nice for everyone.

Taking it further

Two additional use cases present themselves: offline installs and shared builds.

Offline Installs

Let's say we have a system that is not connected to the Internet and, therefore, cannot download packages from PyPi or any other package index. Or perhaps it doesn't have a suitable compiler. Wheels can be built on another *similar*, connected system and transferred by USB drive to the offline system, where they can then be installed. Note that the architecture must be identical and the OS must be very similar between the two systems for this to work.

If the two systems differ too much for a compatible binary build and the offline system has a suitable compiler, then source files can be copied from the pip download cache and transferred from the online system to the offline system for building.

Shared Builds

If many developers are working on the same project, why not share the results of a build with the rest of the team? Here are some ideas to make it work:

- Put wheels on a shared network drive

- Run a private package index server (maybe with `pypiserver`)
- Expose CI built wheels using Apache, Nginx, or `SimpleHTTPServer`

Issues

Here are some of the issues/drawbacks to the methods described above and some possible solutions.

- Configuring pip using environment variables

No worries. Pip uses configuration files too. And a benefit to using them is that it makes all these wheels available to other Python projects you may be working on, and vice versa.

```
# /home/voltrtron/.config/pip/pip.conf
[global]
wheel-dir = /home/voltrtron/.cache/pip/wheelhouse
find-links = file:///home/voltrtron/.cache/pip/wheelhouse
```

Find more on configuring pip [here](#).

- pip does not clean the wheelhouse

This is not a deal-breaker. The wheel directory can just be removed and it will be recreated. Or a script can be used to remove all but the latest versions of packages.

- Requires an additional step or two

That's the price for speed. But it can be mitigated by writing a script or bash alias to perform the steps.

Conclusion

Here is a quick summary of the build times executed above:

Method	Time (minutes)	
	CPU	Total
Each builds on previous		
baseline	8:07	9:02
with download cache	8:05	8:35
ccache, first run	5:09	6:03
ccache, subsequent runs	2:32	3:16
wheel, first run	2:35	3:30
wheel, subsequent runs	0:12	0:15

Not everyone cares about build times, but for those who do, pre-building Python wheels is a great way to improve install times. At a very minimum, every Python developer installing compiled packages will benefit from using ccache.

The techniques used in this document aren't just for VOLTRON, either. They are generally useful for all moderately sized Python projects.

If you haven't installed ccache yet, go do it. There is no excuse.

Acquiring Third Party Agent Code

Third party agents developed from a variety of sources are available from the voltrtron-applications repository (<https://github.com/VOLTRON/voltrtron-applications.git>). The current best practice is to have the main voltrtron and the voltrtron-applications repository within the same common ancestry folder.

```
volttron-repositories/  
|  
|--- volttron/  
|  
|--- volttron-applications/
```

One can clone the latest applications from the repository via the following command:

```
git clone https://github.com/VOLTTRON/volttron-applications.git
```

Required Software: Linux

The following packages will need to be installed if they are not already:

- git
- build-essential
- python-dev
- openssl
- libssl-dev
- libevent-dev

On **Debian-based systems**, these can all be installed with the following command:

```
sudo apt-get update  
sudo apt-get install build-essential python-dev openssl libssl-dev libevent-dev git
```

On **Redhat or CENTOS systems**, these can all be installed with the following command:

```
sudo yum update  
sudo yum install make automake gcc gcc-c++ kernel-devel python-devel openssl openssl-  
↳devel libevent-devel git
```

If you have an agent which requires the pyodbc package, install the following:

- freetds-bin
- unixodbc-dev

```
sudo apt-get install freetds-bin unixodbc-dev
```

Possible issues

The /tmp directory must allow exec. This error could manifest itself during the building of gevent.

```
# Executing mount should have an entry like the following  
mount  
  
tmpfs on /tmp type tmpfs (rw,nosuid,nodev)
```

To change the mount you can use the following code

```
# remount /tmp to allow exec  
sudo mount -o remount,exec /tmp
```

```
# remount /tmp to disallow exec
sudo mount -o remount,noexec /tmp
```

Repository Structure

There are several options for using the VOLTRON code depending on whether you require the most stable version of the code or want the latest updates as they happen. In order of decreasing stability and increasing currency:

For most stable, download the source code for the latest release at: <https://github.com/VOLTRON/voltron/releases>. These are purely source code and are not tied to the git repository. To update them will require downloading the newest source code and re-installing.

The master branch is now the default branch for VOLTRON (meaning this is what you clone if you do not use the “-b” option). This branch will get the latest stable features as they are pushed. 3.x has stopped being updated (except for fixes) and 4.x will track master.

The “develop” branch contains the latest features as they are developed. Once a feature is considered “finished” it is merged back into develop. Develop will be merged into master once it is considered stable and ready for release. This branch can be cloned by those wanting to work from the latest version of the platform but should not be used in deployments.

Features are developed on “feature” branches or developers’ forks of the main repository. It is not recommended to clone these branches except for exploring a new feature.

VOLTRON Restricted Code

Voltron Restricted adds a broader security layer on top of the voltron platform. If you are interested in this package please contact the voltron team at voltron@pnnl.gov.

- NOTE: Once the package is installed all aspects of the package will be enforced. To override the behavior add no-verify, no-mobility, or no-resource-monitor to the *configuration file*.

The Voltron Restricted package contains the following security enhancements:

- The creation and usage of platform specific Certificate Authority (CA) certificates.
- Multi-level signing of agent packages.
- Multi-level verification of signed packages during agent execution.
- Command line and agent based mobility.
- Allows developer to customize an execution contract for required resources on the current and move requested platform.

The following pages describe the functionality exposed by the Volttron Restricted package:

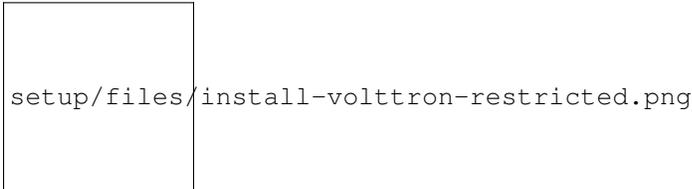
- *Signing and Verification of Agent Packages*
- *Resource Monitor*
- *PingPongAgent*

Note: VOLTRON-Restricted supports VOLTRON 2.x and is being update to VOLTRON 3.5

Installation —

VOLTRON-Restricted requires a software development tool called SWIG ($\geq 2.0.4$). To install VOLTRON-Restricted, follow the steps below. Enter all terminal commands from the VOLTRON directory.

- Extract the VOLTTRON-Restricted code to a new directory. These steps assume the location is `~/volttron-restricted`
- Install VOLTTRON-Restricted dependency:
sudo apt-get install swig
- Activate the VOLTTRON Platform (note the space after the period):
. env/bin/activate
- Install VOLTTRON-Restricted:
pip install -e ~/volttron-restricted



Installing Linux Virtual Machine

VOLTTRON requires a Linux system to run. For Windows users this will require a virtual machine (VM).

Installing Linux Virtual Machine

This section describes the steps necessary to install VOLTTRON using Oracle VirtualBox software. Virtual Box is free and can be downloaded from <https://www.virtualbox.org/wiki/Downloads>.

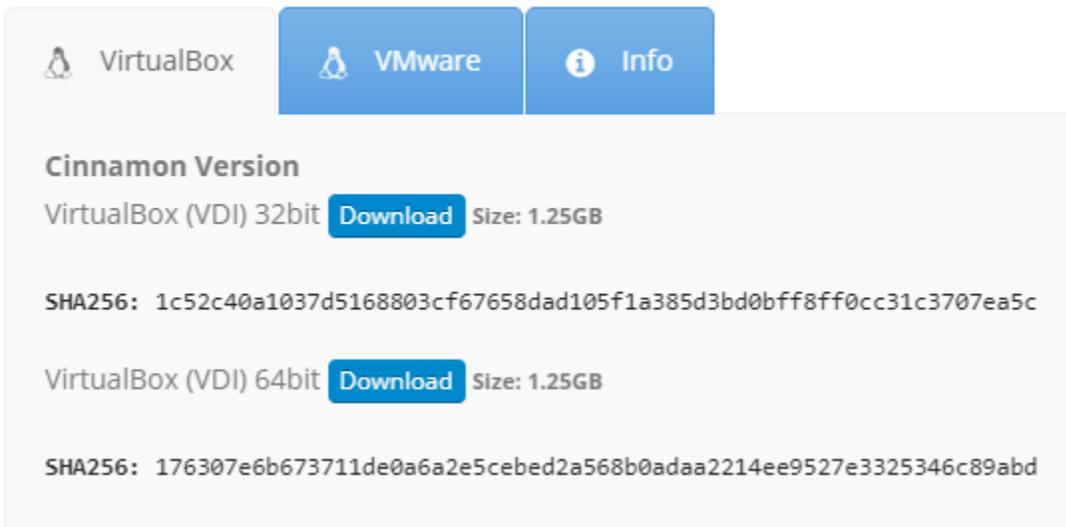
A screenshot of the Oracle VirtualBox website. On the left is a navigation menu with links: About, Screenshots, Downloads, Documentation (with sub-links for End-user docs and Technical docs), Contribute, and Community. The main content area features the VirtualBox logo, a search bar, and links for Login and Preferences. The main heading is "Download VirtualBox", followed by a paragraph: "Here you will find links to VirtualBox binaries and its source code." Below this are sections for "VirtualBox binaries" (with a note about terms and conditions) and "VirtualBox 5.2.12 platform packages" with a bulleted list of links: Windows hosts, OS X hosts, Linux distributions, and Solaris hosts.

After installing VirtualBox download a virtual box appliance from <https://www.osboxes.org/linux-mint/> extract the VDI from the downloaded archive, **or** download a system installation disk. VOLTTRON has been tested using Ubuntu 14.04, 16.04; raspian, debian 8,9; Linux Mint 17, 18; and CentOS 7. However, any modern apt based Linux distribution should work out of the box. Linux Mint 18.3 with the Xfce desktop is used as an example, however platform setup in Ubuntu should be identical.

Note: A 32-bit version of Linux should be used when running VOLTRON on a system with limited hardware (less than 2 GB of RAM).

Adding a VDI Image to VirtualBox Environment

Linux Mint 18.3 Sylvia



VirtualBox VMware Info

Cinnamon Version

VirtualBox (VDI) 32bit [Download](#) Size: 1.25GB

SHA256: 1c52c40a1037d5168803cf67658dad105f1a385d3bd0bff8ff0cc31c3707ea5c

VirtualBox (VDI) 64bit [Download](#) Size: 1.25GB

SHA256: 176307e6b673711de0a6a2e5cebed2a568b0adaa2214ee9527e3325346c89abd

The below info holds the VM's preset username and password.

Linux Mint 18.3 Sylvia



VirtualBox VMware Info

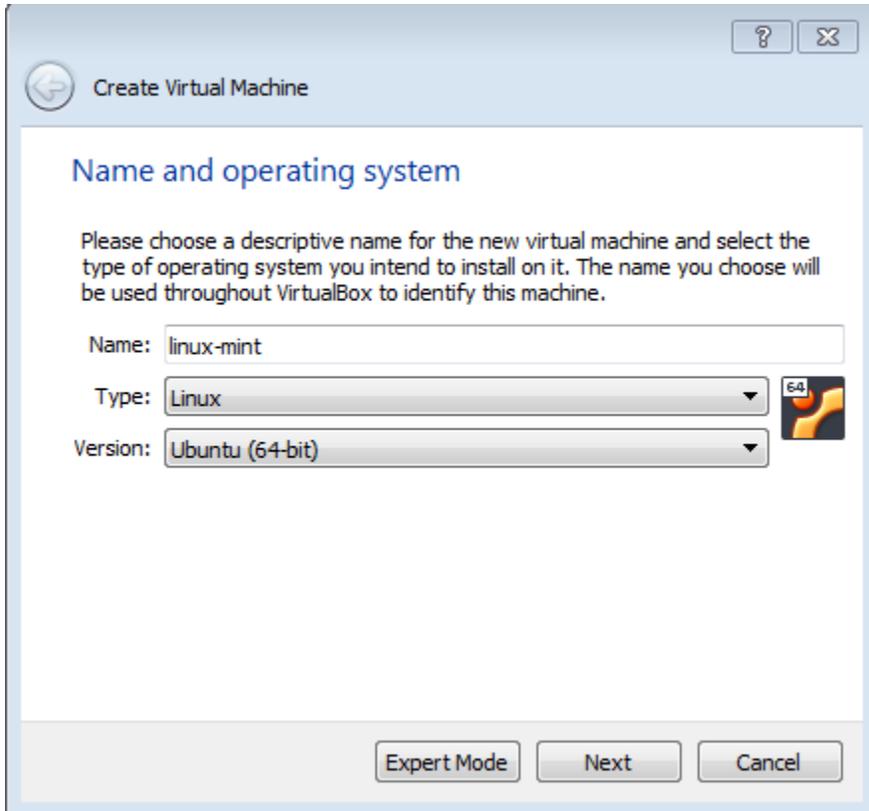
Username: osboxes

Password: osboxes.org

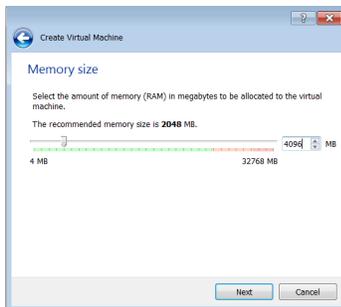
VB Guest Additions & VMware Tools: Not Installed

VMware Compatibility: Version 10+

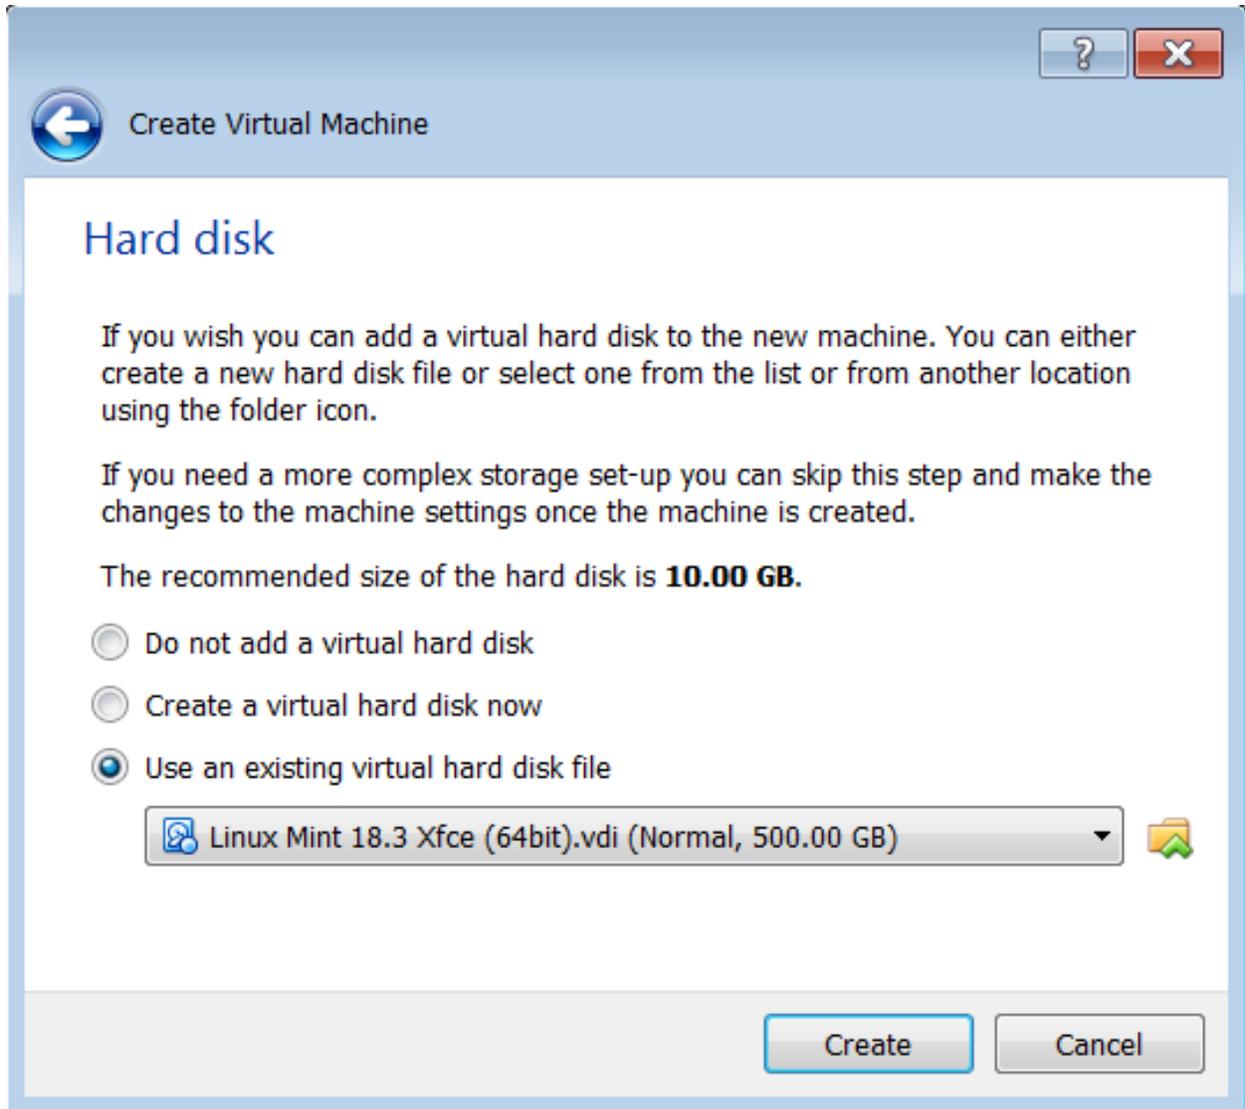
Create a new VirtualBox Image.



Select the amount of RAM for the VM. The recommended minimum is shown in the image below:

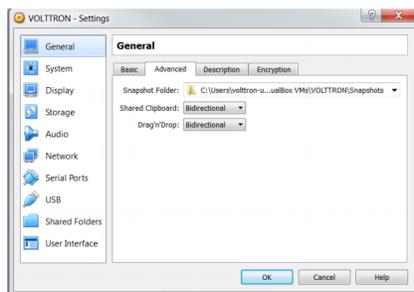


Specify the hard drive image using the extracted VDI file.



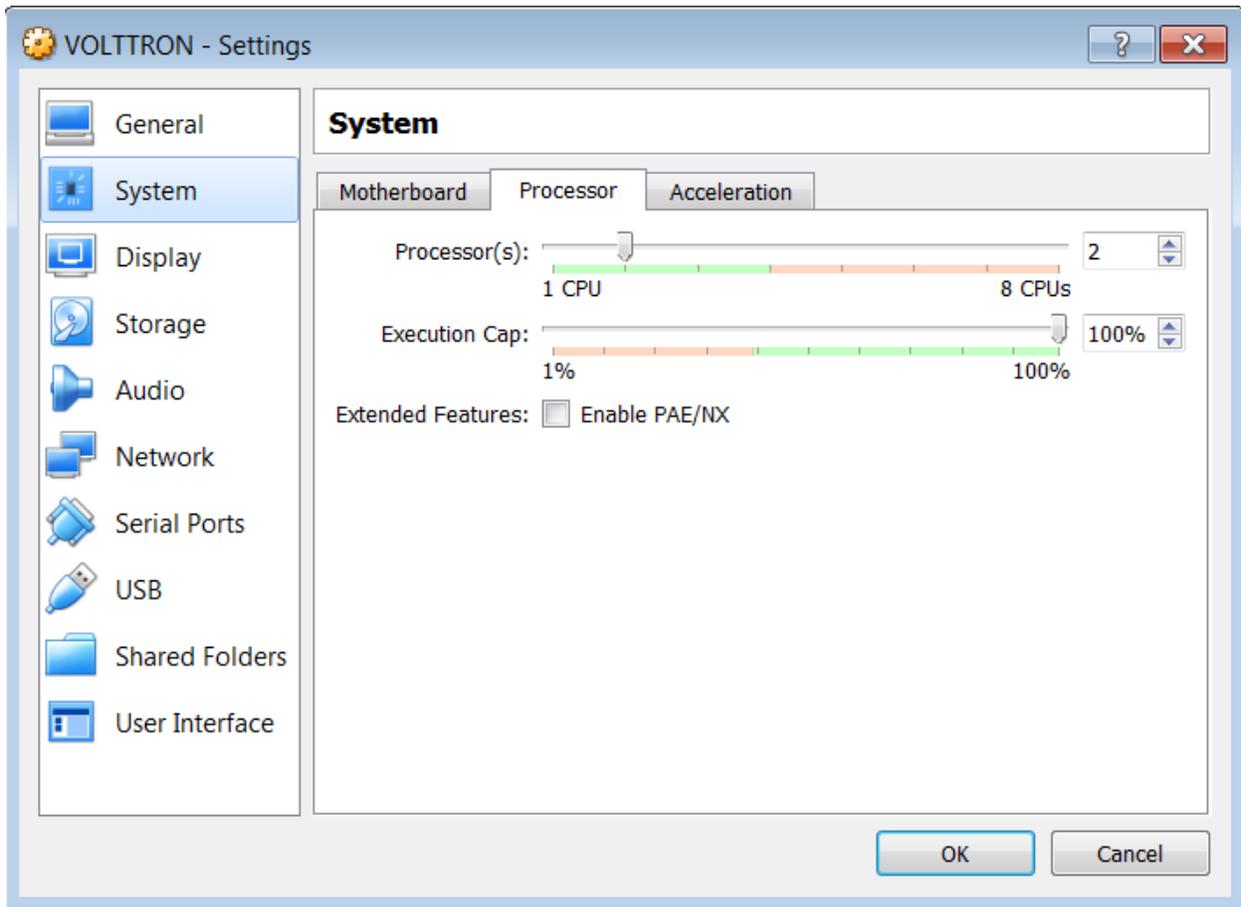
With the newly created VM selected, choose Machine from the VirtualBox menu in the top left corner of the VirtualBox window; from the drop down menu, choose Settings.

To enable bidirectional copy and paste, select the General tab in the VirtualBox Settings. Enable Shared Clipboard and Drag'n'Drop as Bidirectional.



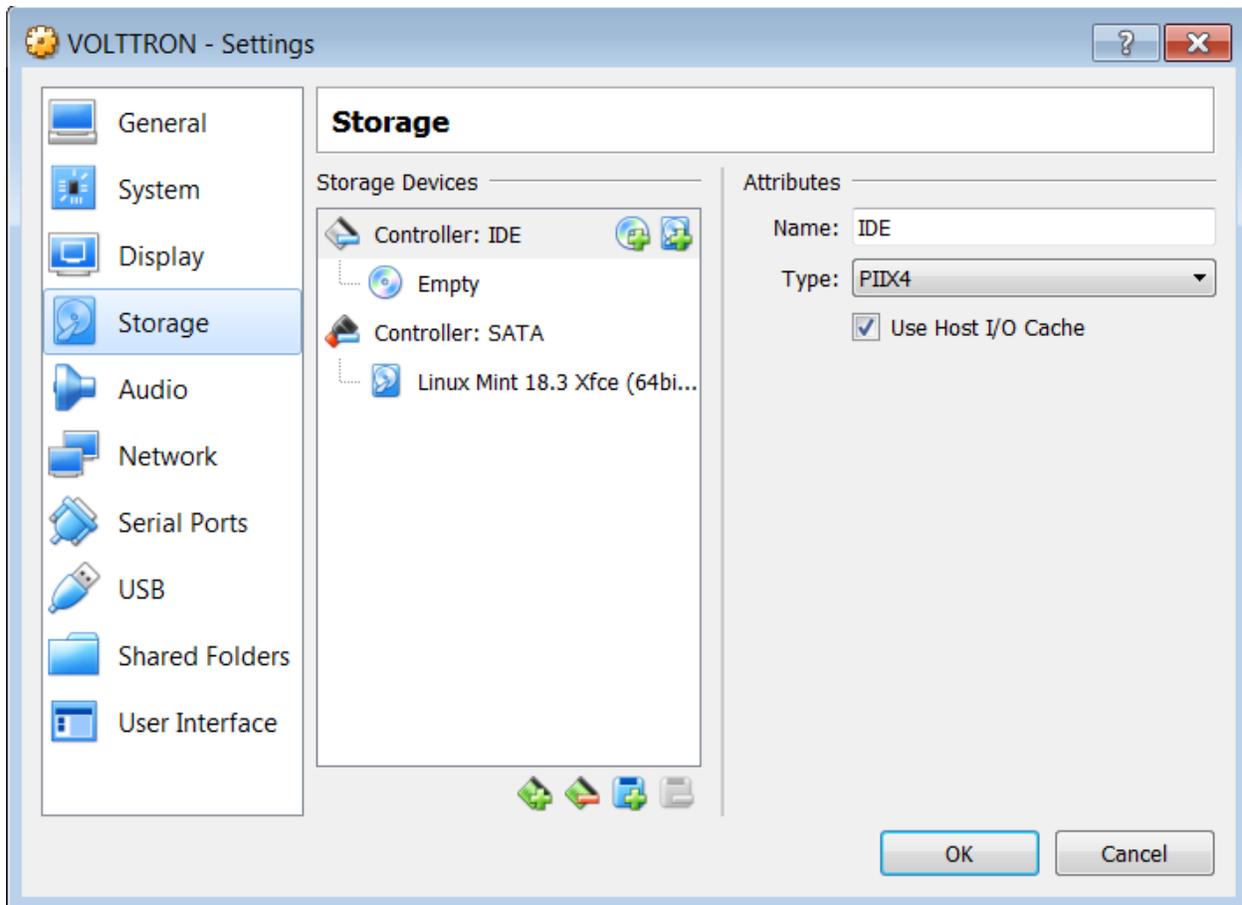
Note: Currently, this feature only works under certain circumstances (e.g. copying / pasting text).

Go to System Settings. In the processor tab, set the number of processors to two.



Go to Storage Settings. Confirm that the Linux Mint VDI is attached to Controller: SATA.

Danger: Do **NOT** mount the Linux Mint iso for Controller: IDE. **Will result in errors.**



Start the machine by saving these changes and clicking the “Start” arrow located on the upper left hand corner of the main VirtualBox window.

Planning a VOLTRON Install

The 3 major installation types for VOLTRON are doing development, doing research using VOLTRON, and collecting and managing physical devices.

Development and Research installation tend to be smaller footprint installations. For development, the data is usually synthetic or copied from another source. The existing documentation covers development installs in significant detail.

Other deployments will have a better installation experience if they consider certain kinds of questions while they plan their installation.

Questions

- Do you want to send commands to the machines ?
- Do you want to store the data centrally ?
- How many machines do you expect to collect data from on each “collector” ?
- How often will the machines collect data ?
- Are all the devices visible to the same network ?
- What types of VOLTRON applications do you want to run ?

Commands

If you wish to send commands to the devices, you will want to install and configure the Volttron Central agent. If you are only using VOLTTRON to securely collect the data, you can turn off the extra agents to reduce the footprint.

Storing Data

VOLTTRON supports multiple historians. MySQL and MongoDB are the most commonly used. As you plan your installation, you should consider how quickly you need access to the data and where. If you are looking at the health and well-being of an entire suite of devices, its likely that you want to do that from a central location. Analytics can be performed at the edge by VOLTTRON applications or can be performed across the data usually from a central data repository. The latency that you can tolerate in your data being available will also determine choices in different agents (ForwardHistorian versus Data Mover)

How Many

The ratio of how many devices-to-collector machine is based on several factors. These include:

- how much memory and network bandwidth the collection machine has. More = More devices
- how fast the local storage is can affect how fast the data cache can be written. Very slow storage devices can fall behind

The second half of the “how many” question is how many collector paltforms are writing to a single VOLTTRON platform to store data - and whether that storage is local, remote, big enough, etc.

If you are storing more than moderate amount of data, you will probably benefit from installing your database on a different machine than your concrete historian machine. Note: This is contra-indicated if you have a slow network connection between you concrete historian and your database machine.

In synthetic testing up to 6 virtual machines hosting 500 devices each (18 points) were easily supported by a single centralized platform writing to a Mongo database - using a high speed network. That central platform experienced very little CPU or memory load when the VOLTTRON Central agent was disabled.

How Often

This question is closely related to the last. A higher sampling frequency will create more data. This wil place more work in the storage phase.

Networks

In many cases, there are constraints on how networks can interact with each other. In many cases, these include security considerations. On some sites, the primary network will be protected from less secure networks and may require different installation considerations. For example, if a data collector machine and the database machine are on the same network with sufficient security, you may choose to have the data collector write directly to the database. If the collector is on an isolated building network then you will likely need to use the ForwardHistorian to bridge the two networks.

Other Considerations

Physical location and maintenance of collector machines must be considered in all live deployments. Although the number of data points may imply a heavy load on a data collection box, the physical constraints may limit the practicality of having more than a single box. The other side of that discussion is deploying many collector boxes may be simpler initially, but may create a maintenance challenge if you don't plan ahead on how you apply patches, etc.

Naming conventions should also be considered. The ability to trace data through the system and identify the collector machine and device can be invaluable in debugging and analysis.

3.4 Developing VOLTRON

3.4.1 Agent Development

Agent Configuration Store Interface

The Agent Configuration Store Subsystem provides an interface for facilitating dynamic configuration via the platform configuration store. It is intended to work alongside the original configuration file to create a backwards compatible system for configuring agents with the bundled configuration file acting as default settings for the agent.

If an Agent Author does not want to take advantage of the platform configuration store they need to make no changes. To completely disable the Agent Configuration Store Subsystem an Agent may pass `enable_store=False` to the `Agent.__init__` method.

The Agent Configuration Store Subsystem caches configurations as the platform sends updates to the agent. Updates from the platform will usually trigger callbacks on the agent.

Agent access to the Configuration Store is managed through the `self.vip.config` object in the Agent class.

The “config” Configuration

The configuration name `config` is considered the canonical name of an Agents main configuration. As such the Agent will always run callbacks for that configuration first at startup and when a change to another configuration triggers any callbacks for `config`.

Configuration Callbacks

Agents may setup callbacks for different configuration events.

The callback method must have the following signature:

```
my_callback(self, config_name, action, contents)
```

Note: The example above is for a class member method, however the method does not need to be a member of the agent class.

- **config_name** - The method to call when a configuration event occurs.
- **action** - The specific configuration event type that triggered the callback. Possible values are “NEW”, “UPDATE”, “DELETE”. See *Configuration Events*

- **contents** - The actual contents of the configuration. Will be a string, list, or dictionary for the actions “NEW” and “UPDATE”. None if the action is “DELETE”.

Note: All callbacks which are connected to the “NEW” event for a configuration will called during agent startup with the initial state of the configuration.

Configuration Events

- **NEW** - This event happens for every existing configuration at Agent startup and whenever a new configuration is added to the Configuration Store.
- **UPDATE** - This event happens every time a configuration is changed.
- **DELETE** - The event happens every time a configuration is removed from the store.

Setting Up a Callback

A callback is setup with the `self.vip.config.subscribe` method.

Note: Subscriptions may be setup at any point in the life cycle of an Agent. Ideally they are setup in `__init__`.

```
subscribe(callback, actions=["NEW", "UPDATE", "DELETE"], pattern="*")
```

- **callback** - The method to call when a configuration event occurs.
- **actions** - The specific configuration event that will trigger the callback. May be a string with the name of a single action or a list of actions.
- **pattern** - The pattern used to match configuration names to trigger the callback.

Configuration Name Pattern Matching

Configuration name matching uses Unix file name matching semantics. Specifically the python module `fnmatch` is used.

Name matching is not case sensitive regardless of the platform VOLTTRON is running on.

For example, the pattern `devices/*` will trigger the supplied callback for any configuration name that starts with `devices/`.

The default pattern matches all configurations.

Getting a Configuration

Once RPC methods are available to an agent (once onstart methods have been called or from any configuration callback) the contents of any configuration may be acquired with the `self.vip.config.get` method.

```
get(config_name="config")
```

If the Configuration Subsystem has not been initialized with the starting values of the agent configuration that will happen in order to satisfy the request. If initialization occurs to satisfy the request callbacks will *not* be called before returning the results.

Typically an Agent will only obtain the contents of a configuration via a callback. This method is included for agents that want to save state in the store and only need to retrieve the contents of a configuration at startup and ignore any changes to the configuration going forward.

Setting a Configuration

Once RPC methods are available to an agent (once onstart methods have been called) the contents of any configuration may be set with the `self.vip.config.set` method.

```
set(config_name, contents, trigger_callback=False, send_update=False)
```

The contents of the configuration may be a string, list, or dictionary.

This method is intended for agents that wish to maintain a copy of their state in the store for retrieval at startup with the `self.vip.config.get` method.

Warning: This method may **not** be called from a configuration callback. The Configuration Subsystem will detect this and raise a `RuntimeError`, even if `trigger_callback` or `send_update` is `False`.

The platform has a locking mechanism to prevent concurrent configuration updates to the Agent. Calling `self.vip.config.set` would cause the Agent and the Platform configuration store for that Agent to deadlock until a timeout occurs.

Optionally an agent may trigger any callbacks by setting `trigger_callback` to `True`. If `trigger_callback` is set to `False` the platform will still send the updated configuration back to the agent. This ensures that a subsequent call to `self.vip.config.get` will still return the correct value. This way the agent's configuration subsystem is kept in sync with the platform's copy of the agent's configuration store at all times.

Optionally the agent may prevent the platform from sending the updated file to the agent by setting `send_update` to `False`. This setting is available strictly for performance tuning.

Warning: This setting will allow the agent's view of the configuration to fall out of sync with the platform. Subsequent calls to `self.vip.config.get` will return an old version of the file if it exists in the agent's view of the configuration store.

This will also affect any configurations that reference the configuration changed with this setting.

Care should be taken to ensure that the configuration is only retrieved at agent startup when using this option.

Setting a Default Configuration

In order to more easily allow agents to use both the Configuration Store while still supporting configuration via the tradition method of a bundled configuration file the `self.vip.config.set_default` method was created.

```
set_default(config_name, contents)
```

Warning: This method may **not** be called once the Agent Configuration Store Subsystem has been initialized. This method should only be called from `__init__` or an `onsetup` method.

The `set_default` method adds a temporary configuration to the Agents Configuration Subsystem. Nothing is sent to the platform. If a configuration with the same name exists in the platform store it will be presented to a callback method in place of the default configuration.

The normal way to use this is to set the contents of the packaged Agent configuration as the default contents for the configuration named `config`. This way the same callback used to process `config` configuration in the Agent will be called when the Configuration Subsystem can be used to process the configuration file packaged with the Agent.

Note: No attempt is made to merge a default configuration with a configuration from the store.

If a configuration is deleted from the store and a default configuration exists with the same name the Agent Configuration Subsystem will call the `UPDATE` callback for that configuration with the contents of the default configuration.

Other Methods

In a well thought out configuration scheme these methods should not be needed but are included for completeness.

List Configurations

A current list of all configurations for the Agent may be called with the `self.vip.config.list` method.

Unsubscribe

All subscriptions can be removed with a call to the `self.vip.config.unsubscribe_all` method.

Delete

A configuration can be deleted with a call to the `self.vip.config.delete` method.

```
delete(config_name, trigger_callback=False)
```

Note: This method may **not** be called from a callback for the same reason as the `self.vip.config.set` method.

Delete Default

A default configuration can be deleted with a call to the `self.vip.config.delete_default` method.

```
delete_default(config_name)
```

Warning: This method may **not** be called once the Agent Configuration Store Subsystem has been initialized. This method should only be called from `__init__` or an `onsetup` method.

Example Agent

The following example shows how to use `set_default` with a basic configuration and how to setup callbacks.

```
def my_agent(config_path, **kwargs):

    config = utils.load_config(config_path) #Now returns {} if config_path does not_
    ↪exist.

    setting1 = config.get("setting1", 42)
    setting2 = config.get("setting2", 2.5)

    return MyAgent(setting1, setting2, **kwargs)

class MyAgent(Agent):
    def __init__(self, setting1=0, setting2=0.0, **kwargs):
        super(MyAgent, self).__init__(**kwargs)

        self.default_config = {"setting1": setting1,
                               "setting2": setting2}

        self.vip.config.set_default("config", self.default_config)
        #Because we have a default config we don't have to worry about "DELETE"
        self.vip.config.subscribe(self.configure_main, actions=["NEW", "UPDATE"],
    ↪pattern="config")
        self.vip.config.subscribe(self.configure_other, actions=["NEW", "UPDATE"],
    ↪pattern="other_config/*")
        self.vip.config.subscribe(self.configure_delete, actions="DELETE", pattern=
    ↪"other_config/*")

    def configure_main(self, config_name, action, contents):
        #Ensure that we use default values from anything missing in the configuration.
        config = self.default_config.copy()
        config.update(contents)

        _log.debug("Configuring MyAgent")

        #Sanity check the types.
        try:
            setting1 = int(config["setting1"])
            setting2 = float(config["setting2"])
        except ValueError as e:
            _log.error("ERROR PROCESSING CONFIGURATION: {}".format(e))
            #TODO: set a health status for the agent
            return

        _log.debug("Using setting1 {}, setting2 {}".format(setting1, setting2))
        #Do something with setting1 and setting2.

    def configure_other(self, config_name, action, contents):
        _log.debug("Configuring From {}".format(config_name))
        #Do something with contents of configuration.

    def configure_delete(self, config_name, action, contents):
        _log.debug("Removing {}".format(config_name))
        #Do something in response to the removed configuration.
```

Agent Creation Walkthrough

The VOLTTRON platform now has utilities to speed the creation and installation of new agents. To use these utilities the VOLTTRON environment must be activated.

From the project directory, activate the VOLTTRON environment with:

```
. env/bin/activate
```

Create Agent Code

Run the following command to start the Agent Creation Wizard:

```
vpkg init TestAgent tester
```

TestAgent is the directory that the agent code will be placed in. The directory must not exist when the command is run.

tester is the name of the agent module created by wizard.

The Wizard will prompt for the following information:

```
Agent version number: [0.1]: 0.5
Agent author: []: VOLTTRON Team
Author's email address: []: volttron@pnnl.gov
Agent homepage: []: https://volttron.org/
Short description of the agent: []: Agent development tutorial.
```

Once the last question is answered the following will print to the console:

```
2018-08-02 12:20:56,604 () volttron.platform.packaging INFO: Creating TestAgent
2018-08-02 12:20:56,604 () volttron.platform.packaging INFO: Creating TestAgent/tester
2018-08-02 12:20:56,604 () volttron.platform.packaging INFO: Creating TestAgent/setup.
↪py
2018-08-02 12:20:56,604 () volttron.platform.packaging INFO: Creating TestAgent/config
2018-08-02 12:20:56,604 () volttron.platform.packaging INFO: Creating TestAgent/
↪tester/agent.py
2018-08-02 12:20:56,604 () volttron.platform.packaging INFO: Creating TestAgent/
↪tester/__init__.py
```

The TestAgent directory is created with the new Agent inside.

Agent Directory

At this point, the contents of the TestAgent directory should look like:

```
TestAgent/
├── setup.py
├── config
├── tester
│   ├── agent.py
│   └── __init__.py
```

Examine the Agent Code

The resulting code is well documented with comments and documentation strings. It gives examples of how to do common tasks in VOLTTRON Agents.

The main agent code is found in *tester/agent.py*

Here we will cover the highlights.

Parse Packaged Configuration and Create Agent Instance

The code to parse a configuration file packaged and installed with the agent is found in the *tester* function:

```
def tester(config_path, **kwargs):
    """Parses the Agent configuration and returns an instance of
    the agent created using that configuration.

    :param config_path: Path to a configuration file.

    :type config_path: str
    :returns: Tester
    :rtype: Tester
    """
    try:
        config = utils.load_config(config_path)
    except StandardError:
        config = {}

    if not config:
        _log.info("Using Agent defaults for starting configuration.")

    setting1 = int(config.get('setting1', 1))
    setting2 = config.get('setting2', "some/random/topic")

    return Tester(setting1,
                  setting2,
                  **kwargs)
```

The configuration is parsed with the *utils.load_config* function and the results are stored in the *config* variable.

An instance of the Agent is created from the parsed values and is returned.

Initialization and Configuration Store Support

The configuration store is a powerful feature introduced in VOLTRON 4. The agent template provides a simple example of setting up default configuration store values and setting up a configuration handler.

```
class Tester(Agent):
    """
    Document agent constructor here.
    """

    def __init__(self, setting1=1, setting2="some/random/topic",
                 **kwargs):
        super(Tester, self).__init__(**kwargs)
        _log.debug("vip_identity: " + self.core.identity)

        self.setting1 = setting1
        self.setting2 = setting2

        self.default_config = {"setting1": setting1,
```

(continues on next page)

(continued from previous page)

```

        "setting2": setting2}

    #Set a default configuration to ensure that self.configure is called
    →immediately to setup
    #the agent.
    self.vip.config.set_default("config", self.default_config)
    #Hook self.configure up to changes to the configuration file "config".
    self.vip.config.subscribe(self.configure, actions=["NEW", "UPDATE"], pattern=
    →"config")

    def configure(self, config_name, action, contents):
        """
        Called after the Agent has connected to the message bus. If a configuration
        →exists at startup
        this will be called before onstart.

        Is called every time the configuration in the store changes.
        """
        config = self.default_config.copy()
        config.update(contents)

        _log.debug("Configuring Agent")

        try:
            setting1 = int(config["setting1"])
            setting2 = str(config["setting2"])
        except ValueError as e:
            _log.error("ERROR PROCESSING CONFIGURATION: {}".format(e))
            return

        self.setting1 = setting1
        self.setting2 = setting2

        self._create_subscriptions(self.setting2)

```

Values in the default config can be built into the agent or come from the packaged configuration file. The subscribe method tells our agent which function to call whenever there is a new or updated config file. For more information on using the configuration store see *Agent Configuration Store*

`_create_subscriptions` (covered in the next section) will use the value in `self.setting2` to create a new subscription.

Setting up a Subscription

The Agent creates a subscription using the value of `self.setting2` in the method `_create_subscription`. The messages for this subscription are handled with the `_handle_publish` method:

```

def _create_subscriptions(self, topic):
    #Unsubscribe from everything.
    self.vip.pubsub.unsubscribe("pubsub", None, None)

    self.vip.pubsub.subscribe(peer='pubsub',
                              prefix=topic,
                              callback=self._handle_publish)

```

(continues on next page)

(continued from previous page)

```
def _handle_publish(self, peer, sender, bus, topic, headers,
                    message):
    pass
```

Agent Lifecycle Events

Methods may be setup to be called at agent startup and shutdown:

```
@Core.receiver("onstart")
def onstart(self, sender, **kwargs):
    """
    This is method is called once the Agent has successfully connected to the
    ↪platform.
    This is a good place to setup subscriptions if they are not dynamic or
    do any other startup activities that require a connection to the message bus.
    Called after any configurations methods that are called at startup.

    Usually not needed if using the configuration store.
    """
    #Example publish to pubsub
    #self.vip.pubsub.publish('pubsub', "some/random/topic", message="HI!")

    #Exmample RPC call
    #self.vip.rpc.call("some_agent", "some_method", arg1, arg2)

@Core.receiver("onstop")
def onstop(self, sender, **kwargs):
    """
    This method is called when the Agent is about to shutdown, but before it
    ↪disconnects from
    the message bus.
    """
    pass
```

As the comment mentions. With the new configuration store feature *onstart* methods are mostly unneeded. However this code does include an example of how to do a Remote Procedure Call to another agent.

Agent Remote Procedure Calls

An agent may receive commands from other agents via a Remote Procedure Call or RPC for short. This is done with the `@RPC.export` decorator:

```
@RPC.export
def rpc_method(self, arg1, arg2, kwarg1=None, kwarg2=None):
    """
    RPC method

    May be called from another agent via self.core.rpc.call """
    return self.setting1 + arg1 - arg2
```

Packaging Configuration

The wizard will automatically create a setup.py file. This file sets up the name, version, required packages, method to execute, etc. for the agent based on your answers to the wizard. The packaging process will also use this information to name the resulting file.

```
from setuptools import setup, find_packages

MAIN_MODULE = 'agent'

# Find the agent package that contains the main module
packages = find_packages('.')
agent_package = 'tester'

# Find the version number from the main module
agent_module = agent_package + '.' + MAIN_MODULE
__temp = __import__(agent_module, globals(), locals(), ['__version__'], -1)
__version__ = __temp.__version__

# Setup
setup(
    name=agent_package + 'agent',
    version=__version__,
    author_email="volttron@pnnl.gov",
    url="https://volttron.org/",
    description="Agent development tutorial.",
    author="VOLTRON Team",
    install_requires=['volttron'],
    packages=packages,
    entry_points={
        'setuptools.installation': [
            'eggsecutable = ' + agent_module + ':main',
        ]
    }
)
```

Launch Configuration

In TestAgent, the wizard will automatically create a file called “config”. It contains configuration information for the agent. This file contains examples every datatype supported by the configuration system:

```
{
  # VOLTRON config files are JSON with support for python style comments.
  "setting1": 2, #Integers
  "setting2": "some/random/topic2", #strings
  "setting3": true, #Booleans: remember that in JSON true and false are not_
↳capitalized.
  "setting4": false,
  "setting5": 5.1, #Floating point numbers.
  "setting6": [1,2,3,4], # Lists
  "setting7": {"setting7a": "a", "setting7b": "b"} #Objects
}
```

Packaging and Installing the Agent

To install the agent the platform must be running. Start the platform with the command:

```
volttron -l volttron.log -vv&
```

Now we must install it into the platform. Use the following command to install it and add a tag for easily referring to the agent. From the project directory, run the following command:

```
python scripts/install-agent.py -s TestAgent/ -c TestAgent/config -t testagent
```

To verify it has been installed, use the following command: `volttron-ctl list`

This will result in output similar to the following:

AGENT	IDENTITY	TAG	STATUS	HEALTH
e testeragent-0.5	testeragent-0.5_1	testagent		

Where the number or letter is the unique portion of the full uuid for the agent. AGENT is the “name” of the agent based on the contents of its class name and the version in its `setup.py`. IDENTITY is the agent’s identity in the platform. This is automatically assigned based on class name and instance number. This agent’s ID is `_1` because it is the first instance. TAG is the name we assigned in the command above. HEALTH is the current health of the agent as reported by the agents health subsystem.

When using lifecycle commands on agents, they can be referred to be UUID (default) or AGENT (name) or TAG.

Testing the Agent

From the Command Line

To test the agent, we will start the platform (if not already running), launch the agent, and check the log file.

- With the VOLTRON environment activated, start the platform by running (if needed):

```
volttron -l volttron.log -vv&
```

- Launch the agent by <uuid> using the result of the list command:

```
vctl start <uuid>
```

- Launch the agent by name with:

```
vctl start --name testeragent-0.1
```

- Launch the agent by tag with:

```
volttron-ctl start --tag testagent
```

- Check that it is *running*:

```
volttron-ctl status
```

- Start the ListenerAgent as in *Building VOLTRON*
- Check the log file for messages indicating the TestAgent is receiving the ListenerAgents messages:

Automated Test cases and documentation

Before contributing a new agent to the VOLTRON source code repository, please consider adding two other essential elements.

1. Integration and unit test cases
2. README file that includes details of pre-requisite software, agent setup details (such as setting up databases, permissions, etc.)and sample configuration

VOLTTRON uses `py.test` as a framework for executing tests. All unit tests should be based on `py.test` framework. `py.test` is not installed with the distribution by default. To install `py.test` and it's dependencies execute the following:

```
python bootstrap.py --testing
```

Note: There are other options for different agent requirements. To see all of the options use:

```
python bootstrap.py --help
```

in the Extra Package Options section.

To run a single test module, use the command

```
pytest <testmodule.py>
```

To run all of the tests in the `voltron` repository execute the following in the root directory using an activated command prompt:

```
./ci-integration/run-tests.sh
```

Agent Development Cheat Sheet

This is a catalogue of features available in `voltron` that are frequently useful in agent development.

Utilities

These functions can be found in the `voltron.platform.agent.utils` module. `logging` also needs to be imported to use the logger.

setup_logging

You'll probably see the following lines near the top of agent files:

```
utils.setup_logging()
_log = logging.getLogger(__name__)
```

This code sets up the logger for this module so it can provide more useful output. In most cases it will be better to use the logger in lieu of simply printing messages with `print`.

load_config

`load_config` does just that. Give it the path to your config file and it will parse the json and return a dictionary.

vip_main

This is the function that is called to start your agent. You'll likely see it in the main methods at the bottom of agents' files. Whatever is passed to it (a class name or a function that returns an instance of your agent) should accept a file path that can be parsed with `load_config`.

Core Agent Functionality

These tools `volttron.platform.vip.agent` module. Try importing

Agent Lifecycle Events

Each agent has four events that are triggered at different stages of its life. These are `onsetup`, `onstart`, `onstop`, and `onfinish`. Registering callbacks to these events are commonplace in agent development, with `onstart` being the most frequently used.

The easiest way to register a callback is with a function decorator:

```
@Core.receiver('onstart')
def function(self, sender, **kwargs):
    function_body
```

Periodic and Scheduled Function Calls

Functions and agent methods can be registered to be called periodically or scheduled to run at a particular time using the `Core.schedule` decorator or by calling an agent's `core.schedule()` method. The latter is especially useful if, for example, a decision needs to be made in an agent's `onstart` method as to whether a call should be scheduled.

```
from volttron.platform.scheduling import cron, periodic

@Core.schedule(t)
def function(self):
    ...

@Core.schedule(periodic(t))
def periodic_function(self):
    ...

@Core.schedule(cron('0 1 * * *'))
def cron_function(self):
    ...
```

or

```
# inside some agent method
self.core.schedule(t, function)
self.core.schedule(periodic(t), periodic_function)
self.core.schedule(cron('0 1 * * *'), cron_function)
```

Subsystem

These features are available to all Agent subclasses. No extra imports are required.

Remote Procedure Calls

Remote Procedure Calls, or RPCs are a powerful way to interact with other agents. To make a function available to call by a remote agent just add the export decorator:

```
@RPC.export
def function(self, ...):
    function_body
```

function can now be called by a remote agent *agent* with

```
# vip_identity is the identity (a string) of the agent
# where function() is defined
agent.vip.rpc.call(vip, 'function').get(timeout=t)
```

Pubsub

Agents can publish and subscribe to topics. Like RPC, pubsub functions can be invoked via decorators or inline through vip. The following function is called whenever the agent sees a message starting with *topic_prefix*.

```
@PubSub.subscribe('pubsub', topic_prefix)
def function(self, peer, sender, bus, topic, headers, message):
    function_body
```

An agent can publish to a topic *topic* with the *self.vip.pubsub.publish* method.

An agent can remove a subscriptions with *self.vip.pubsub.unsubscribe*. Giving None as values for the prefix and callback argument will unsubscribe from everything on that bus. This is handy for subscriptions that must be updated base on a configuration setting.

Configuration Store

Support for the configuration store is done by subscribing to configuration changes with *self.vip.config.subscribe*.

```
self.vip.config.subscribe(self.configure_main, actions=["NEW", "UPDATE"], pattern=
↪"config")
```

See *Agent Configuration Store*

Heartbeat

The heartbeat subsystem provides access to a periodic publish so that others can observe the agent's status. Other agents can subscribe to the *heartbeat* topic to see who is actively publishing to it.

It is turned off by default.

Health

The health subsystem adds extra status information to the an agent's heartbeat. Setting the status will start the heartbeat if it wasn't already.

Agent Skeleton

```

import logging

from volttron.platform.vip.agent import Agent, Core, PubSub, RPC
from volttron.platform.agent import utils

utils.setup_logging()
_log = logging.getLogger(__name__)

class MyAgent(Agent):
    def __init__(self, config_path, **kwargs):
        self.config = utils.load_config(config_path)

    @Core.receiver('onsetup')
    def onsetup(self, sender, **kwargs):
        pass

    @Core.receiver('onstart')
    def onstart(self, sender, **kwargs):
        self.vip.heartbeat.start()

    @Core.receiver('onstop')
    def onstop(self, sender, **kwargs):
        pass

    @Core.receiver('onfinish')
    def onfinish(self, sender, **kwargs):
        pass

    @PubSub.subscribe('pubsub', 'some/topic')
    def on_match(self, peer, sender, bus, topic, headers, message):
        pass

    @RPC.export
    def my_method(self):
        pass

def main():
    utils.vip_main(MyAgent)

if __name__ == '__main__':
    try:
        main()
    except KeyboardInterrupt:
        pass

```

Driver Development

Introduction

All Voltton drivers are implemented through the Master Driver Agent and are technically sub-agents running in the same process as the Master Driver Agent. Each of these driver sub-agents is responsible for creating an interface to a single device. Creating that interface is facilitated by an instance of an interface class. Currently there are two interface classes included: *Modbus* and *BACnet*.

Existing Drivers

In the directory for the Master Driver Agent you'll see a directory called interfaces:



The files `bacnet.py` and `modbus.py` implement the interface class for each respective protocol. (The BACnet interface is mostly just a pass-through to the *BACnet Proxy Agent*, but the Modbus interface is self contained.)

Looking at those two files is a good introduction into how they work.

The file name is used when configuring a driver to determine which interface to use. The name of the interface class in the file must be called `Interface`.

Note: Developing a new driver does not require that your code live with the `MasterDriverAgent` code. You may create the interface file anywhere that you would like and then create a symbolic link to the interface file in the `interfaces` directory. When the `MasterDriverAgent` is packed for distribution a copy of the file represented by the symbolic link is packed into the agent wheel. See *Using Third Party Drivers*

Interface Basics

A complete interface consists of two parts: One or more register classes and the interface class.

Register Class

The Base Interface class uses a Register class to describe the registers of a device to the driver sub-agent. This class is commonly sub-classed to store protocol specific information for the interface class to use. For example, the BACnet interface uses a sub-classed base register to store the instance number, object type, and property name of the point on the device represented by the register class. The Modbus interface uses several different Register classes to deal with the different types of registers on Modbus devices and their different needs.

The register class contains the following attributes:

- **read_only** - True or False
- **register_type** - “bit” or “byte”, used by the driver sub-agent to help deduce some meta data about the point.
- **point_name** - Name of the point on the device. Used by the base interface for reference.
- **units** - units of the value, meta data for the driver
- **description** - meta data for the driver
- **python_type** - python type of the point, used to produce meta data. This must be set explicitly otherwise it default to `int`.

Here is an example of a Registry Class for the BACnet driver:

```
class Register(BaseRegister):
    def __init__(self, instance_number, object_type, property_name, read_only,
↳pointName, units, description = ''):
        super(Register, self).__init__("byte", read_only, pointName, units,
↳description = '')
        self.instance_number = int(instance_number)
        self.object_type = object_type
        self.property = property_name
```

Note that this implementation is incomplete. It does not properly set the `register_type` or `python_type`.

Interface Class

The Interface Class is what is instantiated by the driver sub-agent to do it's work.

`configure(self, config_dict, registry_config_str)`

This method must be implemented by an Interface implementation.

- **config_dict** is a dictionary of key values pairs from the configuration file's "driver_config" section.
- **registry_config_str** is the contents of the "registry_config" entry in the driver configuration file. It is up to the Interface class to parse this file according to the needs of the driver.

Here is an example taken from the BACnet driver:

```
def configure(self, config_dict, registry_config_str):
    self.parse_config(registry_config_str) #Parse the configuration string.
    self.target_address = config_dict["device_address"]
    self.proxy_address = config_dict.get("proxy_address", "platform.bacnet_proxy")
    self.ping_target(self.target_address) #Establish routing to the device if needed.
```

And here is the `parse_config` method (See *BACnet Registry Configuration*):

```
def parse_config(self, config_string):
    if config_string is None:
        return

    f = StringIO(config_string) #Python's CSV file parser wants a file like object.

    configDict = DictReader(f) #Parse the CVS file contents.

    for regDef in configDict:
        #Skip lines that have no address yet.
        if not regDef['Point Name']:
            continue

        io_type = regDef['BACnet Object Type']
        read_only = regDef['Writable'].lower() != 'true'
        point_name = regDef['Volttron Point Name']
        index = int(regDef['Index'])
        description = regDef['Notes']
        units = regDef['Units']
        property_name = regDef['Property']
```

(continues on next page)

(continued from previous page)

```

register = Register(index,
                    io_type,
                    property_name,
                    read_only,
                    point_name,
                    units,
                    description = description)

self.insert_register(register)

```

Once a register is created it must be added with the `insert_register` method.

get_point(self, point_name)

This method must be implemented by an Interface implementation.

Gets the value of a point from a device and returns it.

Here is a simple example from the BACnet driver. In this case it only has to pass the work on to the BACnet Proxy Agent for handling.

```

def get_point(self, point_name):
    register = self.get_register_by_name(point_name)
    point_map = {point_name:[register.object_type,
                             register.instance_number,
                             register.property]}
    result = self.vip.rpc.call(self.proxy_address, 'read_properties',
                              self.target_address, point_map).get()
    return result[point_name]

```

Failure should be indicated by a useful exception being raised. (In this case the we just leave the Exception raised by the BACnet proxy un-handled. This could be improved with better handling when register that does not exist is requested.)

The Register instance for the point can be retrieved with `self.get_register_by_name(point_name)`

set_point(self, point_name, value)

This method must be implemented by an Interface implementation.

Sets the value of a point on a device and ideally returns the actual value set if different.

Here is a simple example from the BACnet driver. In this case it only has to pass the work on to the BACnet Proxy Agent for handling.

```

def set_point(self, point_name, value):
    register = self.get_register_by_name(point_name)
    if register.read_only:
        raise IOError("Trying to write to a point configured read only: "+point_name)
    args = [self.target_address, value,
            register.object_type,
            register.instance_number,
            register.property]
    result = self.vip.rpc.call(self.proxy_address, 'write_property', *args).get()
    return result

```

Failure to raise a useful exception being raised. (In this case the we just leave the Exception raised by the BACnet proxy un-handled unless the point is read only.)

scrape_all(self)

This method must be implemented by an Interface implementation.

This must return a dictionary mapping point names to values for ALL registers.

Here is a simple example from the BACnet driver. In this case it only has to pass the work on to the BACnet Proxy Agent for handling.

```
def scrape_all(self):
    point_map = {}
    read_registers = self.get_registers_by_type("byte", True)
    write_registers = self.get_registers_by_type("byte", False)
    for register in read_registers + write_registers:
        point_map[register.point_name] = [register.object_type,
                                         register.instance_number,
                                         register.property]

    result = self.vip.rpc.call(self.proxy_address, 'read_properties',
                              self.target_address, point_map).get()

    return result
```

self.get_registers_by_type allows you to get lists of registers by their type and if they are read only. (As BACnet currently only uses “byte”, “bit” is ignored.) As the procedure for handling all the different types in BACnet is the same we can bundle them all up into a single request from the proxy.

In the Modbus protocol the distinction is important and so each category must be handled differently.

Developing Historian Agents

VOLTRON provides a convenient base class for developing new historian agents. The base class automatically subscribes to all pertinent topics, cache published data to disk until it is successfully recorded to a historian, create the public facing interface for querying results, and spells out a simple interface for concrete implementation to meet to make a working Historian Agent. The VOLTRON provides support for several historians without modification. Please use one of these if it fits your project criteria, otherwise continue reading.

The base class also breaks data to publish into reasonably sized chunks before handing it off to the concrete implementation for publication. The size of the chunk is configurable.

The base class sets up a separate thread for publication. This way if publication code needs to block for a long period of time (up to 10s of seconds) this will no disrupt the collection of data from the bus or the functioning of the agent itself.

BaseHistorian

All Historians must inherit from the BaseHistorian class in `volttron.platform.agent.base_historian` and implement the following methods:

publish_to_historian(self, to_publish_list)

This method is called by the BaseHistorian class when it has received data from the message bus to be published. to_publish_list is a list of records to publish in the form

```
[
  {
    '_id': 1,
    'timestamp': timestamp,
    'source': 'scrape',
    'topic': 'campus/building/unit/point',
    'value': 90,
    'meta': {'units': 'F'}
  }
  ...
]
```

- `_id` - ID of record. All IDs in the list are unique. This is used for internal record tracking.
- `timestamp` - Python datetime object of the time data was published at timezone UTC
- `source` - Source of the data. Can be scrape, analysis, log, or actuator.
- `topic` - Topic data was published on. Prefix's such as "device" are dropped.
- `value` - Value of the data. Can be any type.
- `meta` - Metadata for the value. Some sources will omit this entirely.

For each item in the list the concrete implementation should attempt to publish (or discard if non-publishable) every item in the list. Publication should be batched if possible. For every successfully published record and every record that is to be discarded because it is non-publishable the agent must call `report_handled` on those records. Records that should be published but were not for whatever reason require no action. Future calls to `publish_to_historian` will include these unpublished records. `publish_to_historian` is always called with the oldest unhandled records. This allows the historian to no lose data due to lost connections or other problems.

As a convenience `report_all_handled` can be called if all of the items in `published_list` were successfully handled.

query_topic_list(self)

Must return a list of all unique topics published.

query_historian(self, topic, start=None, end=None, skip=0, count=None, order=None)

This function must return the results of a query in the form:

```
{"values": [(timestamp1: value1), (timestamp2: value2), ...],
 "metadata": {"key1": value1, "key2": value2, ...}}
```

metadata is not required (The caller will normalize this to {} for you if you leave it out)

- `topic` - the topic the user is querying for.
- `start` - datetime of the start of the query. None for the beginning of time.
- `end` - datetime of the end of of the query. None for the end of time.

- skip - skip this number of results (for pagination)
- count - return at maximum this number of results (for pagination)
- order - “FIRST_TO_LAST” for ascending time stamps, “LAST_TO_FIRST” for descending time stamps.

historian_setup(self)

Implementing this is optional. This function is run on the same thread as the rest of the concrete implementation at startup. It is meant for connection setup.

Developing Market Agents

VOLTRON provides a convenient base class for developing new market agents. The base class automatically subscribes to all pertinent topics, and spells out a simple interface for concrete implementation to make a working Market Agent.

Markets are implemented by the Market Service Agent which is a core service agent. The Market Service Agent publishes information on several topics to which the base agent automatically subscribes. The base agent also provides all the methods you will need to interact with the Market Service Agent to implmnet your market transactions.

MarketAgent

All Market Agents must inherit from the MarketAgent class in `volttron.platform.agent.base_market_agent` and call the following method:

```
self.join_market(market_name, buyer_seller, reservation_callback, offer_callback,
↳aggregate_callback, price_callback, error_callback)
```

This method causes the market agent to join a single market. If the agent wishes to participate in several markets it may be called once for each market. The first argument is the name of the market to join and this name must be unique across the entire volttron instance because all markets are implmented by a single market service agent for each volttron instance. The second argument describes the role that this agent wished to play in this market. The value is imported as:

```
from volttron.platform.agent.base_market_agent.buy_sell import BUYER, SELLER
```

Arguments 3-7 are callback methods that the agent may implement as needed for the agent’s participation in the market.

The Reservation Callback

```
reservation_callback(self, timestamp, market_name, buyer_seller)
```

This method is called when it is time to reserve a slot in the market for the current market cycle. If this callback is not registered a slot is reserved for every market cycle. If this callback is registered it is called for each market cycle and returns True if a reservation is wanted and False if a reservation is not wanted. The name of the market and the roll being played are provided so that a single callback can handle several markets. If the agent joins three markets with the same reservation callback routine it will be called three times with the appropriate market name and buyer/seller role for each call. The MeterAgent example illustrates the use of this of this method and how to determine whether to make an offer when the reservation is refused. A market will only exist if there are reservations for at least one buyer or one seller. If the market fails to achieve the minimum participation the error callback will be called. If only buyers or only sellers make reservations any offers will be rejected with the reason that the market has not formed.

The Offer Callback

```
offer_callback(self, timestamp, market_name, buyer_seller)
```

If the agent has made a reservation for the market and a callback has been registered this callback is called. If the agent wishes to make an offer at this time the market agent computes either a supply or a demand curve as appropriate and offers the curve to the market service by calling the `make_offer` method. The name of the market and the roll being played are provided so that a single callback can handle several markets. For each market joined either an offer callback, an aggregate callback, or a cleared price callback is required.

The Aggregate Callback

```
aggregate_callback(self, timestamp, market_name, buyer_seller, aggregate_curve)
```

When a market has received all its buy offers it calculates an aggregate demand curve. When the market receives all of its sell offers it calculates an aggregate supply curve. This callback delivers the aggregate curve to the market agent whenever the appropriate curve becomes available. If the market agent wants to use this opportunity to make an offer on this or another market it would do that using the `make_offer` method. If the aggregate demand curve is received, obviously you could only make a supply offer on this market. If the aggregate supply curve is received, obviously you could only make a demand offer on this market. You can of course use this information to make an offer on another market. The example AHUAgent does this. The name of the market and the roll being played are provided so that a single callback can handle several markets. For each market joined either an offer callback, an aggregate callback, or a cleared price callback is required.

The Price Callback

```
price_callback(self, timestamp, market_name, buyer_seller, price, quantity)
```

This callback is called when the market clears. If the market agent wants to use this opportunity to make an offer on this or another market it would do that using the `make_offer` method. Once the market has cleared you can't make an offer on that market. You can of course use this information to make an offer on another market. The example AHUAgent does this. The name of the market and the roll being played are provided so that a single callback can handle several markets. For each market joined either an offer callback, an aggregate callback, or a cleared price callback is required.

The Error Callback

```
error_callback(self, timestamp, market_name, buyer_seller, error_code, error_message,   
↳ aux)
```

This callback is called when an error occurs isn't in response to an RPC call. The error codes are documented in:

```
from volttron.platform.agent.base_market_agent.error_codes import NOT_FORMED, SHORT_   
↳ OFFERS, BAD_STATE, NO_INTERSECT
```

- NOT_FORMED - If a market fails to form this will be called at the offer time.
- SHORT_OFFERS - If the market doesn't receive all its offers this will be called while clearing the market.
- BAD_STATE - This indicates a bad state transition while clearing the market and should never happen, but may be called while clearing the market.

- `NO_INTERSECT` - If the market fails to clear this would be called while clearing the market and an auxillary array will be included. The auxillary array contains comparisons between the supply max, supply min, demand max and demand min. They allow the market client to make determinations about why the curves did not intersect that may be useful.

The error callback is optional, but highly recommended.

Agent Development in Eclipse

The Eclipse IDE (integrated development environment), while not required for agent development, can be a powerful developmental tool. Download the IDE from the following links. Choose a download mirror closest to your location.

For 32-bit machines

For 64-bit machines

To go to the main Eclipse webpage, go to <http://eclipse.org/>

Installing Eclipse

To install Eclipse, enter the following commands in a terminal:

1. Install Eclipse dependency:

```
# apt-get install openjdk-7-jdk
```

2. After downloading the eclipse archive file, move the package to the opt directory (enter this command from a terminal in the directory where eclipse was downloaded):

```
$ tar -xvf eclipse-java-mars-R-linux-gtk-x86_64.tar.gz
# mv eclipse /opt/
```

- For 32-bit machines, replace “`gtk-x86_64`” with “`linux-gtk`” in the previous command.

3. Create desktop shortcut:

```
# touch /usr/share/applications/eclipse.desktop
# nano /usr/share/applications/eclipse.desktop
```

Enter the following text, as shown in Figure 1, and save the file. To avoid typos, copy and paste the following:

```
[Desktop Entry]
Name=Eclipse
Type=Application
Exec=/opt/eclipse/eclipse
Terminal=false
Icon=/opt/eclipse/icon.xpm
Comment=Integrated Development Environment
NoDisplay=false
Categories=Development;IDE
Name[en]=eclipse
```

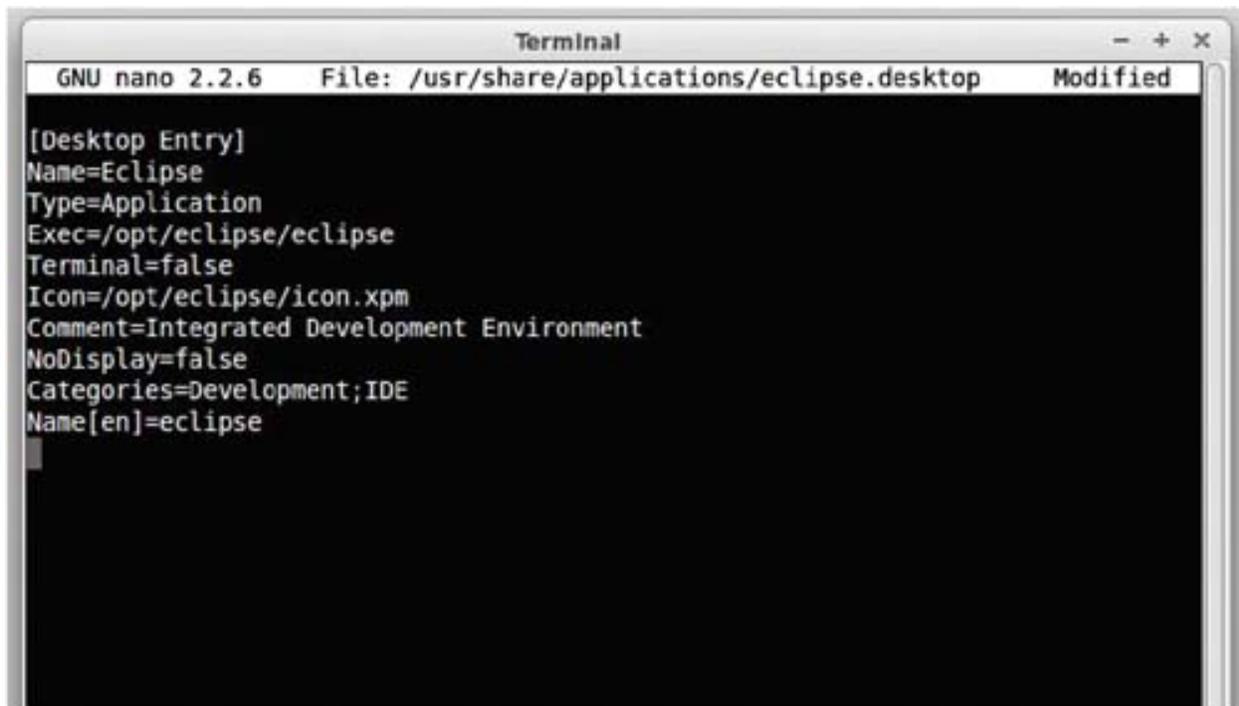


Figure 1. Eclipse Desktop File

4. Copy the shortcut to the desktop:

```
$ cp /usr/share/applications/eclipse.desktop ~/Desktop/
```

Eclipse is now installed and ready to use.

Installing Pydev and EGit Eclipse Plug-ins

The transactional network code is stored in a Git repository. A plug-in is available for Eclipse that makes development more convenient (note: you must have Git installed on the system and have built the project).

1. Select Help. Select Install New Software (Figure 2).

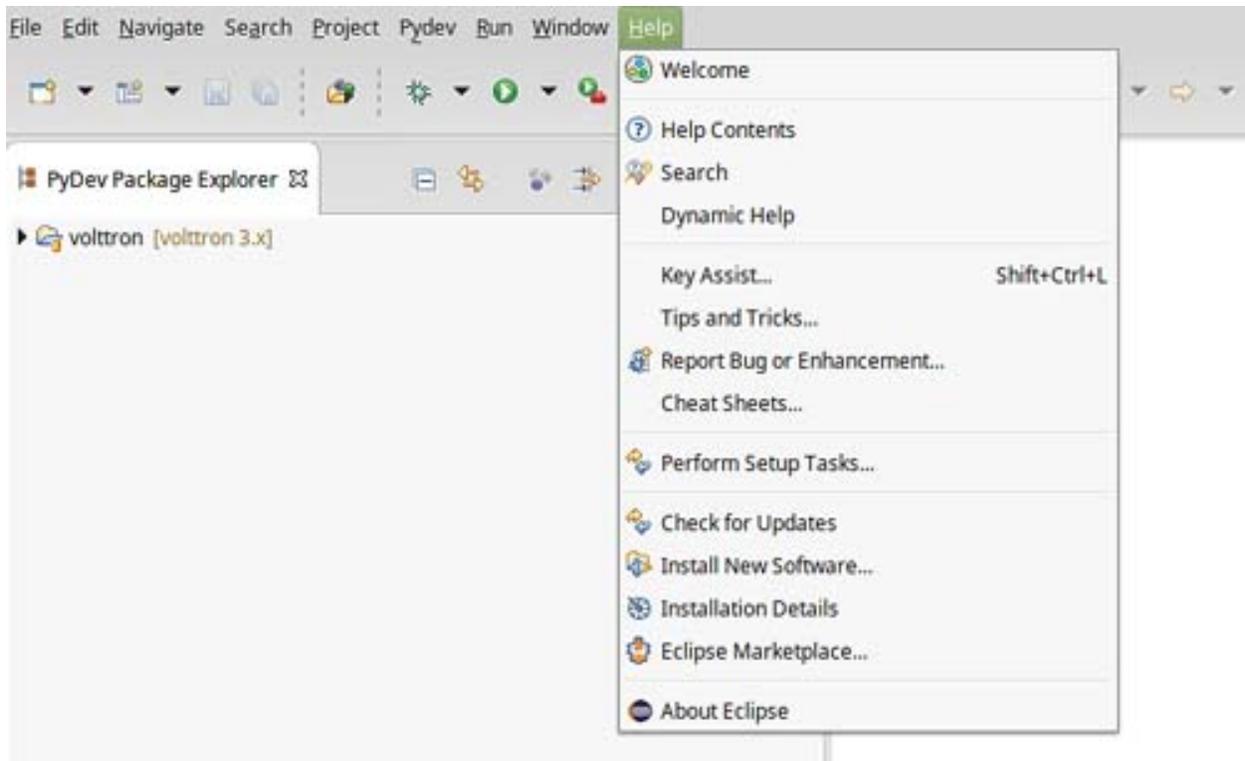


Figure 2. Installing Eclipse EGit Plugin

2. Click the Add button (Figure 3).

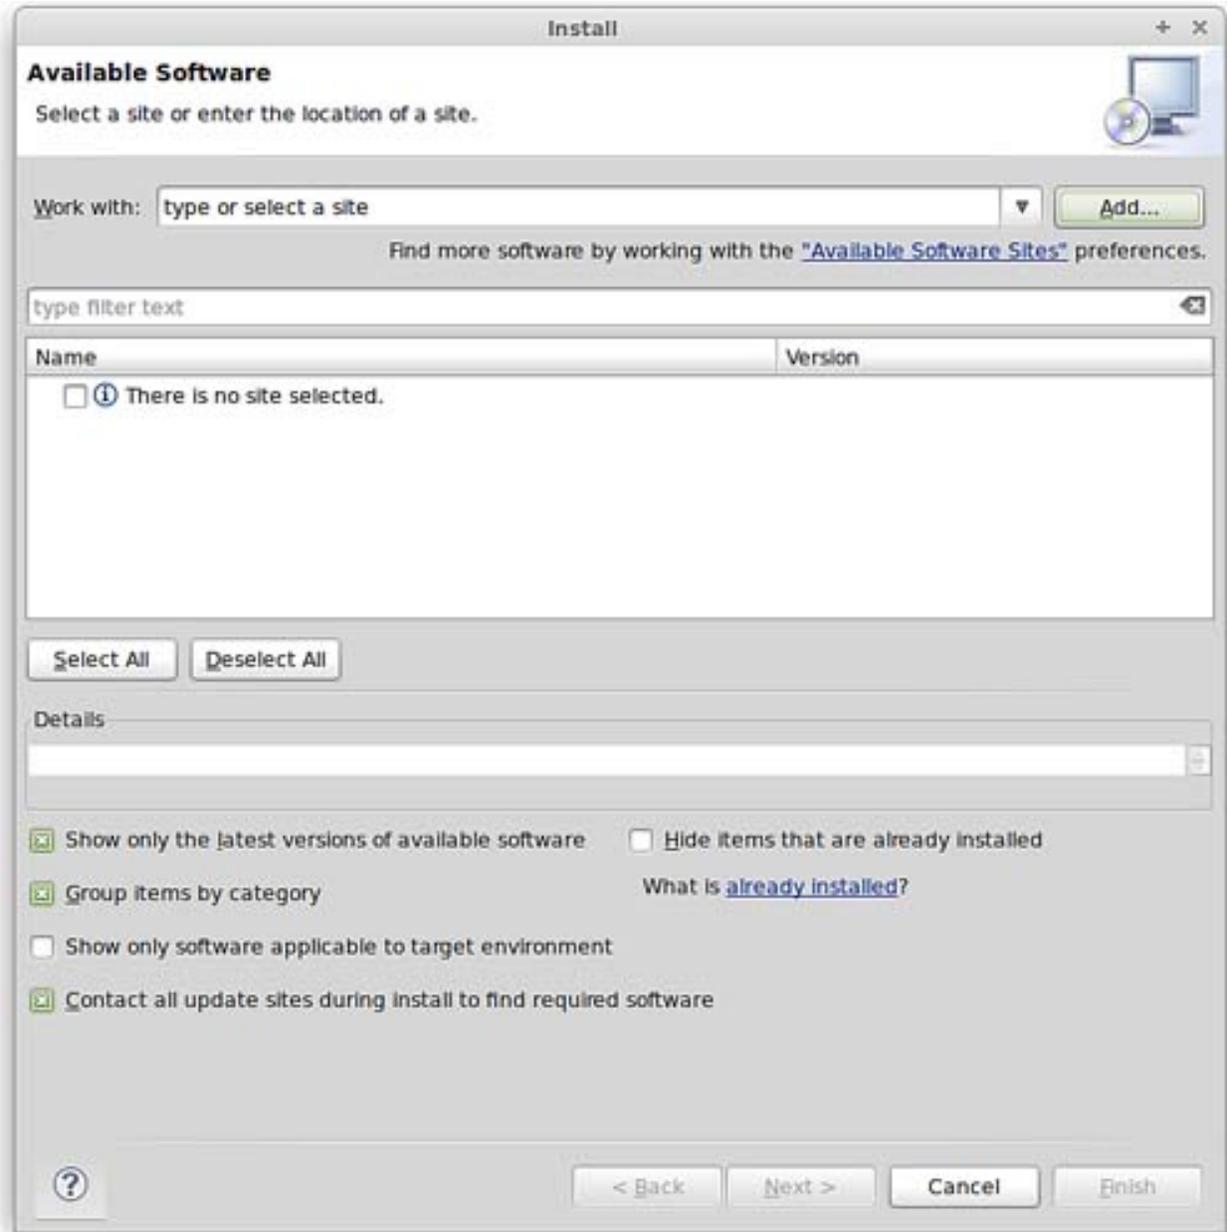


Figure 3. Installing Eclipse EGit Plugin (continued)

3. As shown in Figure 4, enter the following information:

- For name use: EGit
- For location: <http://download.eclipse.org/egit/updates>

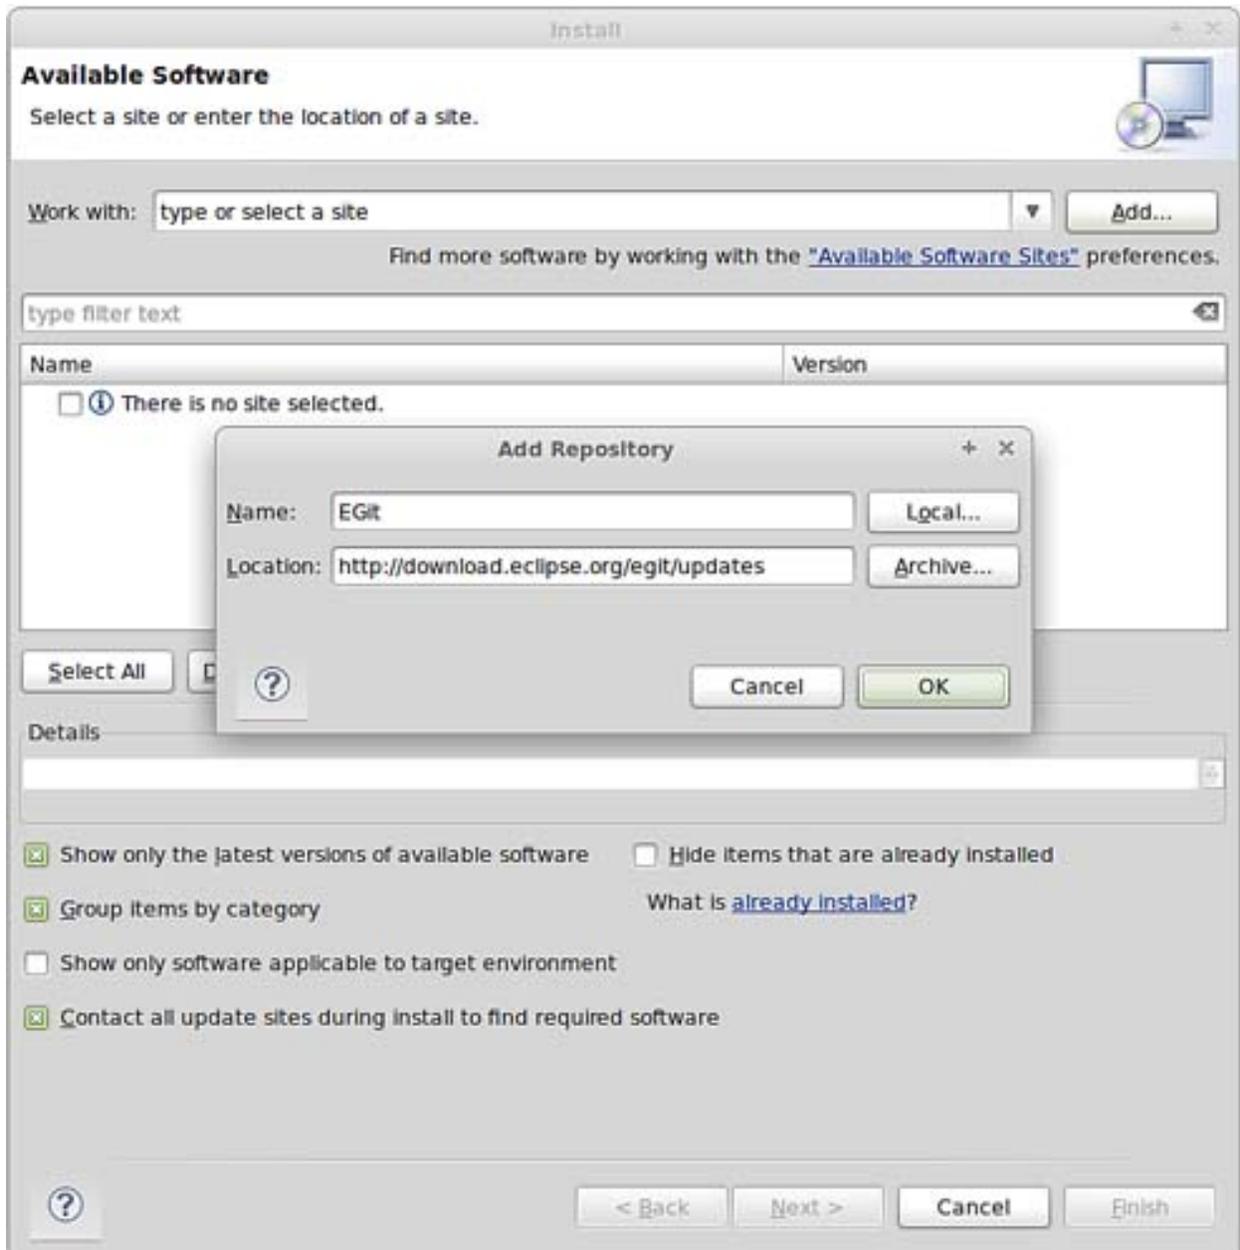


Figure 4. Installing Eclipse Egit Plugin (continued)

4. After clicking OK, check the Select All button.
5. Click through Next > Agree to Terms > Finish. Allow Eclipse to restart.
6. After installing Eclipse, you must add the PyDev plug-in to the environment.

In Eclipse:

- Select Help and select Install New Software.
- Click the Add button.
- **As shown in Figure 5, enter the following information:**
 - For name use: PyDev

- For location: <http://pydev.org/updates>
- Click OK.

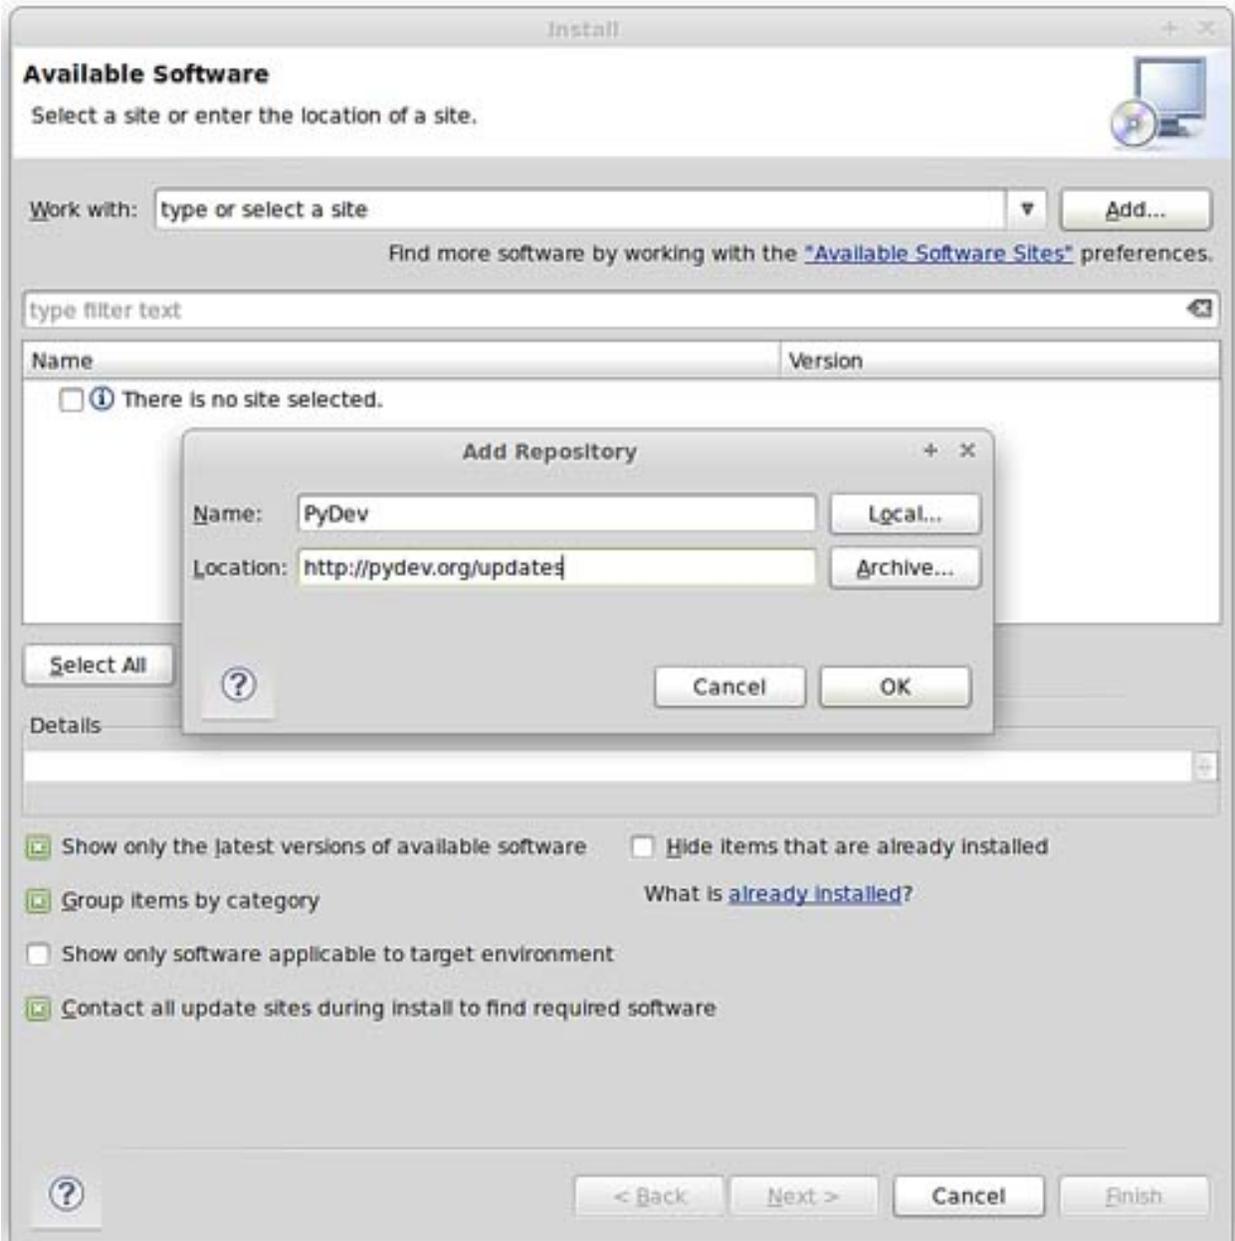


Figure 5. Installing Eclipse PyDev Plugin

7. Check the box for PyDev.
8. Click through Next > Agree to Terms > Finish. Allow Eclipse to restart.

Checkout VOLTRON Project

VOLTRON can be imported into Eclipse from an existing VOLTRON project (VOLTRON was previously checked out from GitHub) or a new download from GitHub.

Import VOLTRON into Eclipse from an Existing Local Repository (Previously Downloaded VOLTRON Project)

To import an existing VOLTRON project into Eclipse, complete the following steps:

1. Select File and select Import (Figure 6).

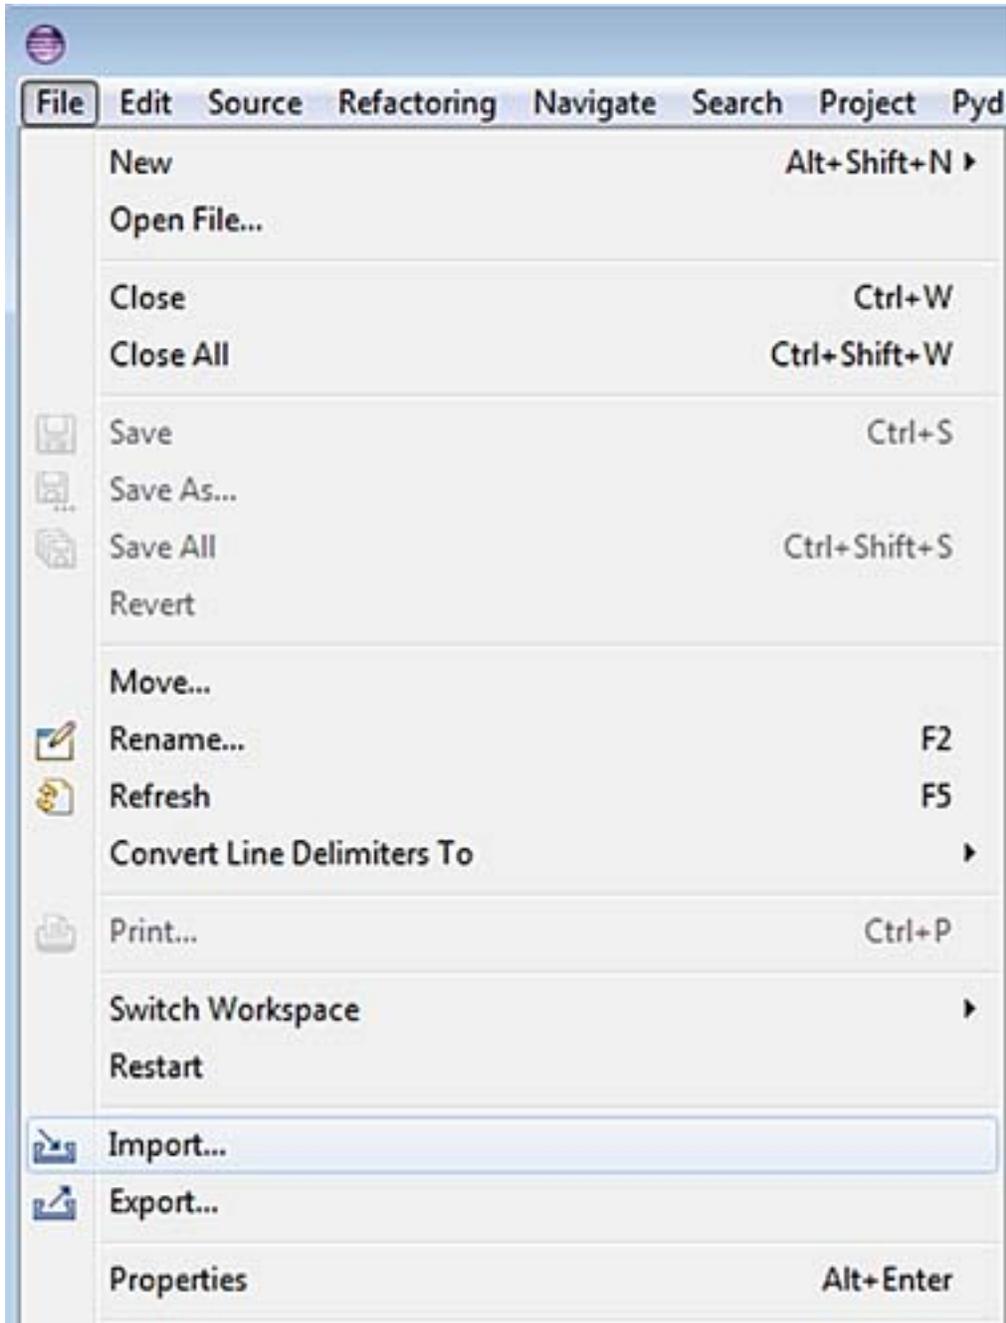


Figure 6. Checking VOLTRON with Eclipse from Local Source

2. Select Git. Select Projects from Git. Click the Next button (Figure 7).

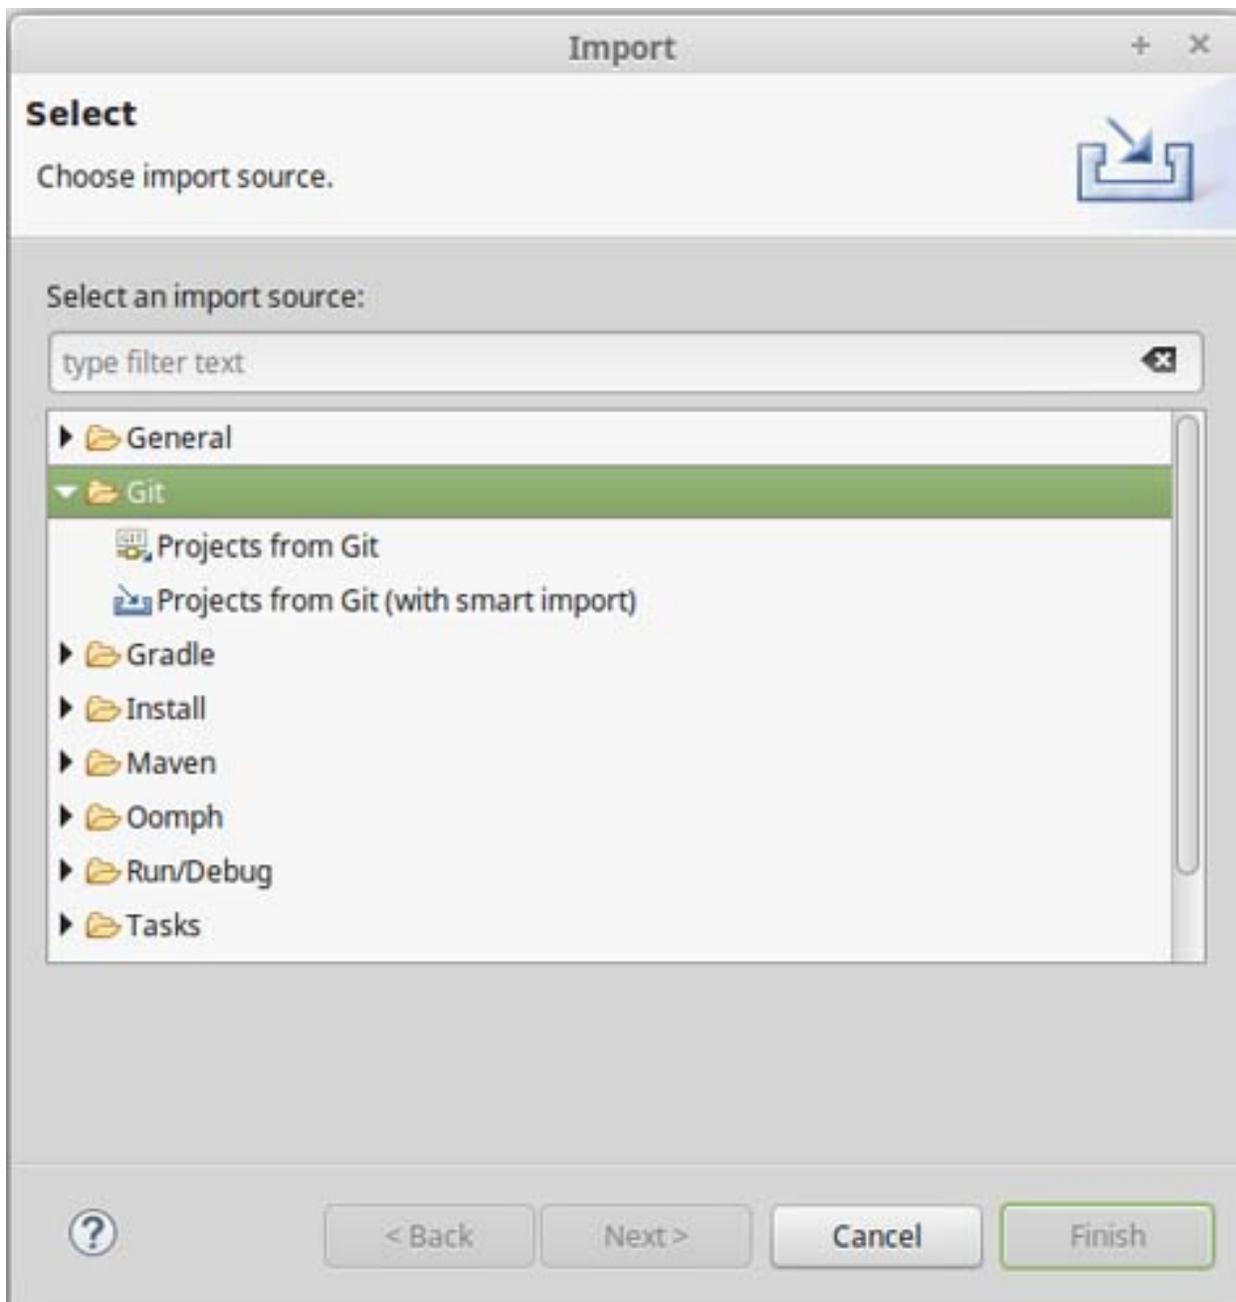


Figure 7. Checking VOLTRON with Eclipse from Local Source (continued)

3. Select Existing local repository and click the Next button (Figure 8).



Figure 8. Checking VOLTRON with Eclipse from Local Source (continued)

4. Select Add (Figure 9).

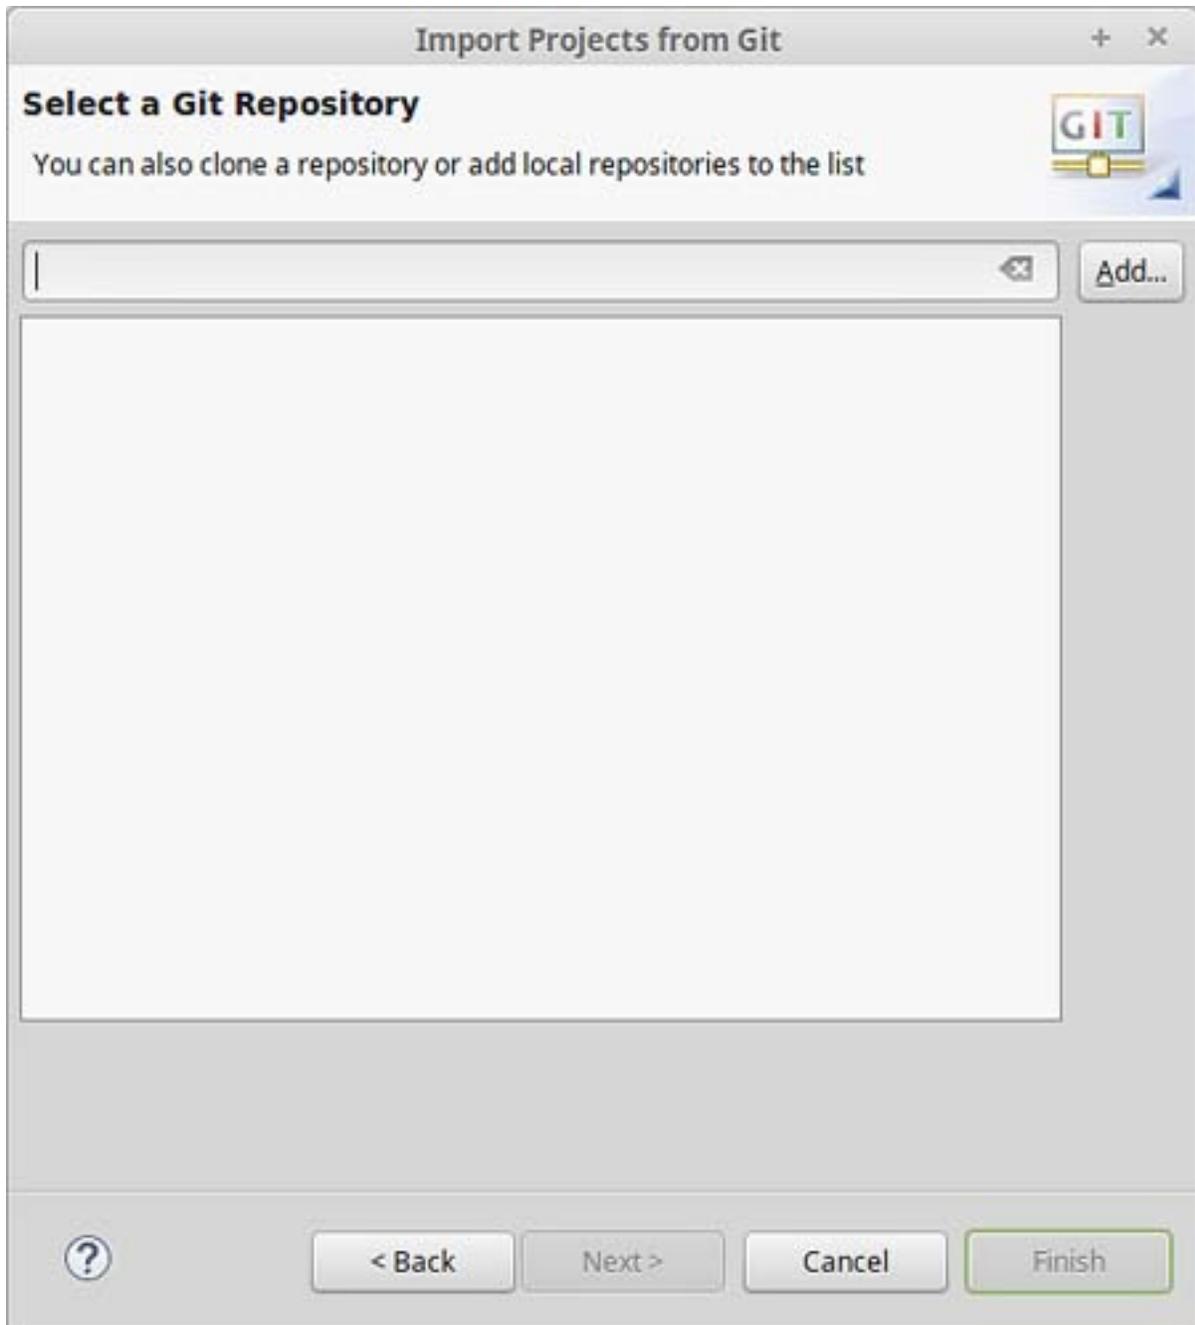


Figure 9. Checking VOLTRON with Eclipse from Local Source (continued)

5. Select Browse. Navigate to the top-level base VOLTRON directory. Select OK (Figure 10).

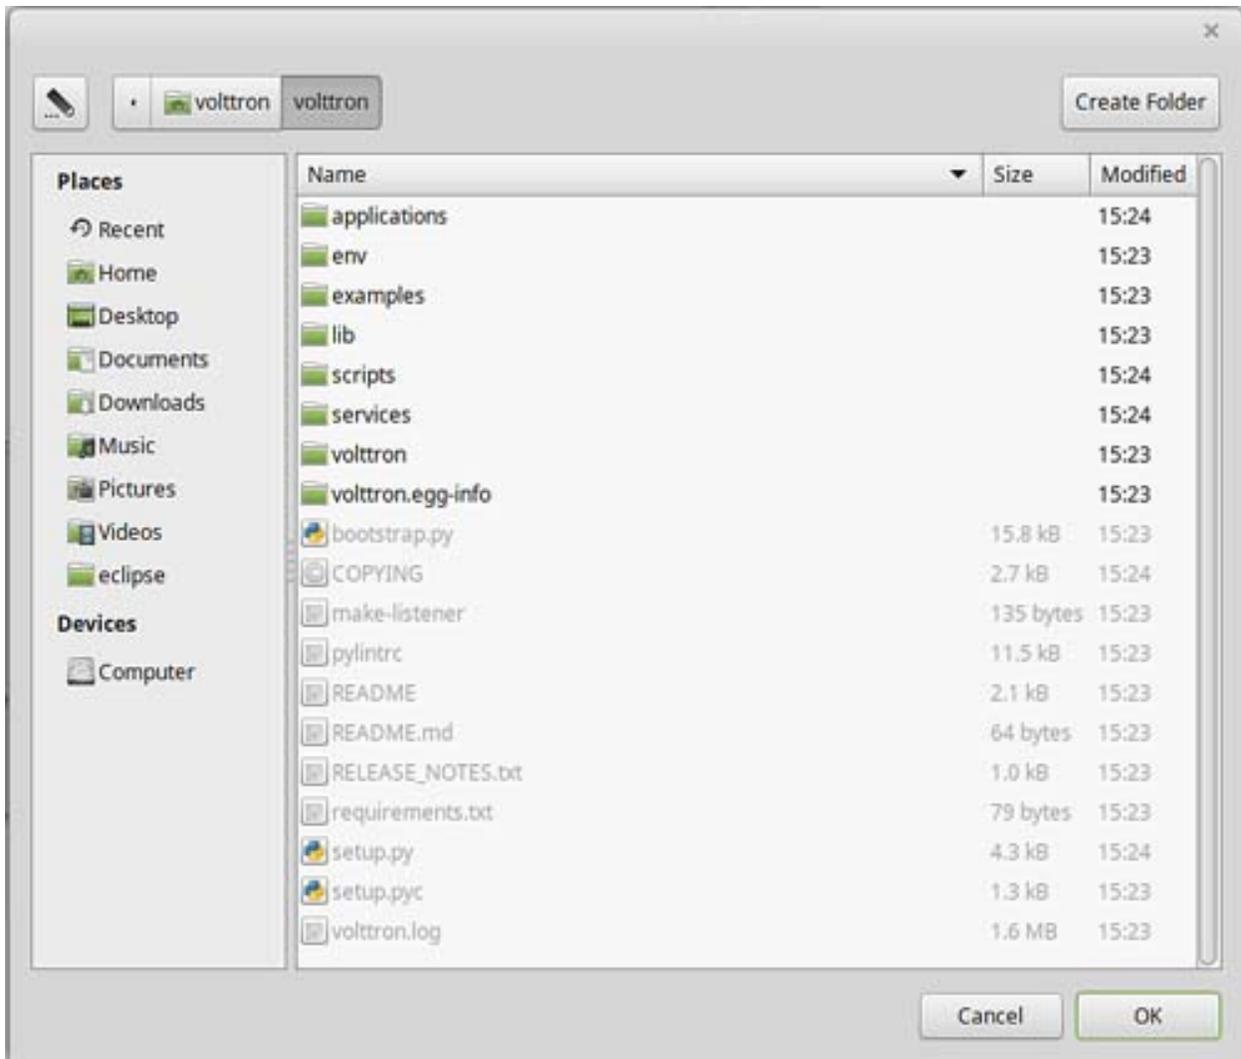


Figure 10. Checking Out VOLTTRON with Eclipse from Local Source (continued)

6. Click Finish (Figure 11).

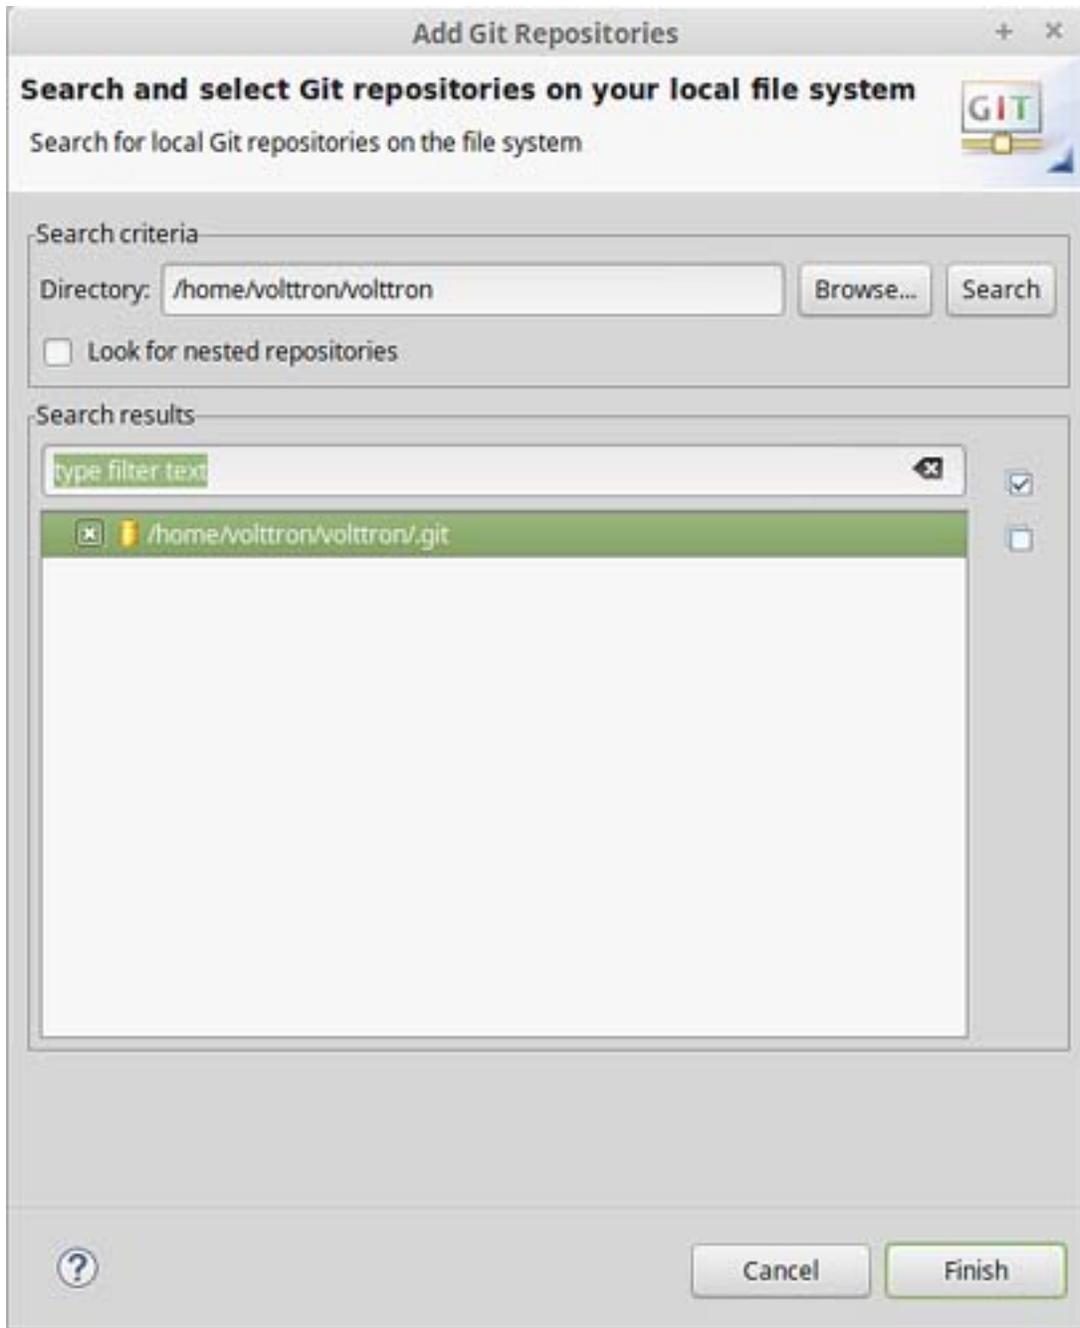


Figure 11. Checking Out VOLTTRON with Eclipse from Local Source (continued)

7. Click Next (Figure 12).

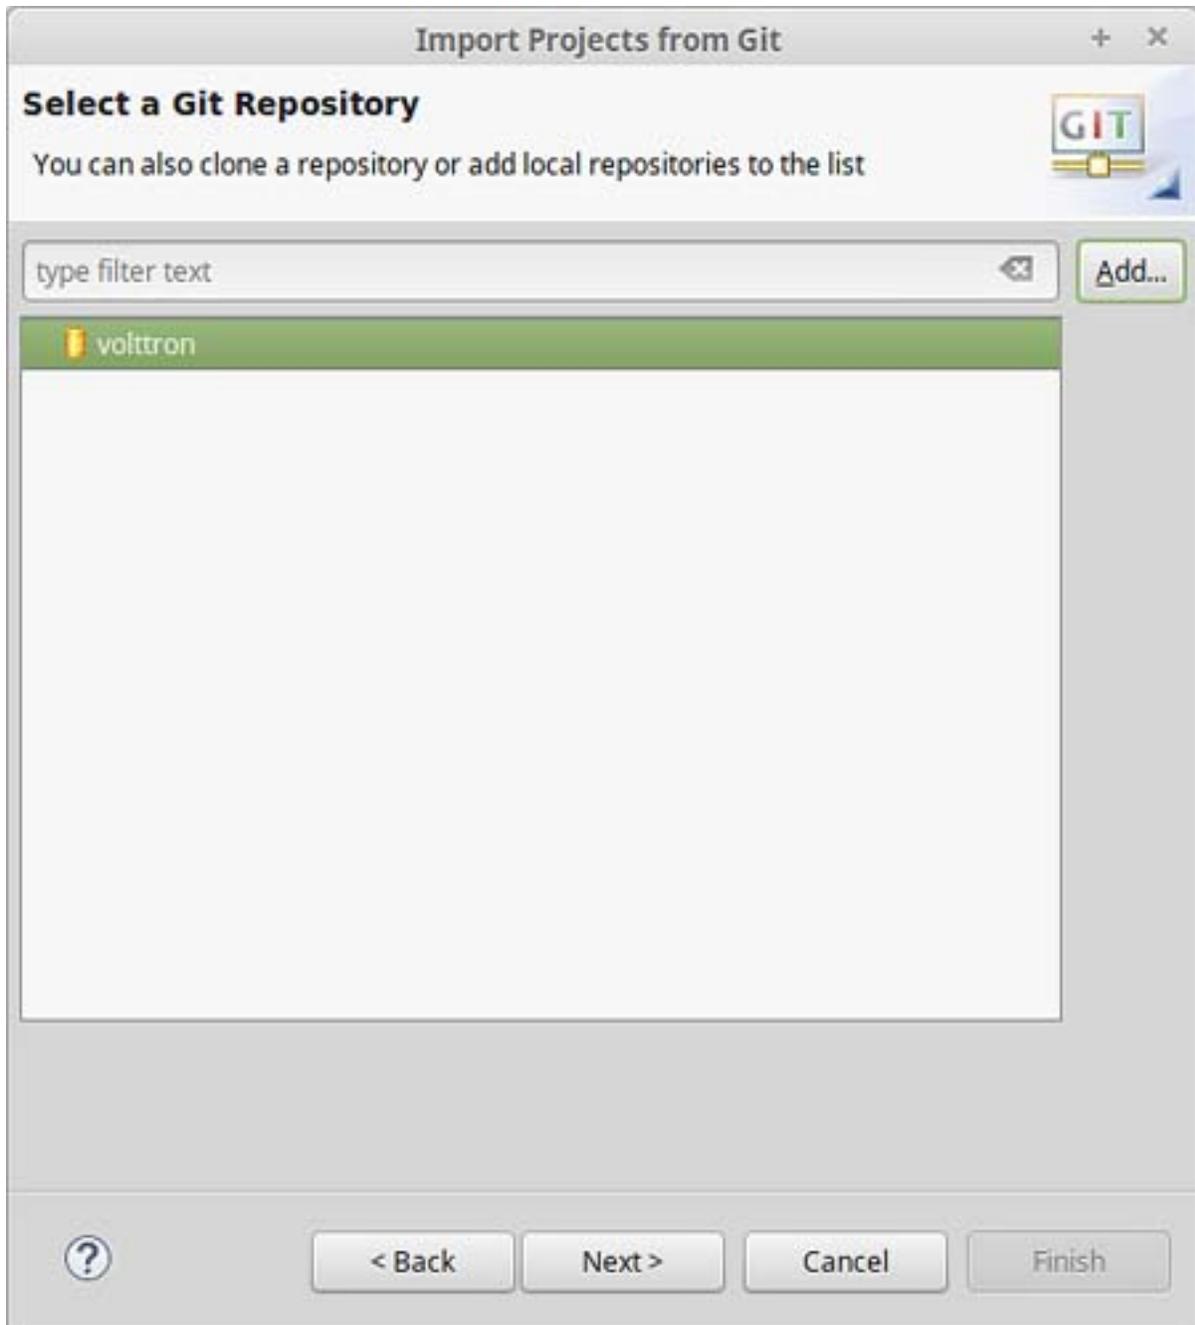


Figure 12. Checking Out VOLTRON with Eclipse from Local Source (continued)

8. Select Import as general project. Click Next. Click Finish (Figure 13). The project will be imported into the workspace.

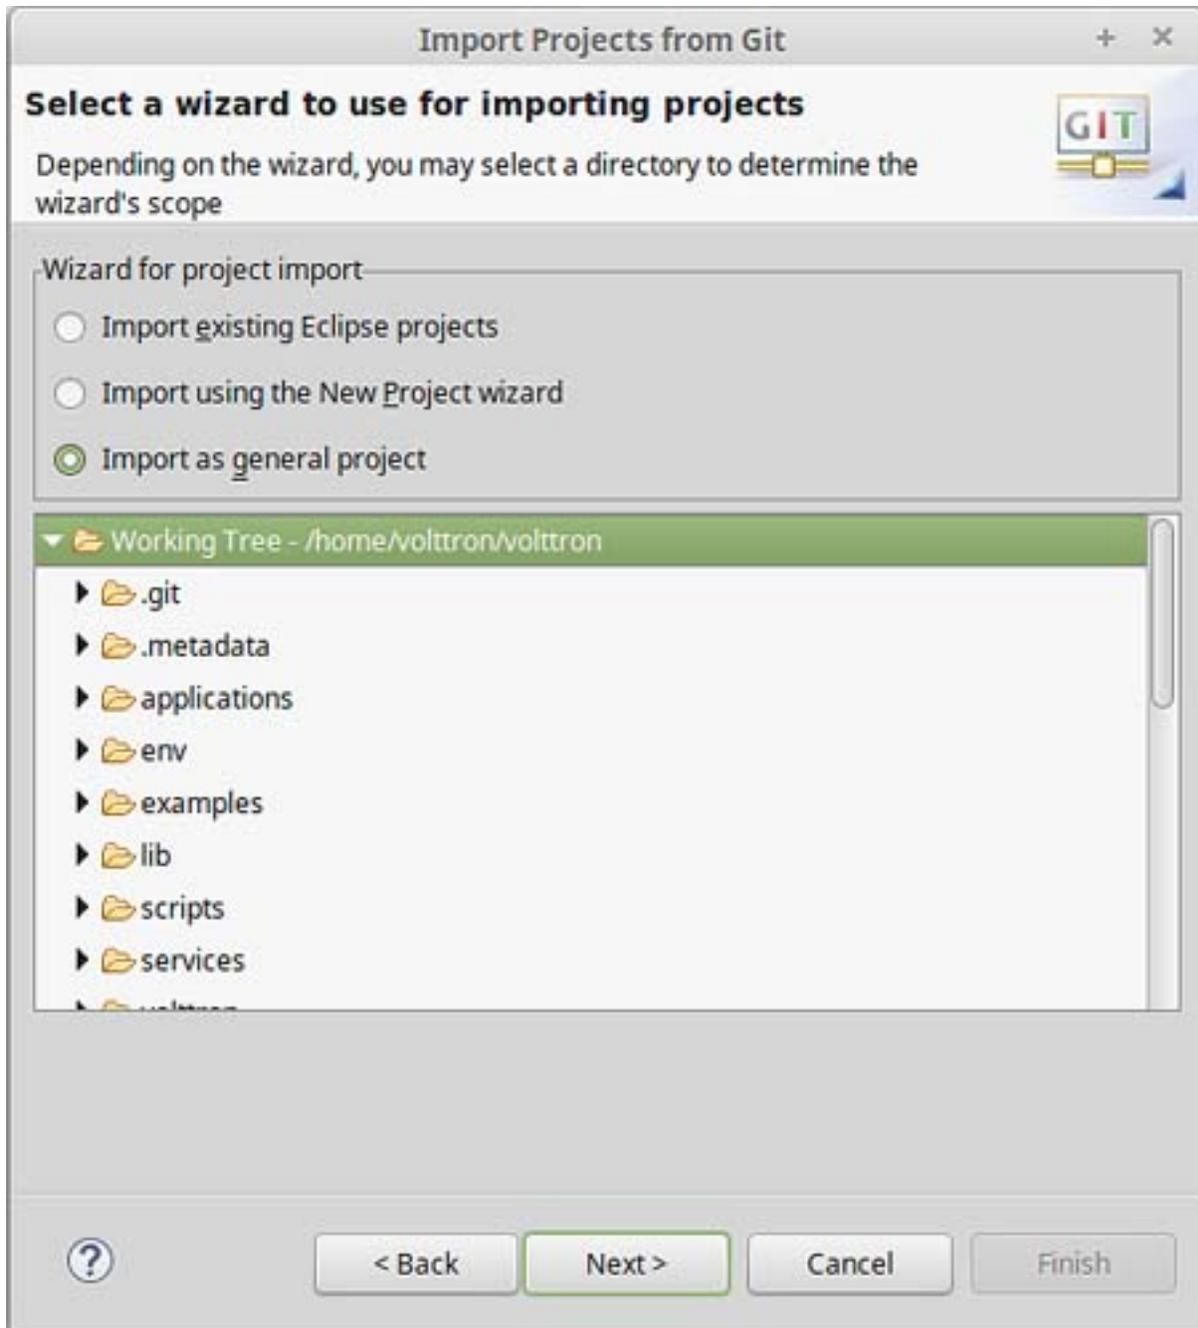


Figure 13. Checking Out VOLTRON with Eclipse from Local Source (continued)

Import New VOLTRON Project from GitHub

To import a new VOLTRON project directly from GitHub into Eclipse, complete the following steps:

1. Select File and select Import (Figure 14).

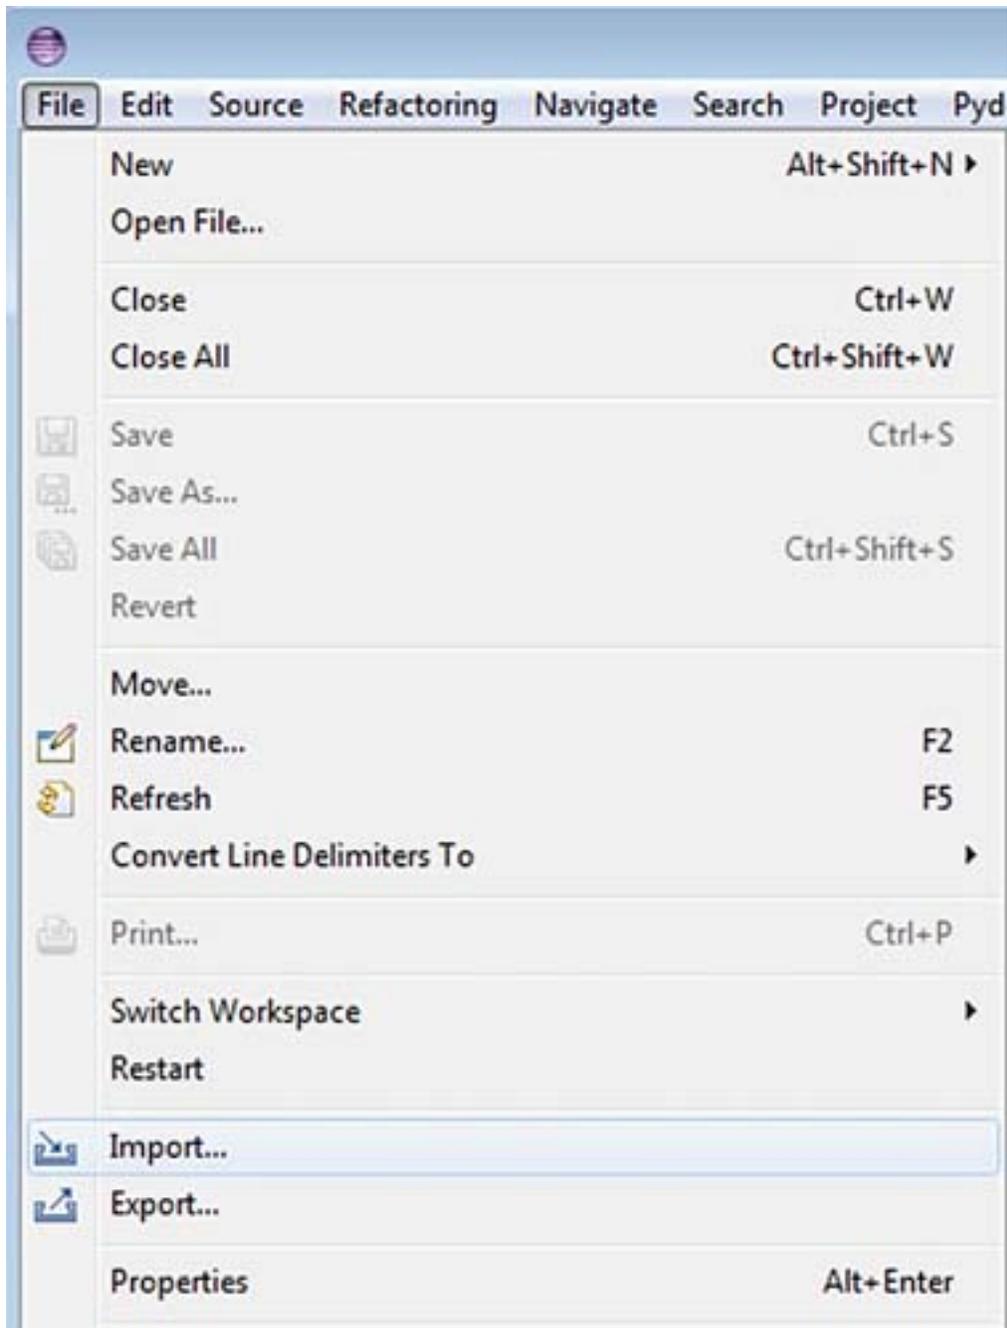


Figure 14. Checking Out VOLTRON with Eclipse from GitHub

2. Select Git. Select Projects from Git. Click the Next button (Figure 15).

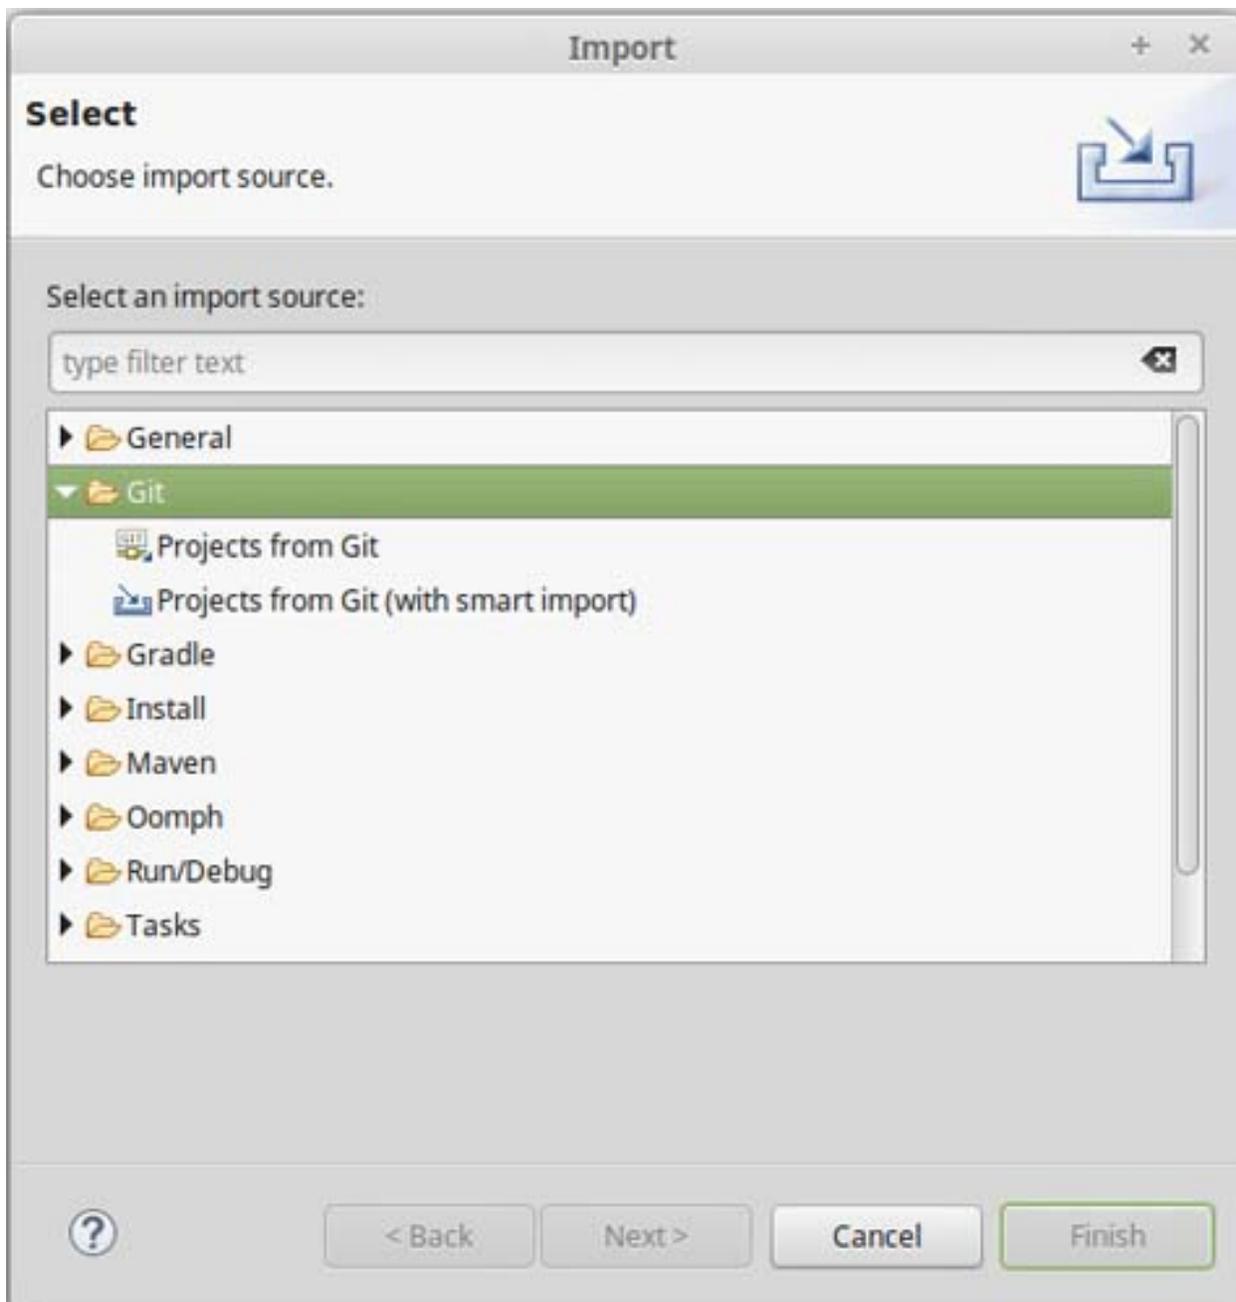


Figure 15. Checking Out VOLTRON with Eclipse from GitHub (continued)

3. Select Clone URI and select Next (Figure 16).

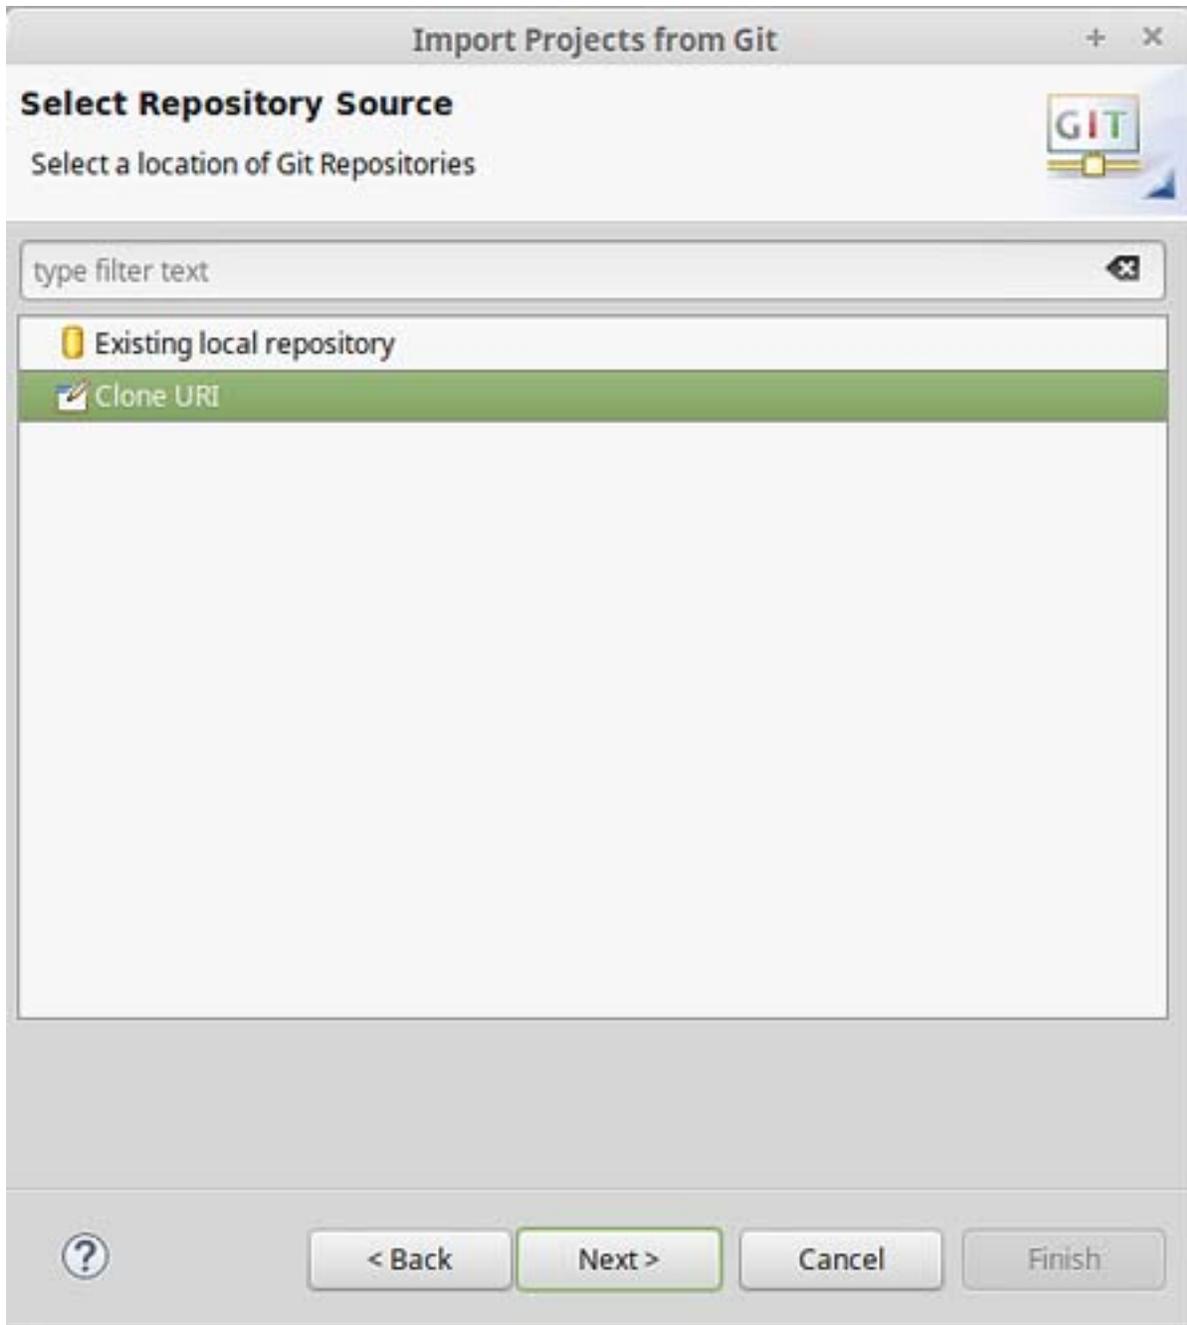


Figure 16. Checking Out VOLTRON with Eclipse GitHub (continued)

4. Fill in <https://github.com/VOLTRON/voltron.git> for the URI. If you have a GitHub account, enter a username and password in the User and Password sections. This is not required but will allow you to receive notifications from GitHub for VOLTRON related news. (Figure 17)

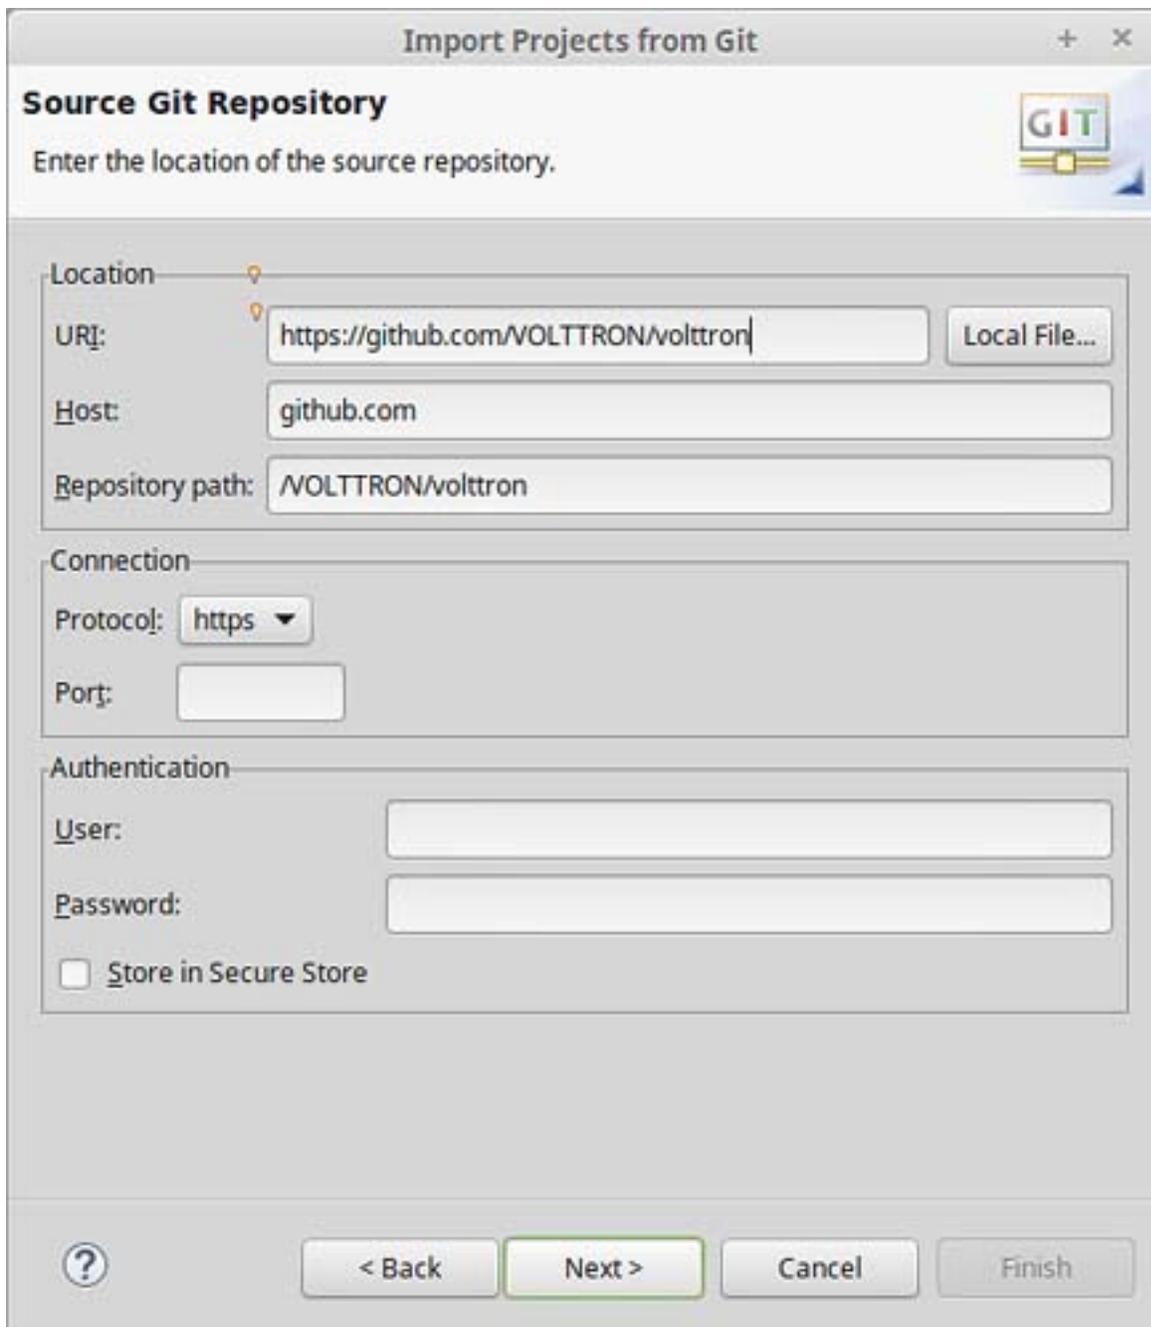


Figure 17. Checking Out VOLTRON with Eclipse from GitHub (continued)

5. Select the master branch (Figure 18).

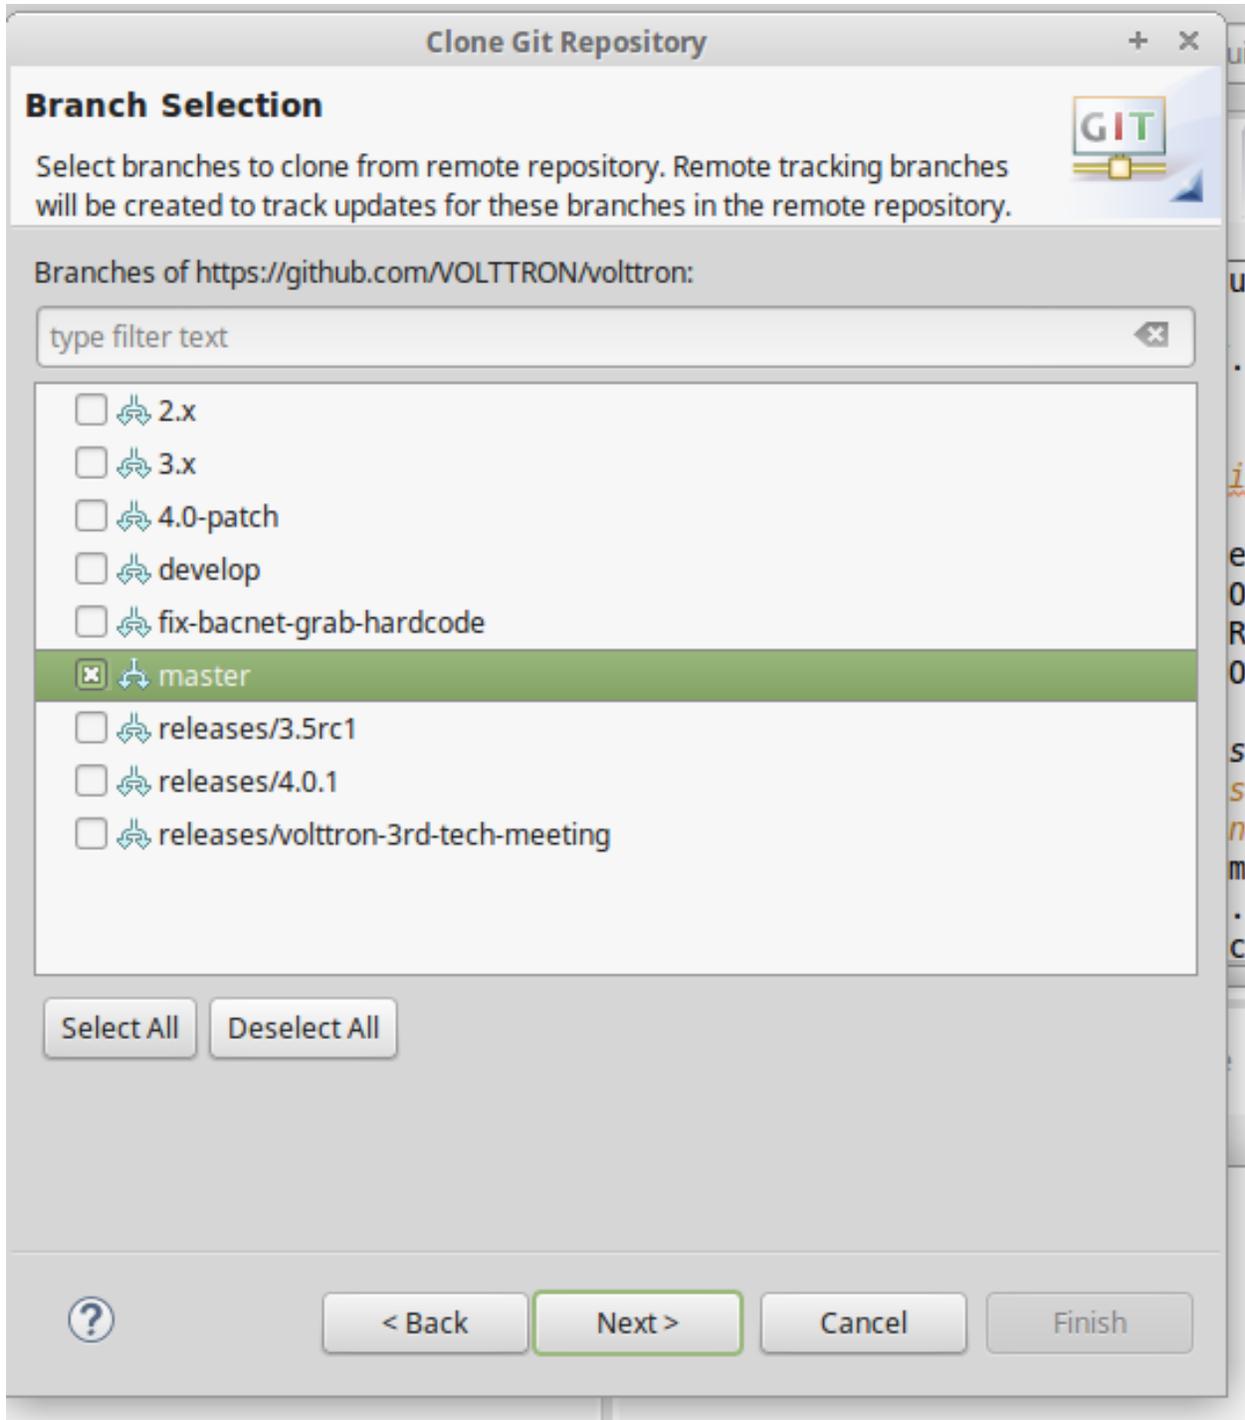


Figure 18. Checking Out VOLTTRON with Eclipse from GitHub (continued)

6. Select a location to save the local repository (Figure 19).

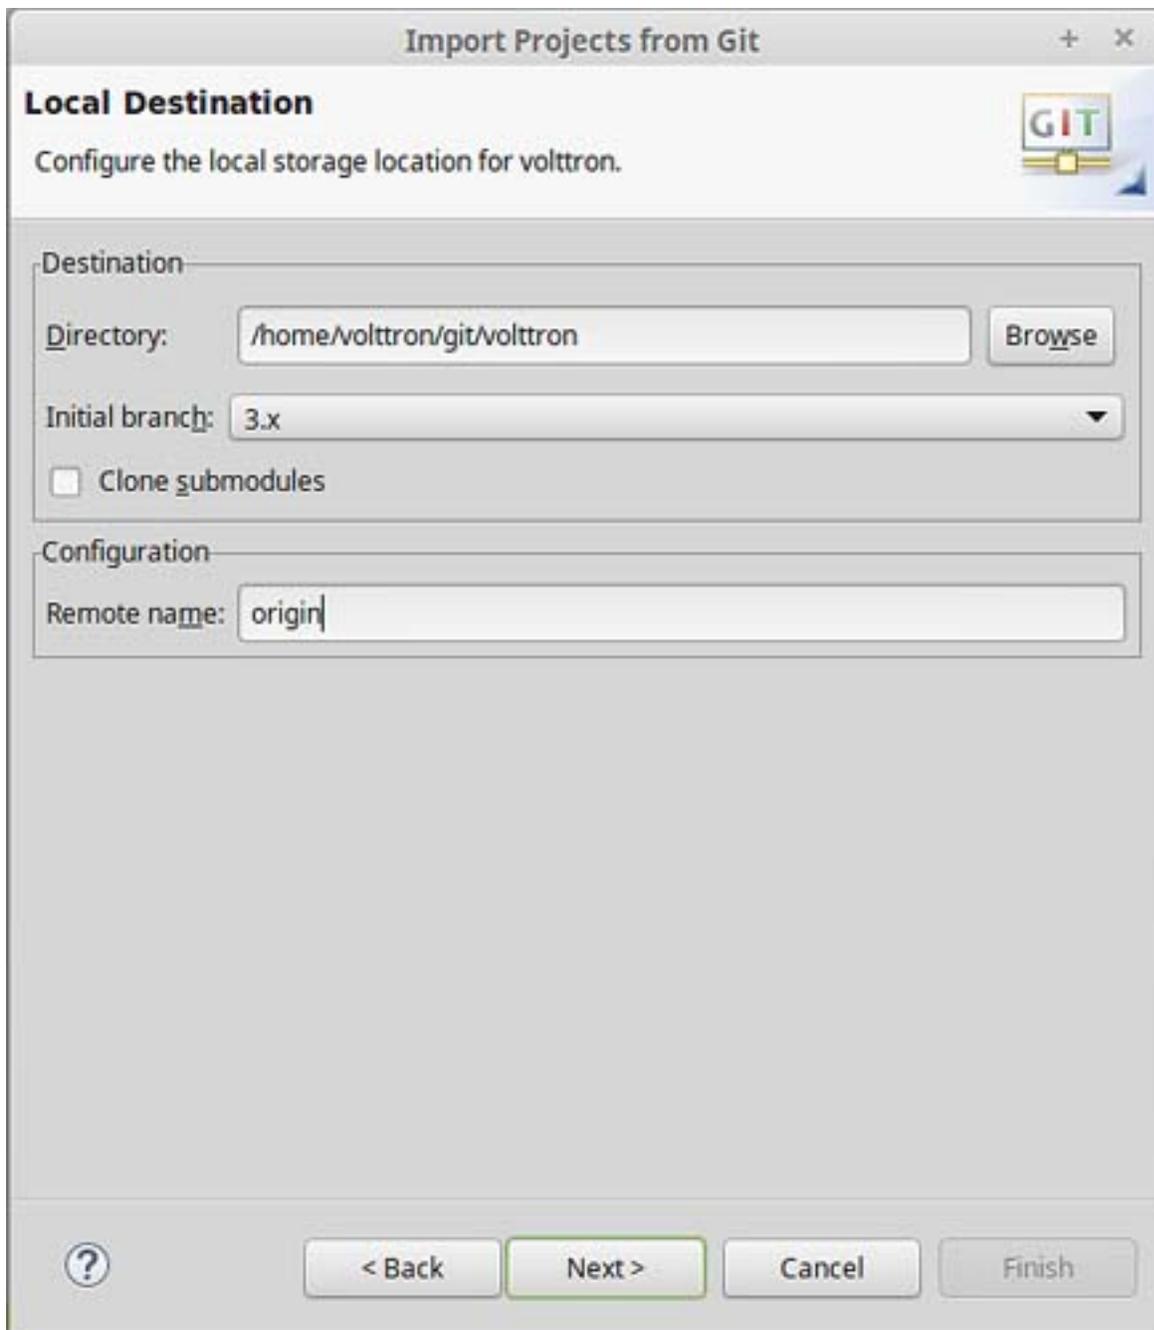


Figure 19. Checking Out VOLTRON with Eclipse from GitHub (continued)

7. Select Import as general project. Select Next. Select Finish (Figure 20). The project will now be imported into the workspace.

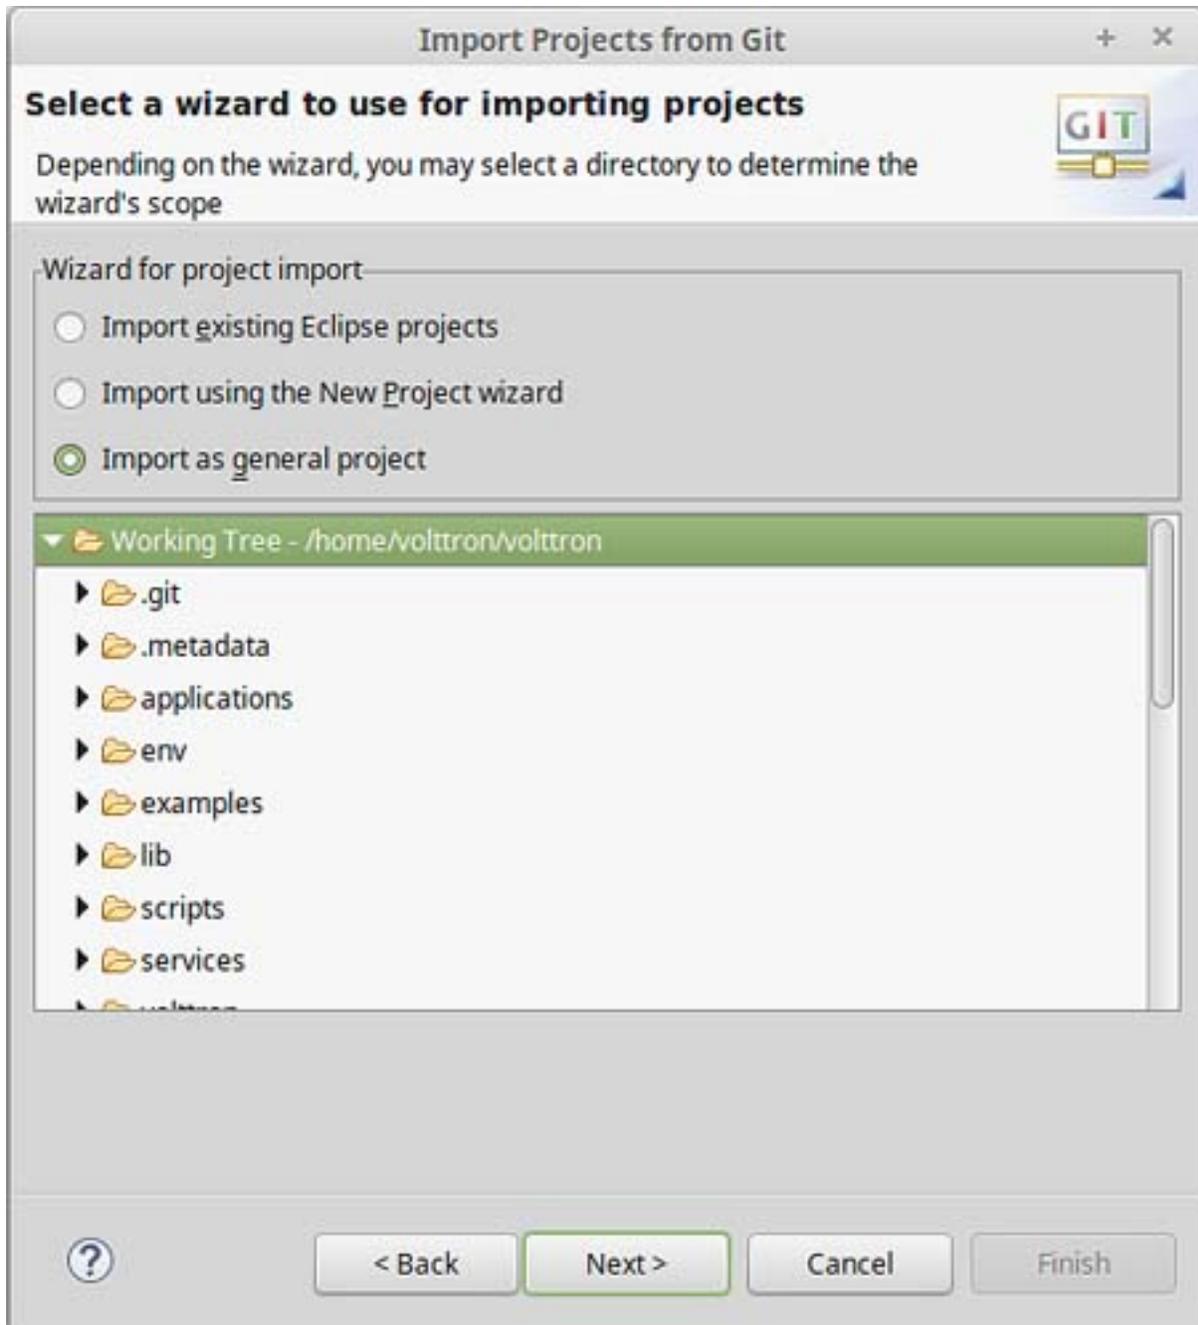


Figure 20. Checking Out VOLTRON with Eclipse from GitHub (continued)

If the VOLTRON project has not been built (<project directory>/bootstrap.py file has not been run), proceed to [Section 2.4 Building the VOLTRON Platform](#) and follow the instruction for running the bootstrap.py script before proceeding to the following sections.

Linking Eclipses

PyDev must now be configured to use the Python interpreter packaged with VOLTRON.

1. Select Window and select Preferences.

2. Expand the PyDev tree.
3. Select Interpreters and select Python interpreter.
4. Select New (Figure 21).

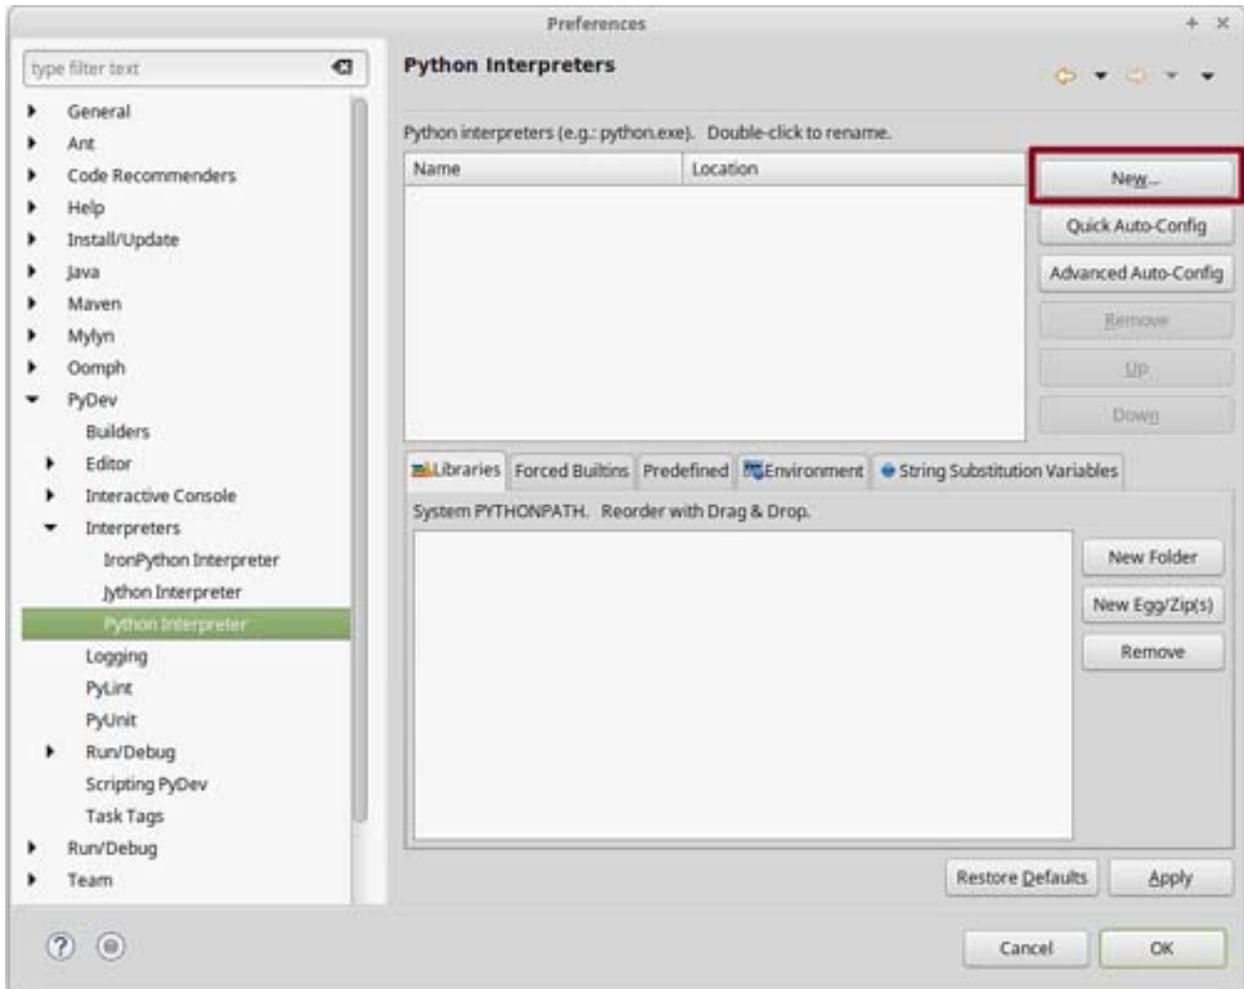


Figure 21. Configuring PyDev

5. Select Browse and navigate to the pydev-python file located at (<project directory>/scripts/pydev-python) (Figure 22).
6. Select OK (Figure 22).



Figure 22. Configuring PyDev (continued)

7. Select All and uncheck the VOLTRON base directory (Figure 23).

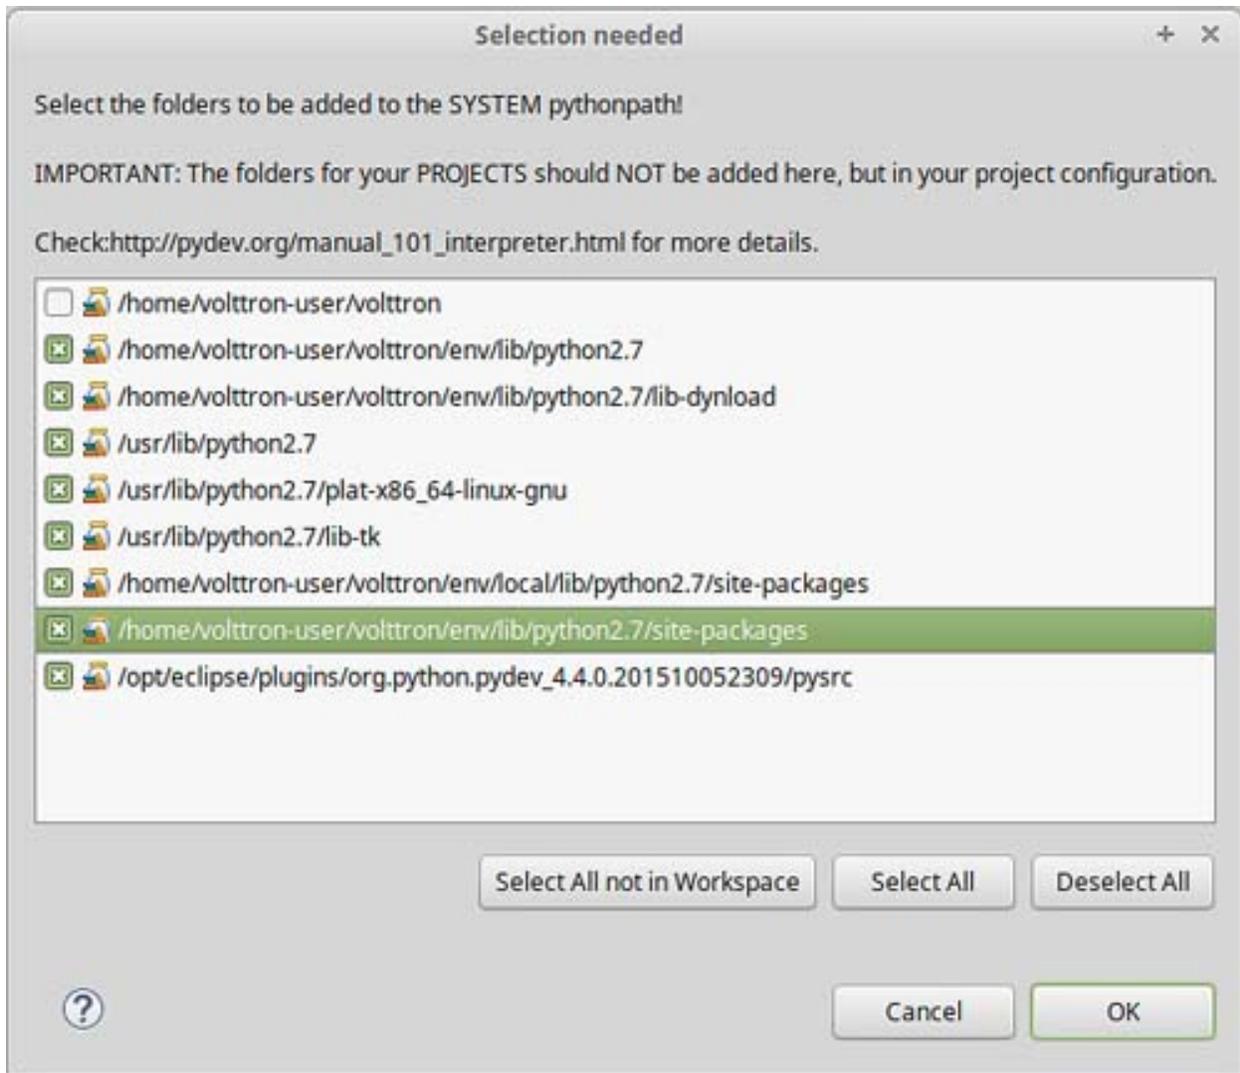


Figure 23. Configuring PyDev (continued)

8. In the Project/PackageExplorer view on the left, right-click on the project, PyDev, and set as PyDev Project (Figure 24).

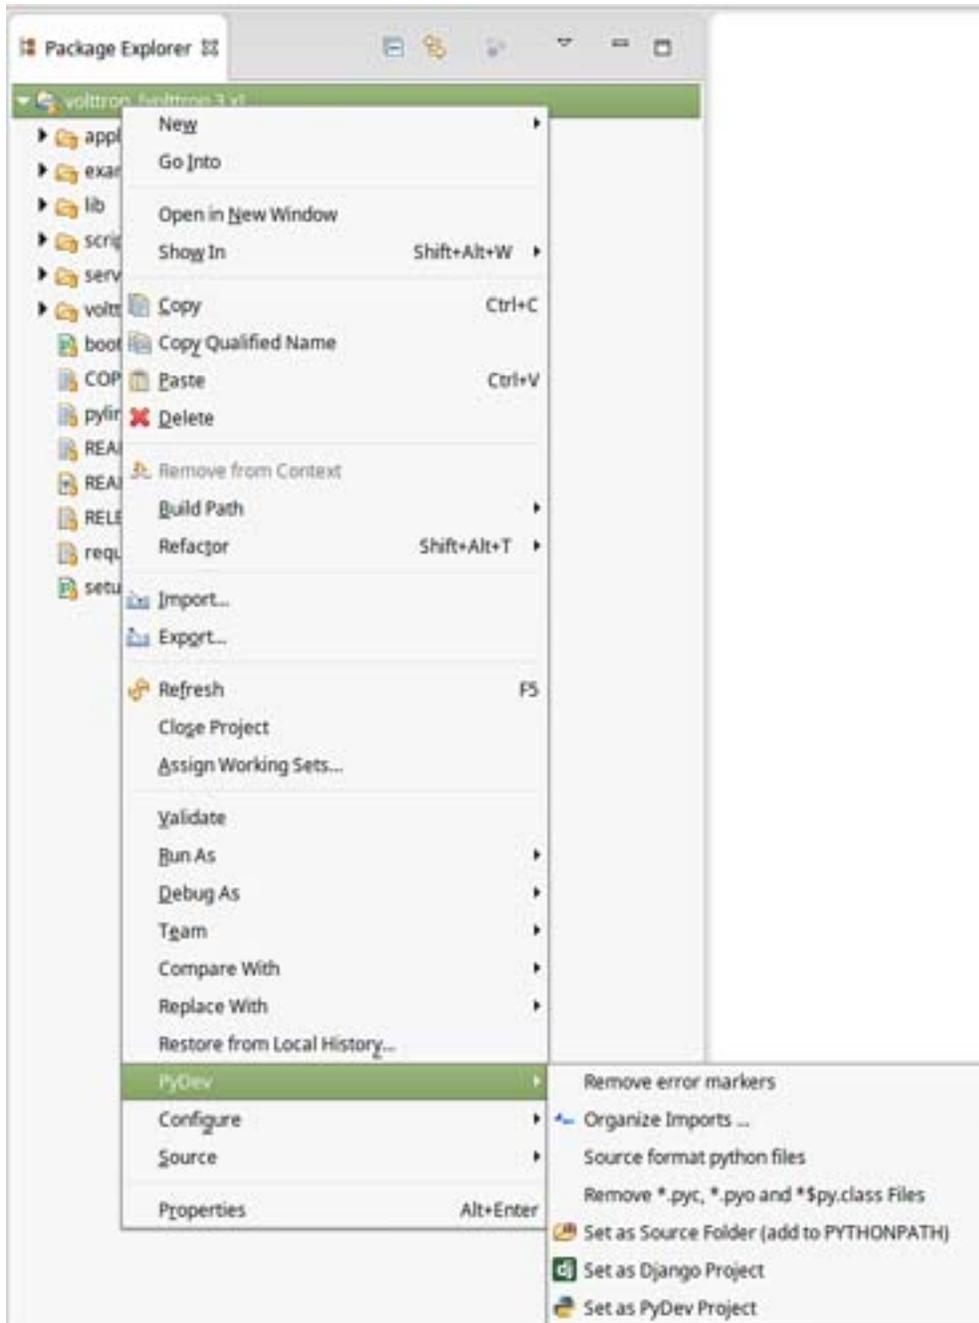


Figure 24. Setting as PyDev Project

9. Switch to the PyDev perspective: Select Window. Select Perspective. Select Open Perspective. Select Other. Select PyDev (Figure 25). Eclipse should now be configured to use the project's environment.

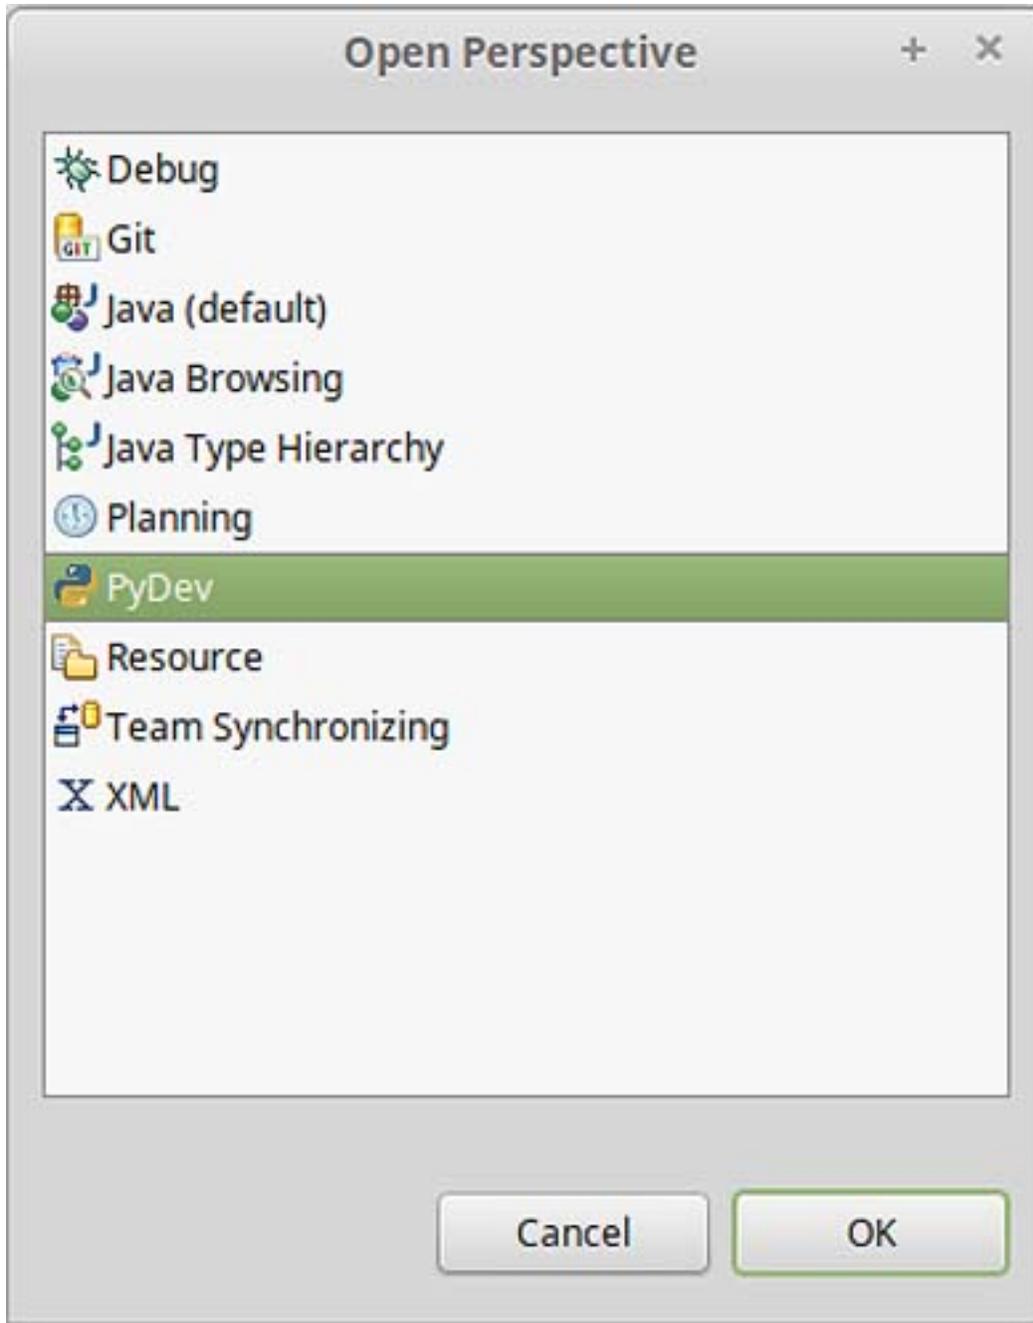


Figure 25. Setting PyDev Perspective in Eclipse

Running the VOLTRON Platform and Agents

VOLTRON and agents within VOLTRON can now be run within Eclipse. This section will describe the process to run VOLTRON and an agent within Eclipse.

Setup a Run Configuration for the Platform

The following steps describe the process for running VOLTRON within Eclipse:

1. Select Run and select Run Configurations (Figure 26).

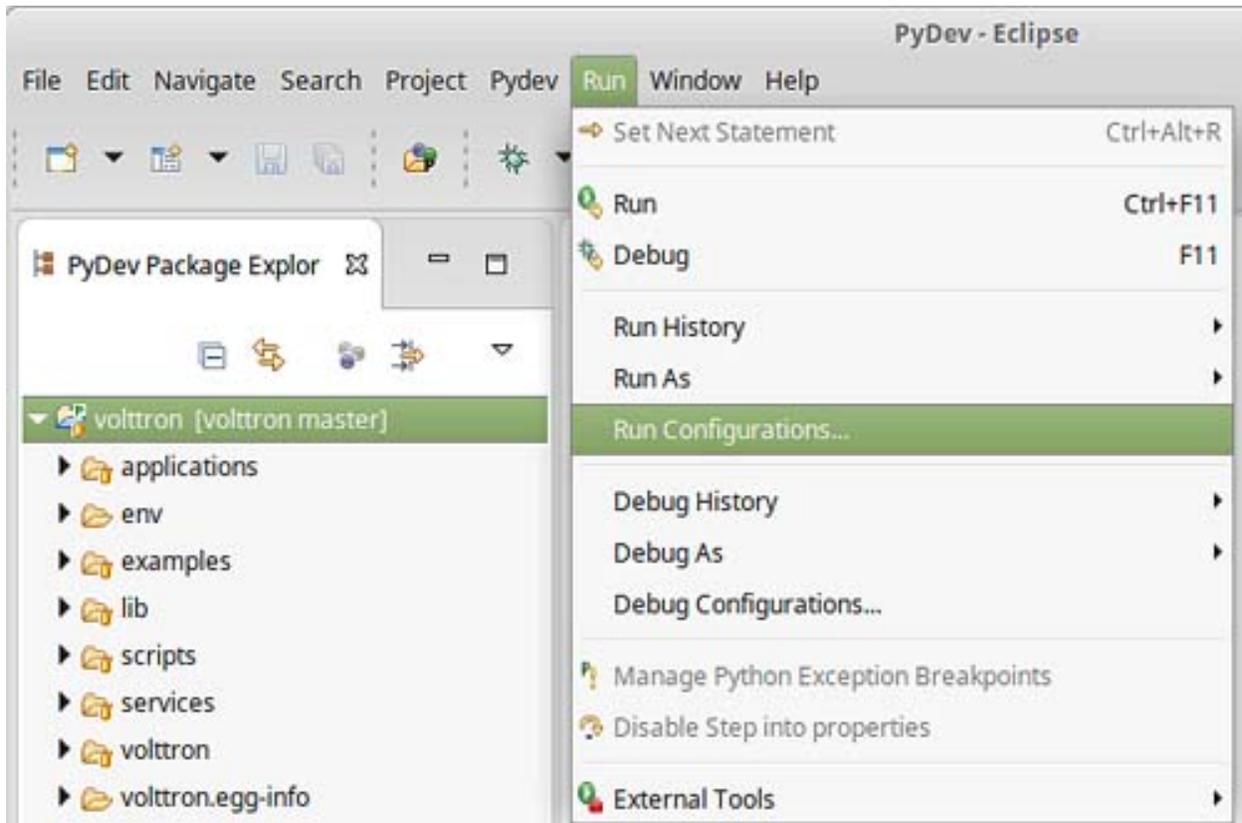


Figure 26. Running VOLTRON Platform, Setting Up a Run Configuration

2. Select Python Run from the menu on left. Click the New launch configuration button (Figure 27).

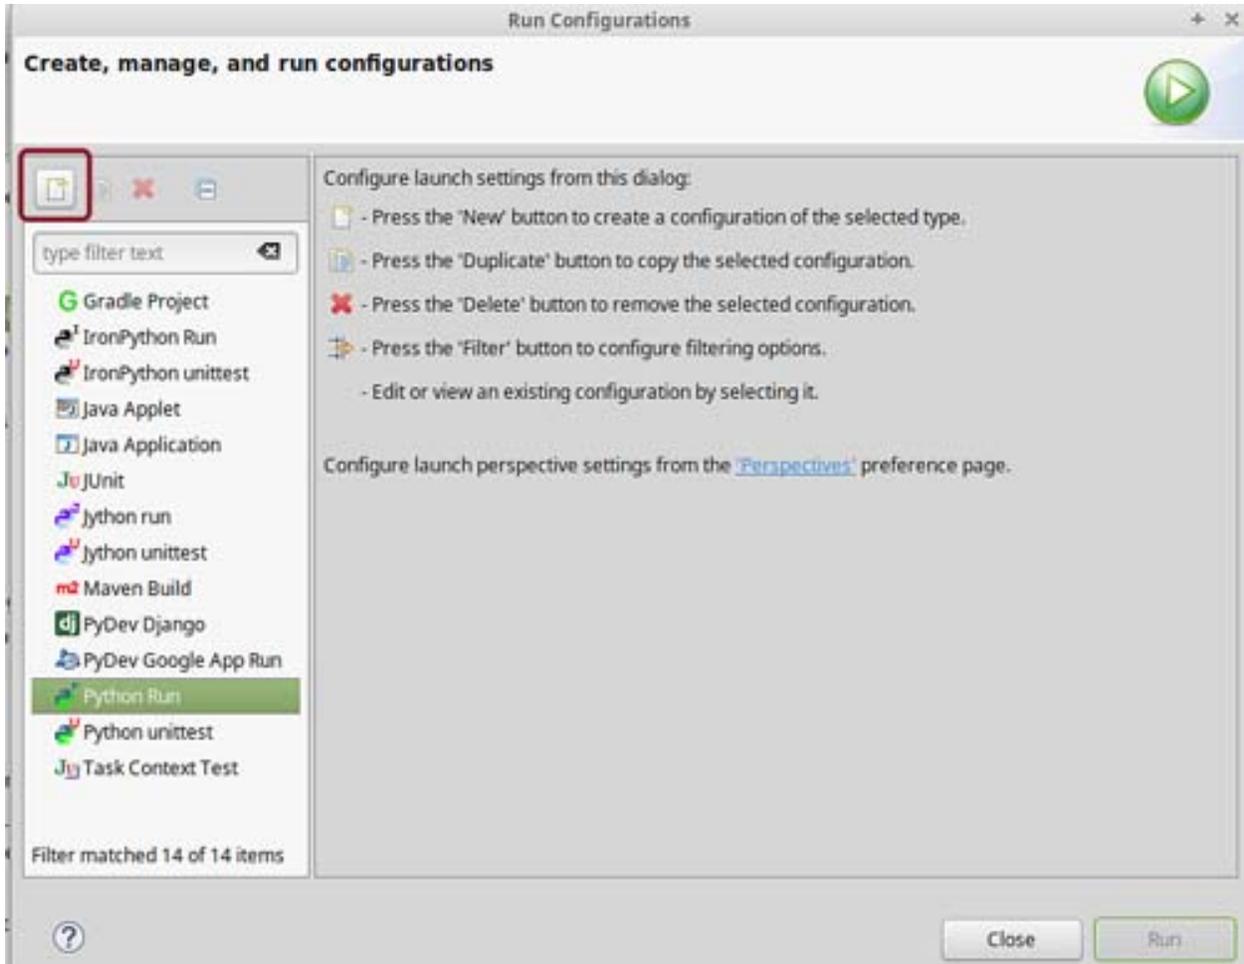


Figure 27. Running VOLTRON Platform, Setting Up a Run Configuration (continued)

3. Change the name (any name may be used but for this example the name VOLTRON was chosen) and select the main module (`<project directory>/volttron/platform/main.py`).
4. Select the Arguments tab and enter `'-vv'` in the Program arguments field (Figure 28) then select the Run button.

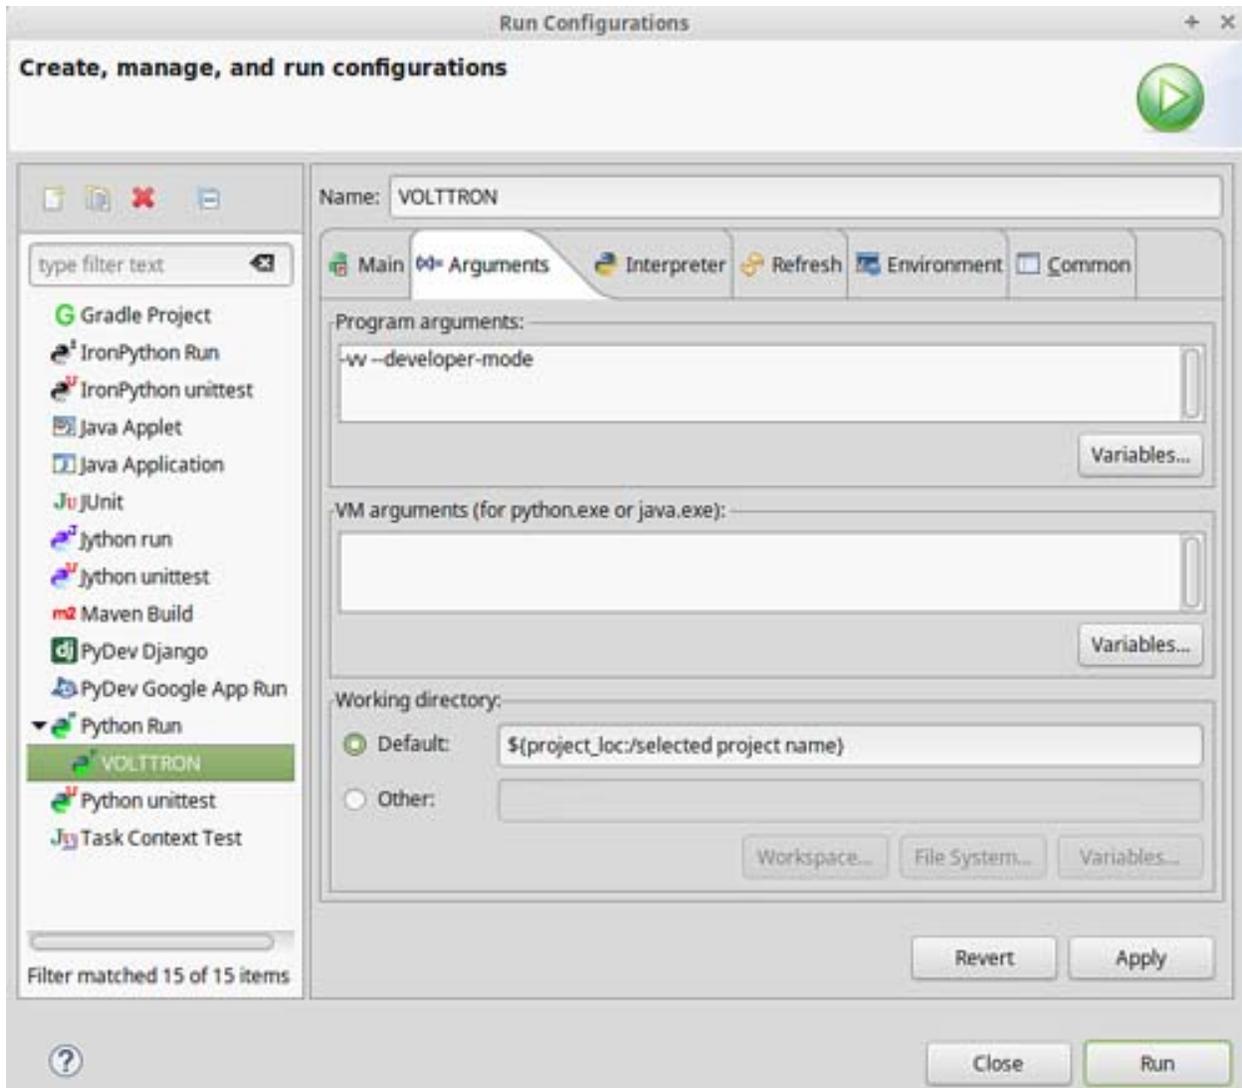


Figure 28. Running VOLTRON Platform, Setting Up a Run Configuration (continued)

5. If the run is successful, the console should appear similar to Figure 29. If the run does not succeed (red text describing why the run failed will populate the console), click the all stop icon (two red boxes overlaid) on the console and then retry.

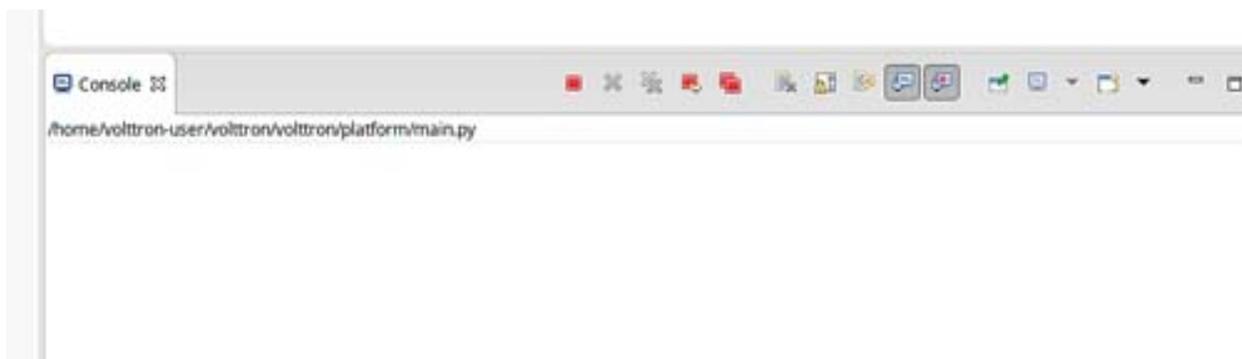


Figure 29. Running VOLTRON Platform, Console View on Successful Run

Configure a Run Configuration for the Listener Agent

The following steps describe the process for configuring an agent within Eclipse:

1. Select Run and select Run Configurations (Figure 30).

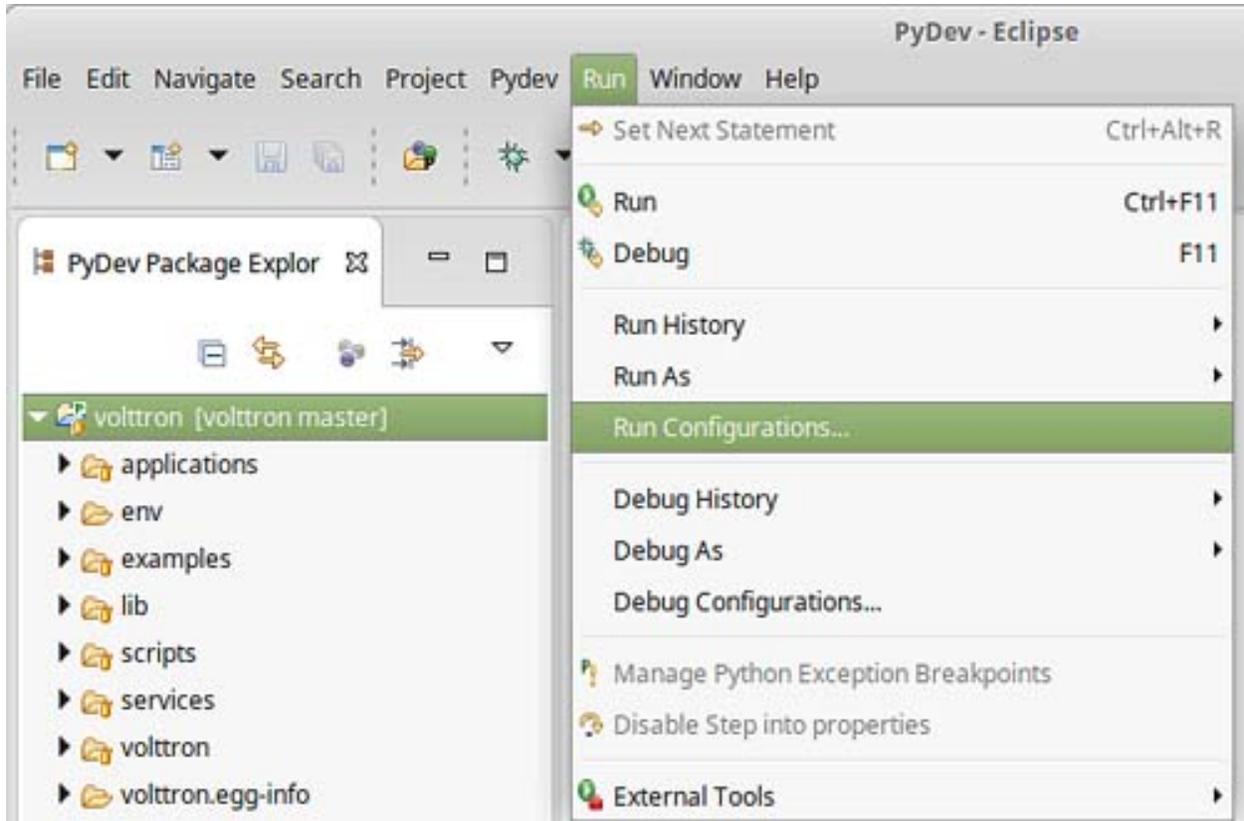


Figure 30. Running the Listener Agent, Setting Up a Run Configuration

2. Select Python Run from the menu on left and click the New launch configuration button (Figure 31).

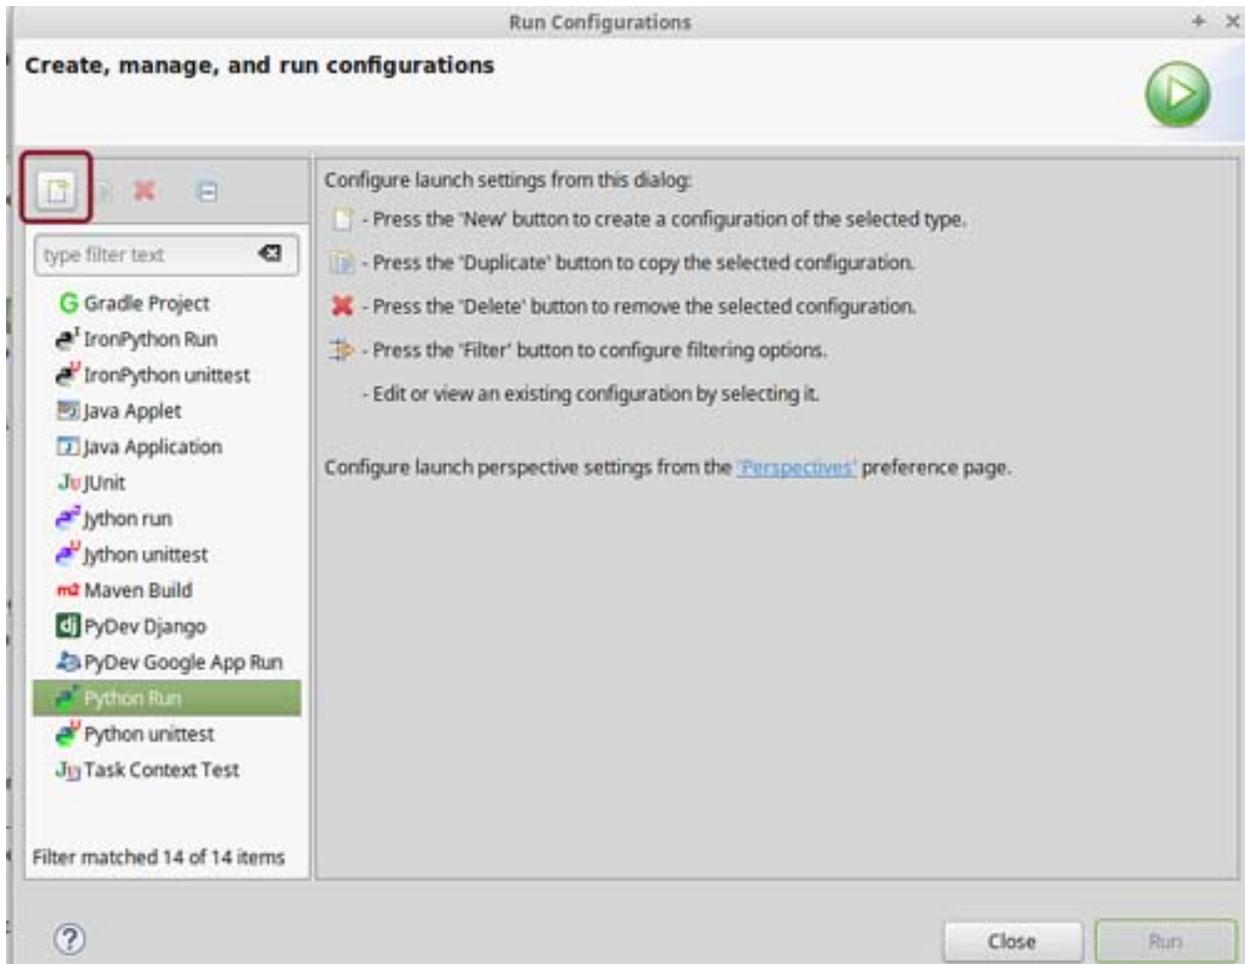


Figure 31. Running the Listener Agent, Setting Up a Run Configuration (continued)

3. Change the name (for this example Listener is used) and select the main module (`<project directory>/examples/ListenerAgent/listener/agent.py`) (Figure 32).

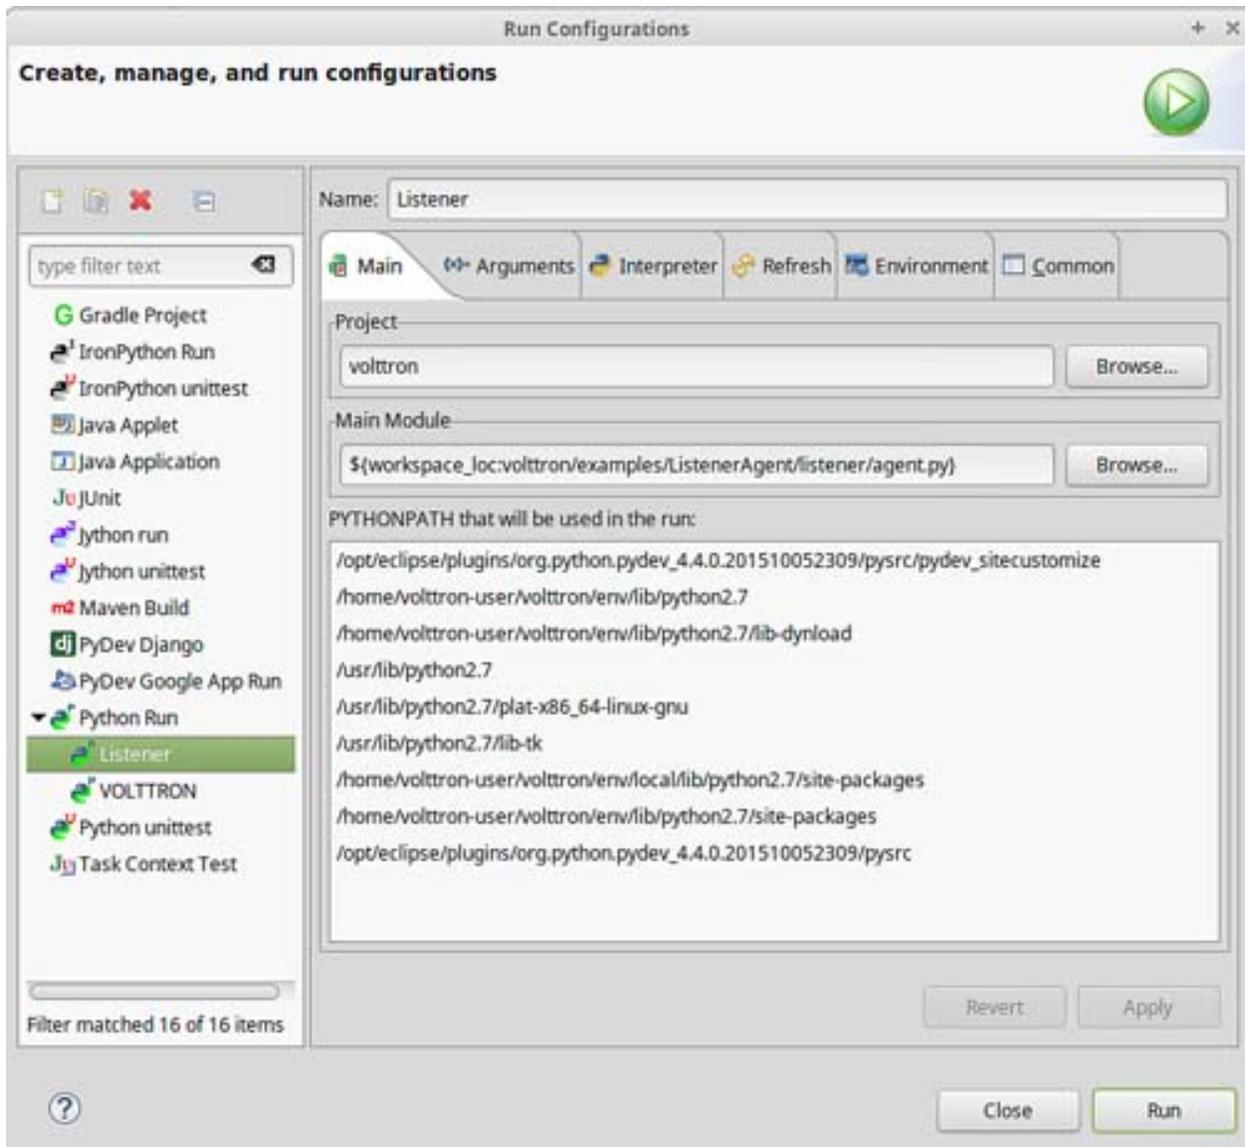


Figure 32. Running the Listener Agent, Setting Up a Run Configuration (continued)

4. Click the Arguments tab and change Working directory to Default (Figure 33).

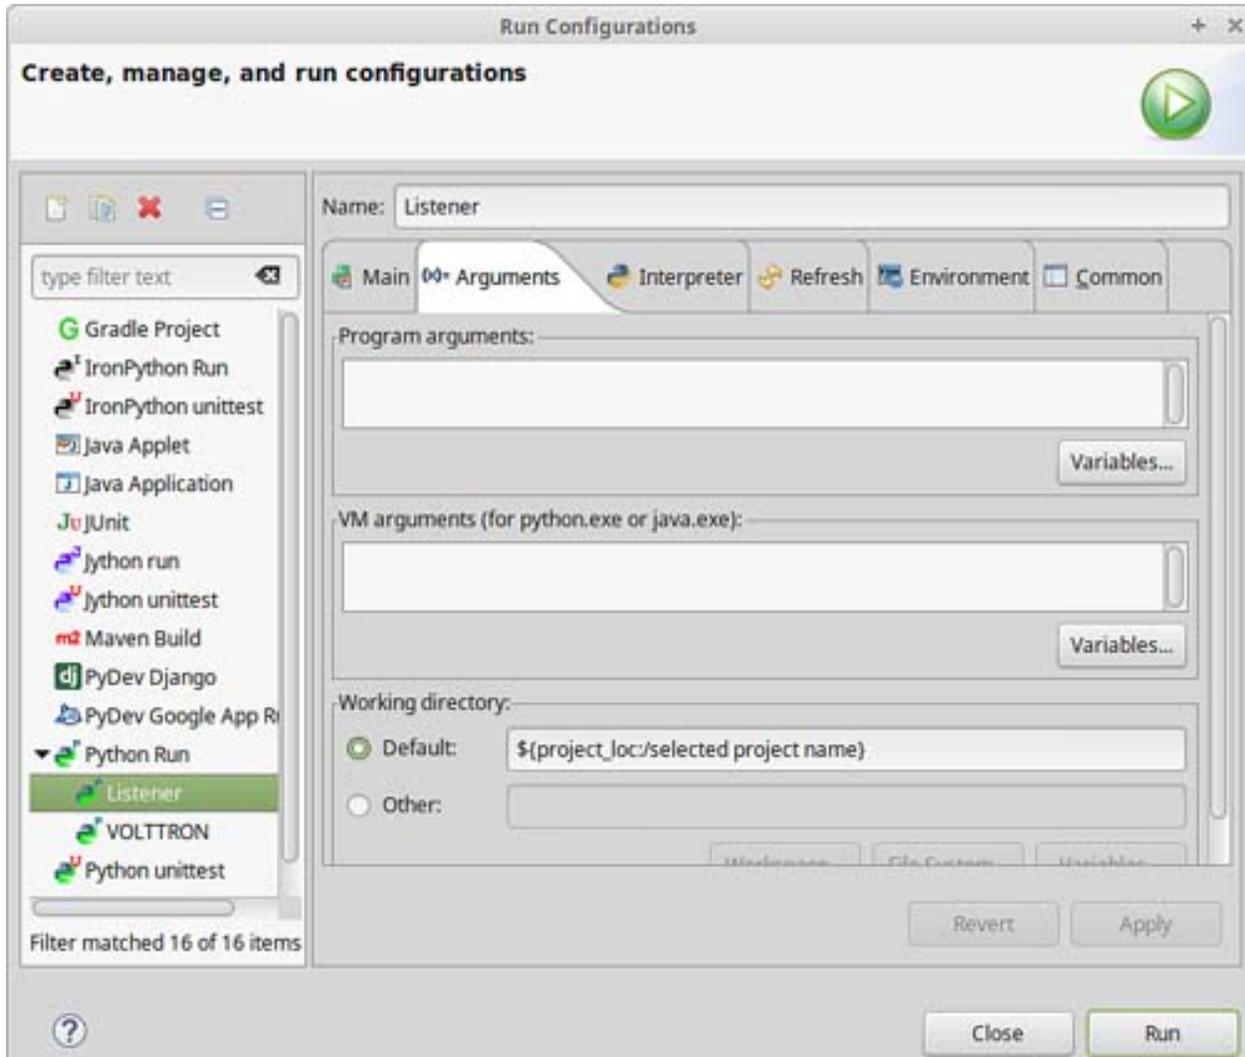


Figure 33. Running the Listener Agent, Setting Up a Run Configuration (continued)

5. In the Environment tab, select New and add the following environment variables (bulleted list below), as shown in Figure 34:

- `AGENT_CONFIG = /home/<USER>/examples /ListenerAgent/config`

`AGENT_CONFIG` is the absolute path the agent's configuration file. To access a remote message bus, use the VIP address as described in [Section 3.5 Platform Management:VOLTRON Management Central](#).

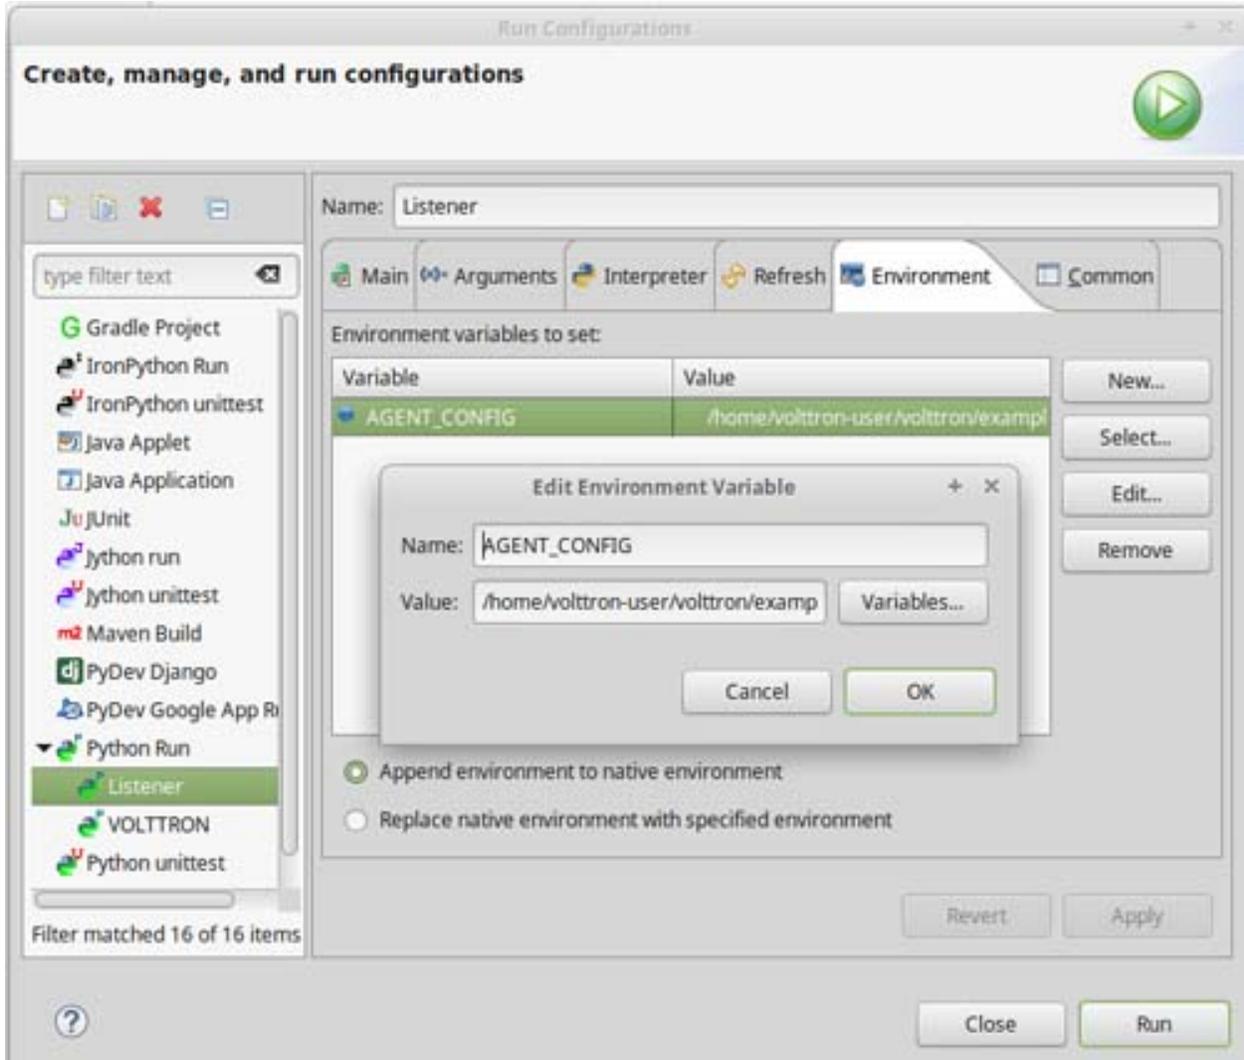


Figure 34. Running the Listener Agent, Setting Up a Run Configuration

6. Click Run. This launches the agent. You should see the agent start to publish and receive its own heartbeat message (Figure 35).

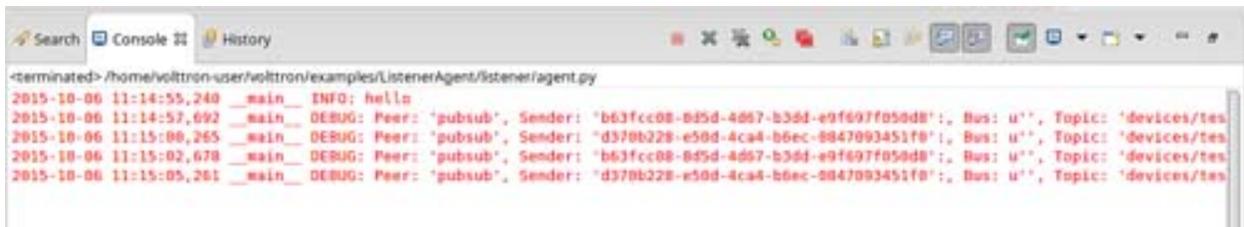


Figure 35. Listener Agent Output on Eclipse Console

The process for running other agents in Eclipse is identical to that of the Listener agent. Several useful development tools are available within Eclipse and PyDev that make development, debugging, and testing of agents much simpler.

Agent Creation Walkthrough

Developers should look at the Listener agent before developing their own agent. The Listener agent illustrates the basic functionality of an agent. The following example demonstrates the steps for creating an agent.

Agent Folder Setup

Create a folder within the workspace to help consolidate the code your agent will utilize.

1. In the VOLTRON base directory, create a new folder TestAgent.
2. In TestAgent, create a new folder tester. This is the package where the Python code will be created (Figure 36).

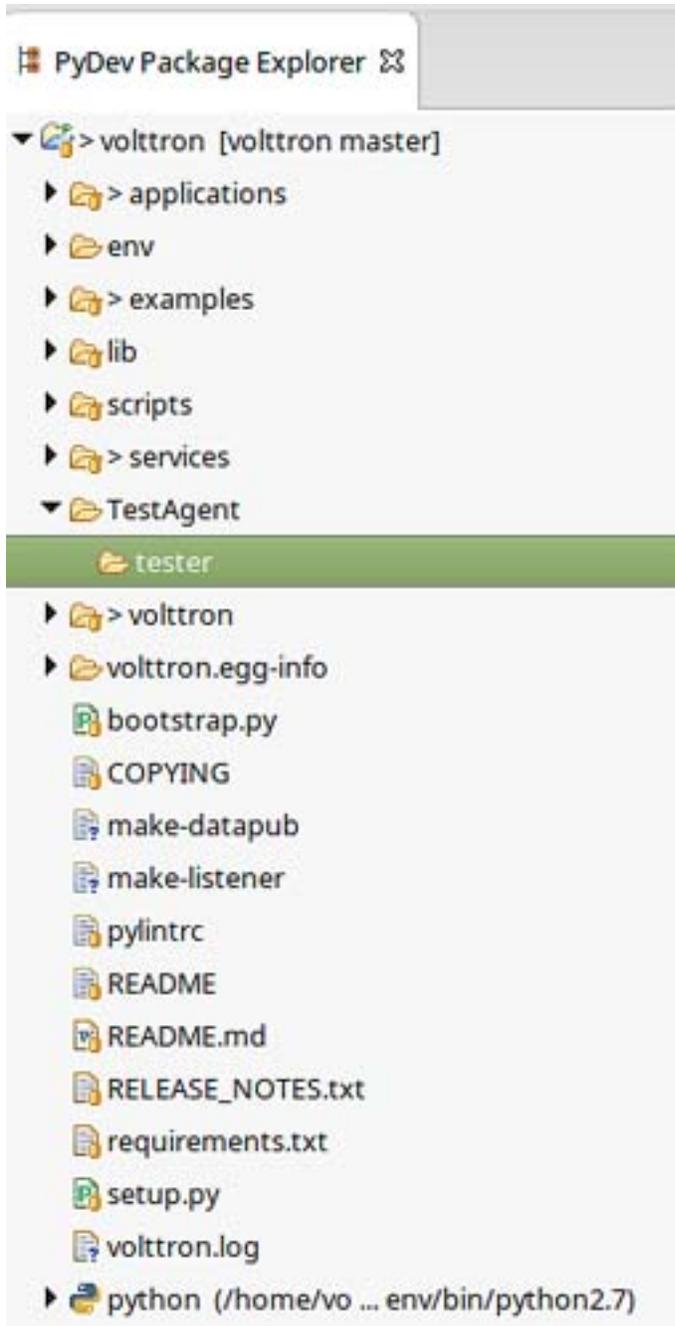


Figure 36. Creating an Agent Test Folder

Create Agent Code

The following steps describe the necessary agent files and modules.

1. In tester, create a file called `__init__.py`, which tells Python to treat this folder as a package.
2. In the tester package folder, create the file `testagent.py`
3. Create a class called TestAgent.

4. Import the packages and classes needed:

```
from __future__ import absolute_import

from datetime import datetime
import logging
import sys

from voltttron.platform.vip.agent import Agent, Core
from voltttron.platform.agent import utils
```

5. Set up a logger. The `utils` module from `voltttron.platform.agent` builds on Python's already robust logging module and is easy to use. Add the following lines after the import statements:

```
utils.setup_logging()
_log = logging.getLogger(__name__)
```

This agent will inherit features from the `Agent` class (base class) extending the agent's default functionality. The class definition for the `TestAgent` will be configured as shown below (with `__init__`).

```
class TestAgent(Agent):
    def __init__(self, config_path, **kwargs):
        super(TestAgent, self).__init__(**kwargs)
```

Setting up a Subscription

1. Create a startup method. This method is tagged with the decorator `@Core.receiver("onstart")`. The startup method will run after the agent is initialized. The `TestAgent`'s startup method will contain a subscription to the `Listener` agent's heartbeat (`heartbeat/listeneragent`). The `TestAgent` will detect when a message with this topic is published on the message bus and will run the method specified with the `callback` keyword argument passed to `self.vip.pubsub.subscribe`.

```
@Core.receiver("onstart")
def starting(self, sender, **kwargs):
    '''
    Subscribes to the platform message bus on
    the heartbeat/listeneragent topic
    '''
    print('TestAgent example agent start-up function')
    self.vip.pubsub.subscribe('pubsub', 'heartbeat/listeneragent',
                              callback=self.on_heartbeat)
```

2. Create the callback method. Typically, the callback is the response to a message (or event). In this simple example, the `TestAgent` will do a print statement and publish a message to the bus:

```
def on_heartbeat(self, peer, sender, bus, topic, headers, message):
    '''TestAgent callback method'''
    print('Matched topic: {}, for bus: {}'.format(topic, bus))
    self.vip.pubsub.publish('pubsub',
                            'testagent/publish',
                            headers=headers,
                            message='test publishing').get(timeout=30)
```

Argument Parsing Main Method

The test agent will need to be able to parse arguments being passed on the command line by the agent launcher. Use the `utils.default_main` method to handle argument parsing and other default behavior.

1. Create a main method that can be called by the launcher:

```
def main(argv=sys.argv):
    '''Main method called by the eggsecutable.'''
    try:
        utils.vip_main(TestAgent)
    except Exception as e:
        _log.exception(e)

if __name__ == '__main__':
    # Entry point for script
    sys.exit(main())
```

Create Support Files for Test Agent

VOLTRON agents need configuration files for packaging, configuration, and launching. The “setup.py” file details the naming and Python package information. The launch configuration file is a JSON-formatted text file used by the platform to launch instances of the agent.

Packaging Configuration

In the TestAgent folder, create a file called “setup.py”. This file sets up the name, version, required packages, method to execute, etc. for the agent. The packaging process will also use this information to name the resulting file.

```
from setuptools import setup, find_packages

#get environ for agent name/identifier
packages = find_packages('.')
package = packages[0]

setup(
    name = package + 'agent',
    version = "0.1",
    install_requires = ['volttron'],
    packages = packages,
    entry_points = {
        'setuptools.installation': [
            'eggsecutable = ' + package + '.testagent:main',
        ]
    }
)
```

Launch Configuration

In TestAgent, create a file called “testagent.launch.json”. This is the file the platform will use to launch the agent. It can also contain configuration parameters for the agent:

```
{
  "agentid": "Test1"
}
```

Testing the Agent

From a terminal, in the base VOLTRON directory, enter the following commands (with the platform activated and VOLTRON running):

1. Run `pack_install` script on TestAgent:

```
$ ./scripts/core/pack_install.sh TestAgent TestAgent/config test-agent
```

- Upon successful completion of this command, the terminal output will show the install directory, the agent UUID (unique identifier for an agent; the UUID shown in red is only an example and each instance of an agent will have a different UUID) and the agent name (blue text):

```
Installed /home/volttron-user/.volttron/packaged/testeragent-0.1-py2-none-any.whl
as d4ca557a-496c-4f02-8ad9-42f5d435868a testeragent-0.1
```

2. Start the agent:

```
$ volttron-ctl start --tag test-agent
```

3. Verify that the agent is running:

```
$ volttron-ctl status
$ tail -f volttron.log
```

If changes are made to the Passive AFDD agent's configuration file after the agent is launched, stop and reload the agent. In a terminal, enter the following commands:

```
$ volttron-ctl stop --tag test-agent
$ volttron-ctl remove --tag test-agent
```

Re-build and start the updated agent (Figure 37).

```
2015-10-06 17:37:33,045 [listeneragent-0.1 27020] listener.agent DEBUG: Peer: 'pubsub', Sender: '036f0934-e065-4c7a-a099-080204bae004', Bus: 'u',
topic: 'testagent/publish', Headers: {'Date': '2015-10-07 09:37:33.563216Z', 'AgentID': 'listener1', 'Content-Type': 'text/plain'}, Message: 'test
blinking'
2015-10-06 17:37:33,567 [testeragent-0.1 27260] tester.testagent INFO: Matched topic: heartbeat/listeneragent, for bus:
```

Figure 37. TestAgent Output In VOLTRON Log

Running the TestAgent in Eclipse

Warning: Before attempting to run an agent in Eclipse, please see the note in: [AgentDevelopment](#)

If you are working in Eclipse, create a run configuration for the TestAgent based on the Listener agent configuration in the Eclipse development environment `##(Section 5.5.5 Running the VOLTRON Platform and Agents)##`.

1. Launch the platform (`##Section 5.5.5.1 Setup a Run Configuration for the Platform##`)
2. Launch the TestAgent by following the steps outlined in *Launching the Listener <Start-Listener-Eclipse>* for launching the Listener agent.

3. Launch the Listener agent. TestAgent should start receiving the heartbeats from Listener agent and the following should be displayed in the console (Figure 38).

```

/home/volttron-user/volttron/TestAgent/tester/testagent.py
2015-10-06 18:09:36,717 __main__ INFO: TestAgent example agent start-up function
2015-10-06 18:09:39,092 __main__ INFO: Matched topic: heartbeat/listeneragent, for bus:
2015-10-06 18:09:44,121 __main__ INFO: Matched topic: heartbeat/listeneragent, for bus:
2015-10-06 18:09:49,114 __main__ INFO: Matched topic: heartbeat/listeneragent, for bus:
2015-10-06 18:09:54,118 __main__ INFO: Matched topic: heartbeat/listeneragent, for bus:
    
```

Figure 38. Console Output for TestAgent

Adding Additional Features to the TestAgent

Additional code can be added to the TestAgent to utilize additional services in the platform. The following sections show how to use the weather and device scheduling service within the TestAgent.

Subscribing to Weather Data

This agent can be modified to listen to weather data from the Weather agent by adding the following line at the end of the TestAgent startup method. This will subscribe the agent to the temperature subtopic. For the full list of topics available, please see:

<https://github.com/VOLTTRON/volttron/wiki/WeatherAgentTopics>

```

self.vip.pubsub.subscribe('pubsub', 'weather/temperature/temp_f',
                           callback=self.on_weather)
    
```

Add the callback method `on_weather`:

```

def on_weather(self, peer, sender, bus, topic, headers, message):
    print("TestAgent got weather\nTopic: {}, Message: {}".format(topic, message))
    
```

The platform log file should appear similar to Figure 39.



Figure 39. TestAgent Output when Subscribing to Weather Topic

Utilizing the Scheduler Agent

The TestAgent can be modified to publish a schedule to the Actuator agent by reserving time on virtual devices. Modify the following code to include current time ranges and include a call to the publish schedule method in setup. The following example posts a simple schedule. For more detailed information on device scheduling, please see:

<https://github.com/VOLTTRON/volttron/wiki/ActuatorAgent>

Ensure the Actuator agent is running as per [Section 3.3 Device Control: Configuring and Launching the Actuator Agent](#). Add the following line to the TestAgent’s import statements:

```

from volttron.platform.messaging import topics
    
```

Add the following lines to the TestAgent's starting method. This sets up a subscription to the **ACTUATOR_RESPONSE** topic and calls the **publish_schedule** method.

```
self.vip.pubsub.subscribe('pubsub', topics.ACTUATOR_RESPONSE,
                          callback=self.on_schedule_result)
self.publish_schedule()
```

The **publish_schedule** method sends a schedule request message to the Actuator agent (Update the schedule with appropriate times):

```
def publish_schedule(self):
    headers = {
        'AgentID': self._agent_id,
        'type': 'NEW_SCHEDULE',
        'requesterID': self._agent_id, # Name of requesting agent
        'taskID': self._agent_id + "-TASK", # Unique task ID
        'priority': 'LOW'               # Task Priority (HIGH, LOW, LOW_PREEMPT)
    }
    msg = [
        ["campus/building/device1", # First time slot.
         "2014-1-31 12:27:00",      # Start of time slot.
         "2014-1-31 12:29:00"],     # End of time slot.
        ["campus/building/device1", # Second time slot.
         "2014-1-31 12:26:00",      # Start of time slot.
         "2014-1-31 12:30:00"],     # End of time slot.
        ["campus/building/device2", # Third time slot.
         "2014-1-31 12:30:00",      # Start of time slot.
         "2014-1-31 12:32:00"],     # End of time slot.
        #etc...
    ]
    self.vip.rpc.call('platform.actuator', # Target agent
                     'request_new_schedule', # Method to call
                     agent_id,              # Requestor
                     "some task",          # TaskID
                     "LOW",                # Priority
                     msg).get(timeout=10)  # Request_
↔message
```

Add the call back method for the schedule request:

```
def on_schedule_result(self, topic, headers, message, match):
    print ("TestAgent schedule result \nTopic: {topic}, "
          "{headers}, Message: {message}")
          .format(topic=topic, headers=headers, message=message))
```

Full TestAgent Code

The following is the full TestAgent code built in the previous steps:

```
from __future__ import absolute_import

from datetime import datetime
import logging
import sys

from volttron.platform.vip.agent import Agent, Core
```

(continues on next page)

```
from volttron.platform.agent import utils
from volttron.platform.messaging import headers as headers_mod

utils.setup_logging()
_log = logging.getLogger(__name__)

class TestAgent(Agent):
    def __init__(self, config_path, **kwargs):
        super(TestAgent, self).__init__(**kwargs)

    @Core.receiver("onstart")
    def starting(self, sender, **kwargs):
        '''
        Subscribes to the platform message bus on
        the heartbeat/listeneragent topic
        '''
        _log.info('TestAgent example agent start-up function')
        self.vip.pubsub.subscribe(peer='pubsub', topic='heartbeat/listeneragent',
                                   callback=self.on_heartbeat)
        self.vip.pubsub.subscribe('pubsub', topics.ACTUATOR_RESPONSE,
                                   callback=self.on_schedule_result)
        self.vip.pubsub.subscribe('pubsub', 'weather/temperature/temp_f',
                                   callback=self.on_weather)

        self.publish_schedule()

    def on_heartbeat(self, peer, sender, bus, topic, headers, message):
        '''TestAgent callback method'''
        _log.info('Matched topic: {}, for bus: {}'.format(topic, bus))
        self.vip.pubsub.publish(peer='pubsub',
                                topic='testagent/publish',
                                headers=headers,
                                message='test publishing').get(timeout=30)

    def on_weather(self, peer, sender, bus, topic, headers, message):
        _log.info(
            "TestAgent got weather\nTopic: {}, Message: {}".format(topic, message))

    def on_schedule_result(self, topic, headers, message, match):
        print(("TestAgent schedule result \nTopic: {topic}, "
              "{headers}, Message: {message}")
              .format(topic=topic, headers=headers, message=message))

def main(argv=sys.argv):
    '''Main method called by the eggsecutable.'''
    try:
        utils.vip_main(TestAgent)
    except Exception as e:
        _log.info(e)

if __name__ == '__main__':
    # Entry point for script
    sys.exit(main())
```

Message Debugging

VOLTRON agent messages are routed over the VOLTRON message bus. The Message Debugger Agent provides enhanced examination of this message stream's contents as an aid to debugging and troubleshooting agents and drivers.

When enabled, the Message Debugger Agent captures and records each message as it is routed. A second process, Message Viewer, provides a user interface that optimizes and filters the resulting data stream, either in real time or retrospectively, and displays its contents.

The Message Viewer can convey information about high-level interactions among VOLTRON agents, representing the message data as conversations that can be filtered and/or expanded. A simple RPC call involving 4 individual message send/receive segments can be displayed as a single row, which can then be expanded to drill down into the message details. This results in a higher-level, easier-to-obtain view of message bus activity than might be gleaned by using `grep` on verbose log files.

Pub/Sub interactions can be summarized by topic, including counts of messages published during a given capture period by sender, receiver and topic.

Another view displays the most-recently-published message, or message exchange, that satisfies the current filter criteria, continuously updated as new messages are routed.

Enabling the Message Debugger

In order to use the Message Debugger, two steps are required:

- VOLTRON must have been started with a `--msgdebug` command line option.
- The Message Debugger Agent must be running.

When VOLTRON has been started with `--msgdebug`, its Router publishes each message to an IPC socket for which the Message Debugger Agent is a subscriber. This is kept disabled by default because it consumes a significant quantity of CPU and memory resources, potentially affecting VOLTRON timing and performance. So as a general rule, the `--msgdebug` option should be employed during development/debugging only, and should not be left enabled in a production environment.

Example of starting VOLTRON with the `--msgdebug` command line option:

```
(volttron) volttron -vv -l log1 ``--msgdebug``
```

If VOLTRON is running in this mode, the stream of routed messages is available to a subscribing Message Debugger Agent. It can be started from `volttron-ctl` in the same fashion as other agents, for example:

```
(volttron) $ volttron-ctl status
AGENT                                IDENTITY                                TAG                                STATUS
fd listeneragent-3.2                 listener                                listener
08 messagedebuggeragent-0.1          platform.messagedebugger               platform.messagedebugger
e1 vcplatformagent-3.5.4             platform.agent                          vcp
47 volttroncentralagent-3.5.5        volttron.central                       vc

(volttron) $ volttron-ctl start 08
Starting 089c53f0-f225-4608-aecb-3e86e0df30eb messagedebuggeragent-0.1

(volttron) $ volttron-ctl status
AGENT                                IDENTITY                                TAG                                STATUS
fd listeneragent-3.2                 listener                                listener
08 messagedebuggeragent-0.1          platform.messagedebugger               platform.messagedebugger_
↪running [43498]
```

(continues on next page)

(continued from previous page)

e1	vcplatformagent-3.5.4	platform.agent	vcp
47	volttroncentralagent-3.5.5	volttron.central	vc

See *Agent Creation Walkthrough* for further details on installing and starting agents from `volttron-ctl`.

Once the Message Debugger Agent is running, it begins capturing message data and writing it to a SQLite database.

Message Viewer

The Message Viewer is a separate process that interacts with the Message Debugger Agent primarily via VOLTRON RPC calls. These calls allow it to request and report on filtered sets of message data.

Since the Agent's RPC methods are available for use by any VOLTRON agent, the Message Viewer is really just one example of a Message Debugger information consumer. Other viewers could be created to satisfy a variety of specific debugging needs. For example, a viewer could support browser-based message debugging with a graphical user interface, or a viewer could transform message data into PCAP format for consumption by WireShark.

The Message Viewer in `services/ops/MessageDebuggerAgent/messageviewer/viewer.py` implements a command-line UI, subclassing Python's `Cmd` class. Most of the command-line options that it displays result in a MessageDebugger-Agent RPC request. The Message Viewer formats and displays the results.

In Linux, the Message Viewer can be started as follows, and displays the following menu:

```
(volttron) $ cd services/ops/MessageDebuggerAgent/messageviewer
(volttron) $ python viewer.py
Welcome to the MessageViewer command line. Supported commands include:
  display_message_stream
  display_messages
  display_exchanges
  display_exchange_details
  display_session_details_by_agent <session_id>
  display_session_details_by_topic <session_id>

  list_sessions
  set_verbosity <level>
  list_filters
  set_filter <filter_name> <value>
  clear_filters
  clear_filter <filter_name>

  start_streaming
  stop_streaming
  start_session
  stop_session
  delete_session <session_id>
  delete_database

  help
  quit
Please enter a command.
Viewer>
```

Command-Line Help

The Message Viewer offers two help levels. Simply typing `help` gives a list of available commands. If a command name is provided as an argument, advice is offered on how to use that command:

```
Viewer> help

Documented commands (type help <topic>):
=====
clear_filter          display_messages          set_filter
clear_filters         display_session_details_by_agent set_verbosity
delete_database       display_session_details_by_topic start_session
delete_session        help                      start_streaming
display_exchange_details list_filters              stop_session
display_exchanges     list_sessions            stop_streaming
display_message_stream quit

Viewer> help set_filter

        Set a filter to a value; syntax is: set_filter <filter_name> <value>

        Some recognized filters include:
        . freq <n>: Use a single-line display, refreshing every <n> seconds (<n>
↳can be floating point)
        . session_id <n>: Display Messages and Exchanges for the indicated
↳debugging session ID only
        . results_only <n>: Display Messages and Exchanges only if they have a
↳result
        . sender <agent_name>
        . recipient <agent_name>
        . device <device_name>
        . point <point_name>
        . topic <topic_name>: Matches all topics that start with the supplied
↳<topic_name>
        . starttime <YYYY-MM-DD HH:MM:SS>: Matches rows with timestamps after the
↳supplied time
        . endtime <YYYY-MM-DD HH:MM:SS>: Matches rows with timestamps before the
↳supplied time
        . (etc. -- see the structures of DebugMessage and DebugMessageExchange)
```

Debug Sessions

The Message Debugger Agent tags each message with a debug session ID (a serial number), which groups a set of messages that are bounded by a start time and an end time. The `list_sessions` command describes each session in the database:

```
Viewer> list_sessions
rowid      start_time                end_time                num_messages
  1         2017-03-20 17:07:13.867951 -                        2243
  2         2017-03-20 17:17:35.725224 -                        1320
  3         2017-03-20 17:33:35.103204 2017-03-20 17:46:15.657487 12388
```

A new session is started by default when the Agent is started. After that, the `stop_session` and `start_session` commands can be used to create new session boundaries. If the Agent is running but no session is active (i.e., because `stop_session` was used to stop it), messages are still written to the database, but they have no session ID.

Filtered Display

The `set_filter <property> <value>` command enables filtered display of messages. A variety of properties can be filtered.

In the following example, message filters are defined by `session_id` and `sender`, and the `display_messages` command displays the results:

```
Viewer> set_filter session_id 4
Set filters to {'session_id': '4'}
Viewer> set_filter sender testagent
Set filters to {'sender': 'testagent', 'session_id': '4'}
Viewer> display_messages
  timestamp      direction      sender      recipient      request_id      device
↪      subsystem      method      topic
↪point          result
  11:51:00      incoming      testagent    messageviewer.connection - -
↪      RPC          pubsub.sync - -
↪      -
  11:51:00      outgoing      testagent    pubsub          - -
↪      RPC          pubsub.push - -
↪      -
  11:51:00      incoming      testagent    platform.driver -
↪1197886248649056372.284581685 RPC      get_point
↪ chargepoint1 Status -
  11:51:01      outgoing      testagent    platform.driver -
↪1197886248649056372.284581685 RPC      -
↪ - - AVAILABLE
  11:51:01      incoming      testagent    pubsub          -
↪1197886248649056373.284581649 RPC      pubsub.publish test_topic/test_
↪subtopic - - -
  11:51:01      outgoing      testagent    pubsub          -
↪1197886248649056373.284581649 RPC      -
↪ - - None
```

Debug Message Exchanges

A VOLTRON message's request ID is not unique to a single message. A group of messages in an “exchange” (essentially a small conversation among agents) will often share a common request ID, for instance during RPC request/response exchanges.

The following example uses the same filters as above, and then uses `display_exchanges` to display a single line for each message exchange, reducing the number of displayed rows from 6 to 2. Note that not all messages have a request ID; messages with no ID are absent from the responses to exchange queries.

```
Viewer> list_filters
{'sender': 'testagent', 'session_id': '4'}
Viewer> display_exchanges
  sender      recipient      sender_time  topic      device
↪point          result
  testagent    platform.driver 11:51:00    -          chargepoint1
↪Status        AVAILABLE
  testagent    pubsub          11:51:01    test_topic/test_subtopic -
↪      None
```

Special Filters

Most filters that can be set with the `set_filter` command are simple string matches on one or another property of a message. Some filters have special characteristics, though. The `set_filter starttime <timestamp>` and `set_filter endtime <timestamp>` filters are inequalities that test for messages after a start time or before an end time.

In the following example, note the use of quotes in the `endtime` value supplied to `set_filter`. Any filter value can be delimited with quotes. Quotes must be used when a value contains embedded spaces, as is the case here:

```
Viewer> list_sessions
  rowid      start_time                end_time                num_messages
  1          2017-03-20 17:07:13.867951    -                      -
  2          2017-03-20 17:17:35.725224    -                      -
  3          2017-03-21 11:48:33.803288    2017-03-21 11:50:57.181136  6436
  4          2017-03-21 11:50:59.656693    2017-03-21 11:51:05.934895  450
  5          2017-03-21 11:51:08.431871    -                      74872
  6          2017-03-21 12:17:30.568260    -                      2331
Viewer> set_filter session_id 5
Set filters to {'session_id': '5'}
Viewer> set_filter sender testagent
Set filters to {'sender': 'testagent', 'session_id': '5'}
Viewer> set_filter endtime '2017-03-21 11:51:30'
Set filters to {'endtime': '2017-03-21 11:51:30', 'sender': 'testagent', 'session_id': '5'}
Viewer> display_exchanges
  sender      recipient      sender_time  topic                device      point_
↪point      result
  testagent   platform.driver  11:51:11    -                    chargepoint1 ↪
↪Status     AVAILABLE
  testagent   pubsub          11:51:11    test_topic/test_subtopic -          -
↪          None
  testagent   platform.driver  11:51:25    -                    chargepoint1 ↪
↪Status     AVAILABLE
  testagent   pubsub          11:51:25    test_topic/test_subtopic -          -
↪          None
  testagent   platform.driver  11:51:26    -                    chargepoint1 ↪
↪Status     AVAILABLE
  testagent   pubsub          11:51:26    test_topic/test_subtopic -          -
↪          None
```

Another filter type with special behavior is `set_filter topic <name>`. Ordinarily, filters do an exact match on a message property. Since message topics are often expressed as hierarchical substrings, though, the `topic` filter does a substring match on the left edge of a message's topic, as in the following example:

```
Viewer> set_filter topic test_topic
Set filters to {'topic': 'test_topic', 'endtime': '2017-03-21 11:51:30', 'sender': 'testagent', 'session_id': '5'}
Viewer> display_exchanges
  sender      recipient      sender_time  topic                device      point_
↪          result
  testagent   pubsub          11:51:11    test_topic/test_subtopic -          - ↪
↪          None
  testagent   pubsub          11:51:25    test_topic/test_subtopic -          - ↪
↪          None
  testagent   pubsub          11:51:26    test_topic/test_subtopic -          - ↪
↪          None
Viewer>
```

Another filter type with special behavior is `set_filter results_only 1`. In the JSON representation of a response to an RPC call, for example an RPC call to a Master Driver interface, the response to the RPC request typically appears as the value of a 'result' tag. The `results_only` filter matches only those messages that have a non-empty value for this tag.

In the following example, note that when the `results_only` filter is set, it is given a value of '1'. This is actually a meaningless value that gets ignored. It must be supplied because the `set_filter` command syntax requires that a value be supplied as a parameter.

In the following example, note the use of `clear_filter <property>` to remove a single named filter from the list of filters that are currently in effect. There is also a `clear_filters` command, which clears all current filters.

```
Viewer> clear_filter topic
Set filters to {'endtime': '2017-03-21 11:51:30', 'sender': 'testagent', 'session_id
↳': '5'}
Viewer> set_filter results_only 1
Set filters to {'endtime': '2017-03-21 11:51:30', 'sender': 'testagent', 'session_id
↳': '5', 'results_only': '1'}
Viewer> display_exchanges
  sender      recipient      sender_time  topic      device      point      ↳
↳result
  testagent   platform.driver  11:51:11    -          chargepoint1  Status    ↳
↳AVAILABLE
  testagent   platform.driver  11:51:25    -          chargepoint1  Status    ↳
↳AVAILABLE
  testagent   platform.driver  11:51:26    -          chargepoint1  Status    ↳
↳AVAILABLE
```

Streamed Display

In addition to exposing a set of RPC calls that allow other agents (like the Message Viewer) to query the Message Debugger Agent's SQLite database of recent messages, the Agent can also publish messages in real time as it receives them.

This feature is disabled by default due to the large quantity of data that it might need to handle. When it is enabled, the Agent applies the filters currently in effect to each message as it is received, and re-publishes the transformed, ready-for-debugging message to a socket if it meets the filter criteria. The Message Viewer can listen on that socket and display the message stream as it arrives.

In the following `display_message_stream` example, the Message Viewer displays all messages sent by the agent named 'testagent', as they arrive. It continues to display messages until execution is interrupted with `ctrl-C`:

```
Viewer> clear_filters
Set filters to {}
Viewer> set_filter sender testagent
Set filters to {'sender': 'testagent'}
Viewer> display_message_stream
Streaming debug messages
  timestamp    direction    sender      recipient    request_id  subsystem  ↳
↳method      topic      device      point      result
  12:28:58     outgoing    testagent   pubsub       -           RPC        ↳
↳pubsub.push -          -          -           -
  12:28:58     incoming    testagent   platform.dr  11978862486  RPC        ↳
↳point      -          chargepoint Status      -
↳          iver      49056826.28
↳          1          4581713
```

(continues on next page)

(continued from previous page)

```

12:28:58    outgoing    testagent    platform.dr  11978862486  RPC    -    ↵
↪          -          -          -          AVAILABLE
          iver          49056826.28
          4581713
12:28:58    incoming    testagent    pubsub      11978862486  RPC    ↵
↪pubsub.publ test_topic/  -          -          49056827.28    ish    ↵
↪          test_subtop          4581685          ↵
↪          ic          12:28:58    outgoing    testagent    pubsub      11978862486  RPC    -    ↵
↪          -          -          -          None
          49056827.28
          4581685
12:28:58    outgoing    testagent    pubsub      -          RPC    ↵
↪pubsub.push -          -          -          -
^CViewer> stop_streaming
Stopped streaming debug messages

```

(Note the use of wrapping in the column formatting. Since these messages aren't known in advance, the Message Viewer has incomplete information about how wide to make each column. Instead, it must make guesses based on header widths, data widths in the first row received, and min/max values, and then wrap the data when it overflows the column boundaries.)

Single-Line Display

Another filter with special behavior is `set_filter freq <seconds>`. This filter, which takes a number *N* as its value, displays only one row, the most recently captured row that satisfies the filter criteria. (Like other filters, this filter can be used with either `display_messages` or `display_exchanges`.) It then waits *N* seconds, reissues the query, and overwrites the old row with the new one. It continues this periodic single-line overwritten display until it is interrupted with `ctrl-C`:

```

Viewer> list_filters
{'sender': 'testagent'}
Viewer> set_filter freq 10
Set filters to {'freq': '10', 'sender': 'testagent'}
Viewer> display_exchanges
 sender      recipient      sender_time  topic          device      point_
↪          result
 testagent   pubsub        12:31:28    test_topic/test_subtopic -          -          ↵
↪          None

```

(Again, the data isn't known in advance, so the Message Viewer has to guess the best width of each column. In this single-line display format, data gets truncated if it doesn't fit, because no wrapping can be performed – only one display line is available.)

Displaying Exchange Details

The `display_exchange_details <request_id>` command provides a way to get more specific details about an exchange, i.e. about all messages that share a common request ID. At low or medium verbosity, when this command is used (supplying the relevant request ID, which can be obtained from the output of other commands), it displays one row for each message:

```

Viewer> set_filter sender testagent
Set filters to {'sender': 'testagent', 'session_id': '4'}
Viewer> display_messages
  timestamp      direction      sender      recipient      request_id      device
↪      subsystem      method      topic
↪point          result
  11:51:00      incoming      testagent    messageviewer.connection - -
↪      RPC          pubsub.sync  - - -
↪
  11:51:00      outgoing      testagent    pubsub          - -
↪      RPC          pubsub.push  - - -
↪
  11:51:00      incoming      testagent    platform.driver
↪1197886248649056372.284581685 RPC      get_point      - -
↪ chargepoint1 Status -
  11:51:01      outgoing      testagent    platform.driver
↪1197886248649056372.284581685 RPC      - -
↪ - AVAILABLE
  11:51:01      incoming      testagent    pubsub
↪1197886248649056373.284581649 RPC      pubsub.publish test_topic/test_
↪subtopic - - -
  11:51:01      outgoing      testagent    pubsub
↪1197886248649056373.284581649 RPC      - -
↪ - None
Viewer> display_exchange_details 1197886248649056373.284581649
timestamp      direction      sender      recipient      request_id      device      point
↪subsystem      method      topic
↪result
  11:51:01      incoming      testagent    pubsub          1197886248649056373.284581649
↪RPC          pubsub.publish test_topic/test_subtopic - -
  11:51:01      outgoing      testagent    pubsub          1197886248649056373.284581649
↪RPC          - - -
↪None

```

At high verbosity, `display_exchange_details` switches display formats, showing all properties for each message in a json-like dictionary format:

```

Viewer> set_verbosity high
Set verbosity to high
Viewer> display_exchange_details 1197886248649056373.284581649
{
  "data": "{\"params\":{\"topic\":\"test_topic/test_subtopic\",\"headers\":{\"Date\
↪\": \"2017-03-21T11:50:56.293830\", \"max_compatible_version\": \"\", \"min_compatible_
↪version\": \"3.0\"}, \"message\": [{\"property_1\": 1, \"property_2\": 2}, {\"property_3\
↪\": 3, \"property_4\": 4}], \"bus\": \"\", \"jsonrpc\": \"2.0\", \"method\": \"pubsub.
↪publish\", \"id\": \"15828311332408898779.284581649\"}},
  "device": "",
  "direction": "incoming",
  "frame7": "",
  "frame8": "",
  "frame9": "",
  "headers": "{u'Date': u'2017-03-21T11:50:56.293830', u'max_compatible_version': u'
↪', u'min_compatible_version': u'3.0'}",
  "message": "[{u'property_1': 1, u'property_2': 2}, {u'property_3': 3, u'property_4
↪': 4}]",
  "message_size": 374,

```

(continues on next page)

(continued from previous page)

```

    "message_value": "{u'property_1': 1, u'property_2': 2}",
    "method": "pubsub.publish",
    "params": "{u'topic': u'test_topic/test_subtopic', u'headers': {u'Date': u'2017-
↪03-21T11:50:56.293830', u'max_compatible_version': u'', u'min_compatible_version': u
↪'3.0'}, u'message': [{u'property_1': 1, u'property_2': 2}, {u'property_3': 3, u
↪'property_4': 4}], u'bus': u''}",
    "point": "",
    "point_value": "",
    "recipient": "pubsub",
    "request_id": "1197886248649056373.284581649",
    "result": "",
    "sender": "testagent",
    "session_id": 4,
    "subsystem": "RPC",
    "timestamp": "2017-03-21 11:51:01.027623",
    "topic": "test_topic/test_subtopic",
    "user_id": "",
    "vip_signature": "VIP1"
}

{
  "data": "{\\"params\\":{\\"topic\\":\\"test_topic/test_subtopic\\",\\"headers\\":{\\"Date\\
↪\\":\\"2017-03-21T11:50:56.293830\\",\\"max_compatible_version\\":\\"\\",\\"min_compatible_
↪version\\":\\"3.0\\"},\\"message\\":[{\\"property_1\\":1,\\"property_2\\":2},{\\"property_3\\
↪\\":3,\\"property_4\\":4}],\\"bus\\":\\"\\",\\"jsonrpc\\":\\"2.0\\",\\"method\\":\\"pubsub.
↪publish\\",\\"id\\":\\"15828311332408898779.284581649\\"}",
  "device": "",
  "direction": "outgoing",
  "frame7": "",
  "frame8": "",
  "frame9": "",
  "headers": "{u'Date': u'2017-03-21T11:50:56.293830', u'max_compatible_version': u'
↪', u'min_compatible_version': u'3.0'}",
  "message": "{[u'property_1': 1, u'property_2': 2}, {u'property_3': 3, u'property_4
↪': 4}]",
  "message_size": 383,
  "message_value": "{u'property_1': 1, u'property_2': 2}",
  "method": "pubsub.publish",
  "params": "{u'topic': u'test_topic/test_subtopic', u'headers': {u'Date': u'2017-
↪03-21T11:50:56.293830', u'max_compatible_version': u'', u'min_compatible_version': u
↪'3.0'}, u'message': [{u'property_1': 1, u'property_2': 2}, {u'property_3': 3, u
↪'property_4': 4}], u'bus': u''}",
  "point": "",
  "point_value": "",
  "recipient": "testagent",
  "request_id": "1197886248649056373.284581649",
  "result": "",
  "sender": "pubsub",
  "session_id": 4,
  "subsystem": "RPC",
  "timestamp": "2017-03-21 11:51:01.031183",
  "topic": "test_topic/test_subtopic",
  "user_id": "testagent",
  "vip_signature": "VIP1"
}

```

Verbosity

As mentioned in the previous section, Agent and Viewer behavior can be adjusted by changing the current verbosity with the `set_verbosity <level>` command. The default verbosity is low. low, medium and high levels are available:

```
Viewer> set_verbosity high
Set verbosity to high
Viewer> set_verbosity none
Invalid verbosity choice none; valid choices are ['low', 'medium', 'high']
```

At high verbosity, the following query formatting rules are in effect:

- When displaying timestamps, display the full date and time (including microseconds), not just HH:MM:SS.
- In responses to `display_message_exchanges`, use dictionary format (see example in previous section).
- Display all columns, not just “interesting” columns (see the list below).
- Don’t exclude messages/exchanges based on excluded senders/receivers (see the list below).

At medium or low verbosity:

- When displaying timestamps, display HH:MM:SS only.
- In responses to `display_message_exchanges`, use table format.
- Display “interesting” columns only (see the list below).
- Exclude messages/exchanges for certain senders/receivers (see the list below).

At low verbosity:

- If > 1000 objects are returned by a query, display the count only.

The following “interesting” columns are displayed at low and medium verbosity levels (at high verbosity levels, all properties are displayed):

Debug Message	Debug Message Exchange	Debug Session
timestamp	sender_time	rowid
direction		start_time
sender	sender	end_time
recipient	recipient	num_messages
request_id		
subsystem		
method		
topic	topic	
device	device	
point	point	
result	result	

Messages from the following senders, or to the following receivers, are excluded at low and medium verbosity levels:

Sender	Receiver
(empty)	(empty)
None	
control	control
config.store	config.store
pubsub	

(continues on next page)

(continued from previous page)

```
control.connection
messageviewer.connection
platform.messagedebugger
platform.messagedebugger.loopback_rpc
```

These choices about which columns are “interesting” and which senders/receivers are excluded are defined as parameters in Message Viewer, and can be adjusted as necessary by changing global value lists in viewer.py.

Session Statistics

One useful tactic for starting at a summary level and drilling down is to capture a set of messages for a session and then examine the counts of sending and receiving agents, or sending agents and topics. This gives hints on which values might serve as useful filters for more specific queries.

The `display_session_details_by_agent <session_id>` command displays statistics by sending and receiving agent. Sending agents are table columns, and receiving agents are table rows. This query also applies whatever filters are currently in effect; the filters can reduce the counts and can also reduce the number of columns and rows.

The following example shows the command being used to list all senders and receivers for messages sent during debug session 7:

```
Viewer> list_sessions
  rowid      start_time                end_time                num_messages
  1          2017-03-20 17:07:13.867951 - -
  2          2017-03-20 17:17:35.725224 - -
  3          2017-03-21 11:48:33.803288 2017-03-21 11:50:57.181136 6436
  4          2017-03-21 11:50:59.656693 2017-03-21 11:51:05.934895 450
  5          2017-03-21 11:51:08.431871 - 74872
  6          2017-03-21 12:17:30.568260 2017-03-21 12:38:29.070000 60384
  7          2017-03-21 12:38:31.617099 2017-03-21 12:39:53.174712 3966

Viewer> clear_filters
Set filters to {}
Viewer> display_session_details_by_agent 7
  Receiving Agent      control      listener      messageviewer.connection
  ↳platform.driver    platform.messagedebugger      pubsub      testagent
  (No Receiving Agent) - - - 2
  ↳ control - - - -
  ↳ listener - 2 - -
  ↳ messageviewer.connection - 679 - -
  ↳ platform.driver - 3 - -
  ↳ platform.messagedebugger - - 1249 16
  ↳ pubsub - 2 679 -
  ↳ 1249 - - 4 31
  ↳ testagent - - - -
  ↳ 16 - 31 -
```

The `display_session_details_by_topic <session_id>` command is similar to `display_session_details_by_agent`, but each row contains statistics for a topic instead of for a receiving agent:

```

Viewer> display_session_details_by_topic 7
Topic                control      listener  messageviewer.
↳connection platform.driver platform.messagedebugger      pubsub      testagent
(No Topic)
↳ 5                640                3                1314                39
devices/chargepoint1/Address
↳ -                6                -                6                -
devices/chargepoint1/City
↳ -                6                -                6                -
devices/chargepoint1/Connector
↳ -                5                -                5                -
devices/chargepoint1/Country
↳ -                5                -                5                -
devices/chargepoint1/Current
↳ -                6                -                6                -
devices/chargepoint1/Description
↳ -                6                -                6                -
devices/chargepoint1/Energy
↳ -                5                -                5                -
devices/chargepoint1/Lat
↳ -                6                -                6                -
devices/chargepoint1/Level
↳ -                5                -                5                -
devices/chargepoint1/Long
↳ -                6                -                6                -
devices/chargepoint1/Mode
↳ -                5                -                5                -
devices/chargepoint1/Power
↳ -                6                -                6                -
devices/chargepoint1/Reservable
↳ -                5                -                5                -
devices/chargepoint1/State
↳ -                6                -                6                -
devices/chargepoint1/Status
↳ -                5                -                5                -
devices/chargepoint1/Status.TimeSta
↳ -                6                -                6                -
mp
devices/chargepoint1/Type
↳ -                6                -                6                -
devices/chargepoint1/Voltage
↳ -                5                -                5                -
devices/chargepoint1/alarmTime
↳ -                6                -                6                -
devices/chargepoint1/alarmType
↳ -                6                -                6                -
devices/chargepoint1/all
↳ -                5                -                5                -
devices/chargepoint1/allowedLoad
↳ -                6                -                6                -
devices/chargepoint1/clearAlarms
↳ -                6                -                6                -
devices/chargepoint1/currencyCode
↳ -                6                -                6                -
devices/chargepoint1/driverAccountN
↳ -                5                -                5                -
umber

```

(continues on next page)

(continued from previous page)

devices/chargepoint1/driverName	-	-	-	-	↳
↳ -	5	-	5	-	
devices/chargepoint1/endTime	-	-	-	-	↳
↳ -	5	-	5	-	
devices/chargepoint1/mainPhone	-	-	-	-	↳
↳ -	6	-	6	-	
devices/chargepoint1/maxPrice	-	-	-	-	↳
↳ -	5	-	5	-	
devices/chargepoint1/minPrice	-	-	-	-	↳
↳ -	5	-	5	-	
devices/chargepoint1/numPorts	-	-	-	-	↳
↳ -	6	-	6	-	
devices/chargepoint1/orgID	-	-	-	-	↳
↳ -	5	-	5	-	
devices/chargepoint1/organizationName	-	-	-	-	↳
↳ -	5	-	5	-	
me					
devices/chargepoint1/percentShed	-	-	-	-	↳
↳ -	6	-	6	-	
devices/chargepoint1/portLoad	-	-	-	-	↳
↳ -	6	-	6	-	
devices/chargepoint1/portNumber	-	-	-	-	↳
↳ -	6	-	6	-	
devices/chargepoint1/sessionID	-	-	-	-	↳
↳ -	5	-	5	-	
devices/chargepoint1/sessionTime	-	-	-	-	↳
↳ -	6	-	6	-	
devices/chargepoint1/sgID	-	-	-	-	↳
↳ -	6	-	6	-	
devices/chargepoint1/sgName	-	-	-	-	↳
↳ -	6	-	6	-	
devices/chargepoint1/shedState	-	-	-	-	↳
↳ -	5	-	5	-	
devices/chargepoint1/startTime	-	-	-	-	↳
↳ -	6	-	6	-	
devices/chargepoint1/stationID	-	-	-	-	↳
↳ -	5	-	5	-	
devices/chargepoint1/stationMacAddr	-	-	-	-	↳
↳ -	6	-	6	-	
devices/chargepoint1/stationManufacturer	-	-	-	-	↳
↳ -	5	-	5	-	
turer					
devices/chargepoint1/stationModel	-	-	-	-	↳
↳ -	6	-	6	-	
devices/chargepoint1/stationName	-	-	-	-	↳
↳ -	5	-	5	-	
devices/chargepoint1/stationRightsProfile	-	-	-	-	↳
↳ -	6	-	6	-	
rofile					
devices/chargepoint1/stationSerialNumber	-	-	-	-	↳
↳ -	6	-	6	-	
um					
heartbeat/control		1	-	-	↳
↳ -	-	-	1	-	
heartbeat/listener		-	15	-	↳
↳ -	-	-	15	-	
heartbeat/platform.driver		-	-	-	↳
↳ -	1	-	1	-	

(continues on next page)

(continued from previous page)

heartbeat/pubsub	-	-	-	-
↔ -	-	2	-	↳
test_topic/test_subtopic	-	-	-	-
↔ -	-	8	8	↳

Database Administration

The Message Debugger Agent stores message data in a SQLite database's DebugMessage, DebugMessageExchange and DebugSession tables. If the database isn't present already when the Agent is started, it is created automatically.

The SQLite database can consume a lot of disk space in a relatively short time, so the Message Viewer has command-line options that recover that space by deleting the database or by deleting all messages belonging to a given debug session.

The `delete_session <session_id>` command deletes the database's DebugSession row with the indicated ID, and also deletes all DebugMessage and DebugMessageExchange rows with that session ID. In the following example, `delete_session` deletes the 60,000 DebugMessages that were captured during a 20-minute period as session 6:

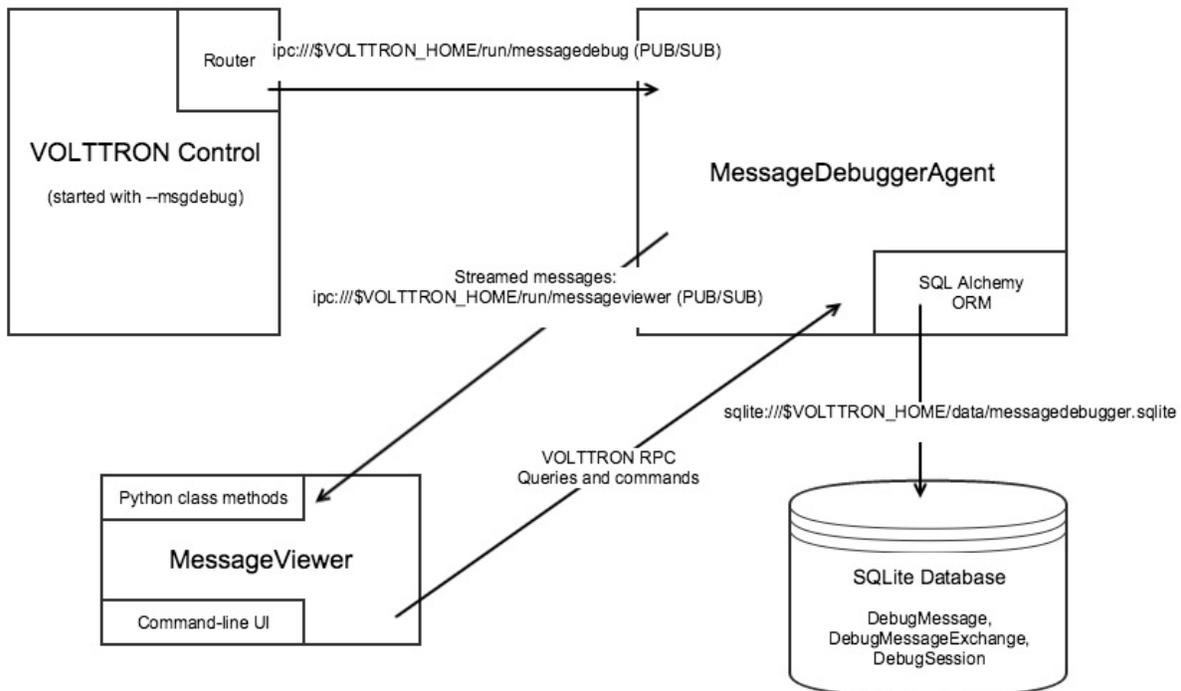
```
Viewer> list_sessions
rowid      start_time          end_time           num_messages
1          2017-03-20 17:07:13.867951 - -
2          2017-03-20 17:17:35.725224 - -
3          2017-03-21 11:48:33.803288 2017-03-21 11:50:57.181136 6436
4          2017-03-21 11:50:59.656693 2017-03-21 11:51:05.934895 450
5          2017-03-21 11:51:08.431871 - 74872
6          2017-03-21 12:17:30.568260 2017-03-21 12:38:29.070000 60384
7          2017-03-21 12:38:31.617099 2017-03-21 12:39:53.174712 3966
8          2017-03-21 12:42:08.482936 - 3427
Viewer> delete_session 6
Deleted debug session 6
Viewer> list_sessions
rowid      start_time          end_time           num_messages
1          2017-03-20 17:07:13.867951 - -
2          2017-03-20 17:17:35.725224 - -
3          2017-03-21 11:48:33.803288 2017-03-21 11:50:57.181136 6436
4          2017-03-21 11:50:59.656693 2017-03-21 11:51:05.934895 450
5          2017-03-21 11:51:08.431871 - 74872
7          2017-03-21 12:38:31.617099 2017-03-21 12:39:53.174712 3966
8          2017-03-21 12:42:08.482936 - 4370
```

The `delete_database` command deletes the entire SQLite database, removing all records of previously-captured DebugMessages, DebugMessageExchanges and DebugSessions. The database will be re-created the next time a debug session is started.

```
Viewer> delete_database
Database deleted
Viewer> list_sessions
No query results
Viewer> start_session
Message debugger session 1 started
Viewer> list_sessions
rowid      start_time          end_time           num_messages
1          2017-03-22 12:39:40.320252 - 180
```

It's recommended that the database be deleted if changes are made to the `DebugMessage`, `DebugMessageExchange` or `DebugSession` object structures that are defined in `agent.py`. A skew between these data structures in Python code vs. the ones in the database can cause instability in the Message Debugger Agent, perhaps causing it to fail. If a failure of this kind prevents use of the Message Viewer's `delete_database` command, the database can be deleted directly from the filesystem. By default, it is located in `$VOLTTTRON_HOME`'s `run` directory.

Implementation Details



Router changes: `MessageDebuggerAgent` reads and stores all messages that pass through the VIP router. This is accomplished by subscribing to the messages on a new socket published by the platform's `Router.issue()` method.

The “direction” property: Most agent interactions result in at least two messages, an incoming request and an outgoing response. `Router.issue()` has a `topic` parameter with values `INCOMING`, `OUTGOING`, `ERROR` and `UNROUTABLE`. The publication on the socket that happens in `issue()` includes this “issue topic” (not to be confused with a message's `topic`) along with each message. `MessageDebuggerAgent` records it as a `DebugMessage` property called `direction`, since its value for almost all messages is either `INCOMING` or `OUTGOING`.

SQLite Database and SQL Alchemy: `MessageDebuggerAgent` records each message as a `DebugMessage` row in a relational database. `SQLite` is used since it's packaged with Python and is already being used by other VOLTTTRON agents. Database semantics are kept simple through the use of a `SQL Alchemy` object-relational mapping framework. Python's “`SQLAlchemy`” plug-in must be loaded in order for `MessageDebuggerAgent` to run.

Calling MessageViewer Directly: The `viewer.py` module that starts the Message Viewer command line also contains a `MessageViewer` class. It exposes class methods which can be used to make direct Python calls that, in turn, make Message Debugger Agent's RPC calls. The `MessageViewer` class-method API includes the following calls:

- `delete_debugging_db()`
- `delete_debugging_session(session_id)`

- `disable_message_debugging()`
- `display_db_objects(db_object_name, filters=None)`
- `display_message_stream()`
- `enable_message_debugging()`
- `message_exchange_details(message_id)`
- `session_details_by_agent(session_id)`
- `session_details_by_topic(session_id)`
- `set_filters(filters)`
- `set_verbosity(verbosity_level)`
- `start_streaming(filters=None)`
- `stop_streaming()`

The command-line UI's `display_messages` and `display_exchanges` commands are implemented here as `display_db_objects('DebugMessage')` and `display_db_objects(DebugMessageExchange)`. These calls return json-encoded representations of `DebugMessages` and `DebugMessageExchanges`, which are formatted for display by `MessageViewerCmd`.

MessageViewer connection: `MessageViewer` is not actually a VOLTRON agent. In order for it make `MessageDebuggerAgent` RPC calls, which are agent-agent interactions, it builds a “connection” that manages a temporary agent. This is a standard VOLTRON pattern that is also used, for instance, by `Voltron Central`.

TestAgent Source Code

Full code of agent detailed in `AgentDevelopment`:

```
"""
Agent documentation goes here.
"""

__docformat__ = 'reStructuredText'

import logging
import sys
from voltron.platform.agent import utils
from voltron.platform.vip.agent import Agent, Core, RPC

_log = logging.getLogger(__name__)
utils.setup_logging()
__version__ = "0.5"

def tester(config_path, **kwargs):
    """Parses the Agent configuration and returns an instance of
    the agent created using that configuration.

    :param config_path: Path to a configuration file.

    :type config_path: str
    :returns: Tester
    :rtype: Tester
    """
```

(continues on next page)

(continued from previous page)

```

try:
    config = utils.load_config(config_path)
except StandardError:
    config = {}

if not config:
    _log.info("Using Agent defaults for starting configuration.")

    setting1 = int(config.get('setting1', 1))
    setting2 = config.get('setting2', "some/random/topic")

return Tester(setting1,
                setting2,
                **kwargs)

class Tester(Agent):
    """
    Document agent constructor here.
    """

    def __init__(self, setting1=1, setting2="some/random/topic",
                **kwargs):
        super(Tester, self).__init__(**kwargs)
        _log.debug("vip_identity: " + self.core.identity)

        self.setting1 = setting1
        self.setting2 = setting2

        self.default_config = {"setting1": setting1,
                               "setting2": setting2}

        #Set a default configuration to ensure that self.configure is called
        ↪ immediately to setup
        #the agent.
        self.vip.config.set_default("config", self.default_config)
        #Hook self.configure up to changes to the configuration file "config".
        self.vip.config.subscribe(self.configure, actions=["NEW", "UPDATE"], pattern=
        ↪ "config")

    def configure(self, config_name, action, contents):
        """
        Called after the Agent has connected to the message bus. If a configuration
        ↪ exists at startup
        this will be called before onstart.

        Is called every time the configuration in the store changes.
        """
        config = self.default_config.copy()
        config.update(contents)

        _log.debug("Configuring Agent")

        try:
            setting1 = int(config["setting1"])
            setting2 = str(config["setting2"])

```

(continues on next page)

(continued from previous page)

```

except ValueError as e:
    _log.error("ERROR PROCESSING CONFIGURATION: {}".format(e))
    return

self.setting1 = setting1
self.setting2 = setting2

self._create_subscriptions(self.setting2)

def _create_subscriptions(self, topic):
    #Unsubscribe from everything.
    self.vip.pubsub.unsubscribe("pubsub", None, None)

    self.vip.pubsub.subscribe(peer='pubsub',
                               prefix=topic,
                               callback=self._handle_publish)

def _handle_publish(self, peer, sender, bus, topic, headers,
                    message):
    pass

@Core.receiver("onstart")
def onstart(self, sender, **kwargs):
    """
    This is method is called once the Agent has successfully connected to the_
↪platform.
    This is a good place to setup subscriptions if they are not dynamic or
    do any other startup activities that require a connection to the message bus.
    Called after any configurations methods that are called at startup.

    Usually not needed if using the configuration store.
    """
    #Example publish to pubsub
    #self.vip.pubsub.publish('pubsub', "some/random/topic", message="HI!")

    #Exmample RPC call
    #self.vip.rpc.call("some_agent", "some_method", arg1, arg2)

@Core.receiver("onstop")
def onstop(self, sender, **kwargs):
    """
    This method is called when the Agent is about to shutdown, but before it_
↪disconnects from
    the message bus.
    """
    pass

@RPC.export
def rpc_method(self, arg1, arg2, kwarg1=None, kwarg2=None):
    """
    RPC method

    May be called from another agent via self.core.rpc.call """
    return self.setting1 + arg1 - arg2

def main():
    """Main method called to start the agent."""

```

(continues on next page)

(continued from previous page)

```

utils.vip_main(tester,
                version=__version__)

if __name__ == '__main__':
    # Entry point for script
    try:
        sys.exit(main())
    except KeyboardInterrupt:
        pass

```

Contents of setup.py for TestAgent:

```

from setuptools import setup, find_packages

MAIN_MODULE = 'agent'

# Find the agent package that contains the main module
packages = find_packages('.')
agent_package = 'tester'

# Find the version number from the main module
agent_module = agent_package + '.' + MAIN_MODULE
_temp = __import__(agent_module, globals(), locals(), ['__version__'], -1)
__version__ = _temp.__version__

# Setup
setup(
    name=agent_package + 'agent',
    version=__version__,
    author_email="volttron@pnnl.gov",
    url="https://volttron.org/",
    description="Agent development tutorial.",
    author="VOLTRON Team",
    install_requires=['volttron'],
    packages=packages,
    entry_points={
        'setuptools.installation': [
            'eggsecutable = ' + agent_module + ':main',
        ]
    }
)

```

Contents of config:

```

{
    # VOLTRON config files are JSON with support for python style comments.
    "setting1": 2, #Integers
    "setting2": "some/random/topic2", #strings
    "setting3": true, #Booleans: remember that in JSON true and false are not
    ↪capitalized.
    "setting4": false,
    "setting5": 5.1, #Floating point numbers.
    "setting6": [1,2,3,4], # Lists
    "setting7": {"setting7a": "a", "setting7b": "b"} #Objects
}

```

3.4.2 Deployment Advice

Deployment Options

There are several ways to deploy the VOLTTTRON platform in a Linux environment. It is up to the user to determine which is right for them. The following assumes that the platform has already been bootstrapped and is ready to run.

Simple Command Line

With the VOLTTTRON environment activated the platform can be started simply by running VOLTTTRON on the command line.

```
$volttron -vv
```

This will start the platform in the current terminal with very verbose logging turned on. This is most appropriate for testing Agents or testing a deployment for problems before switching to a more long term solution. This will print all log messages to the console in real time.

This should not be used for long term deployment. As soon as an SSH session is terminated for whatever reason the processes attached to that session will be killed. This also will not capture log message to a file.

Running VOLTTTRON as a Background Process

A simple, more long term solution, is to run volttron in the background and disown it from the current terminal.

Warning: If you plan on running VOLTTTRON in the background and detaching it from the terminal with the `disown` command be sure to redirect `stderr` and `stdout` to `/dev/null`. Even if logging to a file is used some libraries which VOLTTTRON relies on output directly to `stdout` and `stderr`. This will cause problems if those file descriptors are not redirected to `/dev/null`.

```
$volttron -vv -l volttron.log > /dev/null 2>&1&

#If there are other jobs running in your terminal be sure to disown the correct one.
$jobs
[1]+  Running                  something else
[2]+  Running                  volttron -vv -l volttron.log > /dev/null 2>&1 &

#Disown VOLTTTRON
$disown %2
```

This will run the VOLTTTRON platform in the background and turn it into a daemon. The log output will be directed to a file called `volttron.log` in the current directory.

To keep the size of the log under control for more longer term deployments us the rotating log configuration file `examples/rotatinglog.py`.

```
$volttron -vv --log-config examples/rotatinglog.py > /dev/null 2>&1&
```

This will start a rotate the log file at midnight and limit the total log data to seven days worth.

The main downside to this approach is that the VOLTTTRON platform will not automatically resume if the system is restarted. It will need to be restarted manually after reboot.

Setting up VOLTRON as a System Service

Systemd

An example service file `scripts/admin/volttron.service` for systemd can be used as a starting point for setting up VOLTRON as a service. Note that as this will redirect all the output that would be going to stdout - to the syslog. This can be accessed using `journalctl`. For systems that run all the time or have a high level of debugging turned on, we recommend checking the system's logrotate settings.

```
[Unit]
Description=VOLTRON Platform Service
After=network.target

[Service]
Type=simple

#Change this to the user that VOLTRON will run as.
User=volttron
Group=volttron

#Uncomment and change this to specify a different VOLTRON_HOME
#Environment="VOLTRON_HOME=/home/volttron/.volttron"

#Change these to settings to reflect the install location of VOLTRON
WorkingDirectory=/var/lib/volttron
ExecStart=/var/lib/volttron/env/bin/volttron -vv
ExecStop=/var/lib/volttron/env/bin/volttron-ctl shutdown --platform

[Install]
WantedBy=multi-user.target
```

After the file has been modified to reflect the setup of the platform you can install it with the following commands. These need to be run as root or with `sudo` as appropriate.

```
#Copy the service file into place
cp scripts/admin/volttron.service /etc/systemd/system/

#Set the correct permissions if needed
chmod 644 /etc/systemd/system/volttron.service

#Notify systemd that a new service file exists (this is crucial!)
systemctl daemon-reload

#Start the service
systemctl start volttron.service
```

Init.d

An example init script `scripts/admin/volttron` can be used as a starting point for setting up VOLTRON as a service on `init.d` based systems.

Minor changes may be needed for the file to work on the target system. Specifically the `USER`, `VLHOME`, and `VOLTRON_HOME` variables may need to be changed.

```
...
#Change this to the user VOLTTRON will run as.
USER=volttron
#Change this to the install location of VOLTTRON
VLHOME=/var/lib/volttron

...

#Uncomment and change this to specify a different VOLTTRON_HOME
#export VOLTTRON_HOME=/home/volttron/.volttron
```

The script can be installed with the following commands. These need to be run as root or with sudo as appropriate.

```
#Copy the script into place
cp scripts/admin/volttron /etc/init.d/

#Make the file executable
chmod 755 /etc/init.d/volttron

#Change the owner to root
chown root:root /etc/init.d/volttron

#These will set it to startup automatically at boot
update-rc.d volttron defaults

#Start the service
/etc/init.d/volttron start
```

Platform Hardening for VOLTTRON

Rev. 0 | 1/29/2015 | Initial Document Development

Rev. 1 | 2/5/2015 | Integrate comments from extended VOLTTRON team.

Introduction

VOLTTRON is an agent-based application development platform for distributed control systems. VOLTTRON itself is built with modern security principles in mind [security-wp] and implements many security features for hosted agents. However, VOLTTRON is built on top of Linux and the underlying Linux platform also needs to be secured in order to declare the resulting control system as “secure.” Any system is only as secure as its weakest link. The rest of this note is dedicated to making recommendations for hardening of the underlying Linux platform that VOLTTRON uses. Note that no system can be 100% secure and the cyber security strategy that is recommended in this document is based on risk management.

Linux System Hardening

Here are the non-exhaustive recommendations for Linux hardening from the VOLTTRON team:

- **Physical Security:** Keep the system in locked cabinets or a locked room. Limit physical access to systems and to the networks to which they are attached. The goal should be to avoid physical access by untrusted personnel. This could be extended to blocking or locking USB ports, removable media drives, etc. Drive encryption could be used to avoid access via alternate-media booting (off USB stick or DVD) if physical access can't be guaranteed. Downside of drive encryption would be needing to enter a passphrase to start system. Alternately, the Trusted Platform Module (TPM) may be used, but the drive might still be accessible to those with physical

access. Enable chassis intrusion detection and reporting if supported. If available, use a physical tamper seal along with or in place of an interior switch.

- **Low level device Security:** Keep firmware of all devices (including BIOS) up-to-date. Password-protect the BIOS. Disable unneeded/unnecessary devices including serial, parallel, USB, Firewire, etc. ports; optical drives; wireless devices, such as Wi-Fi and Bluetooth. Leaving a USB port enabled may be helpful if a breach occurs to allow saving forensic data to an external drive.
- **Boot security:** Disable automounting of external devices. Restrict the boot device. Disable PXE and other network boot options (unless that is the primary boot method). Disable booting from USB and other removable drives. Secure the boot loader. Require an administrator password to do anything but start the default kernel. Do not allow editing of kernel parameters. Disable, remove, or password-protect emergency/recovery boot entries.
- **Security Updates:** First and foremost, configure the system to automatically download security updates. Most security updates can be installed without rebooting the system, but some updated (e.g. shared libraries, kernel, etc) require the system to be rebooted. If possible, configure the system to install the security updates automatically and reboot at a particular time. We also recommend reserving the reboot time (e.g. 1:30AM on a Saturday morning) using the Actuator Agent so that no control actions can happen during that time.
- **System Access only via Secured Protocols:** Disallow all clear text access to VOLTRON systems. No telnet, no rsh, no ftp and no exceptions. Use ssh to gain console access, and scp/sftp to get files in and out of the system. Disconnect excessively idle SSH Sessions.
- **Disable remote login for “root” users.** Do not allow a user to directly access the system as the “root” user from a remote network location. Root access to privileged operations can be accomplished using “sudo” This adds an extra level of security by restricting access to privileged operations and tracking those operations through the system log.
- **Manage users and usernames.** Limit the number of user accounts. Use complex usernames rather than first names.
- **Authentication.** If possible, use two factor authentication to allow access to the system. Informally, two factor authentication uses a combination of “something you know” and “something you have” to allow access to the system. RSA SecurID tokens are commonly used for two factor authentication but other tools are available. When not using two-factor authentication, use strong passwords and do not share accounts.
- **Scan for weak passwords.** Use password cracking tools such as John the Ripper (<http://www.openwall.com/john/>) or nmap with password cracking modules (<http://nmap.org>) to look for weak passwords.
- **Utilize Pluggable Authentication Modules (PAM) to strengthen passwords and the login process.** We recommend:
 - pam_abl: Automated blacklisting on repeated failed authentication attempts
 - pam_captcha: A visual text-based CAPTCHA challenge module for PAM
 - pam_passwdqc: A password strength checking module for PAM-aware password changing programs
 - pam_cracklib: PAM module to check the password against dictionary words
 - pam_pwhistory: PAM module to remember last passwords
- **Disable unwanted services.** Most desktop and server Linux distributions come with many unnecessary services enabled. Disable all unnecessary services. Refer to your distribution’s documentation to discover how to check and disable these services.
- **Just as scanning for weak passwords is a step to more secure systems, regular network scans using Nmap (www.nmap.org) to find what network services are being offered is another step towards a more secure system.** Note, use nmap or similar tools very carefully on BACnet and modbus environments. These scanning tools are known to crash/reset BACnet and modbus devices.

- Control incoming and outgoing network traffic. Use the built-in host-based firewall to control who/what can connect to this system. Many iptables frontends offer a set of predefined rules that provide a default deny policy for incoming connections and provide rules to prevent or limit other well known attacks (i.e. rules that limit certain responses that might amplify a DDoS attack). ufw (uncomplicated firewall) is a good example. For example, if the system administrators for the VOLTRON device are all located in 10.10.10.0/24 subnetwork, then allow SSH and SCP logins from only that IP address range. If VOLTRON system exports data to a historian at 10.20.20.1 using TCP port 443, allow outgoing traffic to that port on that server. The idea here is to limit the attack surface of the system. The smaller the surface, the better we can analyze the communication patterns of the system and detect anomalies. One word of caution. While some system administrators disable network-based diagnostic tools such as ICMP ECHO responses, VOLTRON team believes that this hampers usability. As an example, monitoring which incoming and outgoing firewall rules are triggering can be accomplished with this command: `watch --interval=5 'iptables -nvL | grep -v "0 0"'`.
- Rate limit incoming connections to discourage brute force hacking attempts. Use a tool such as fail2ban (http://www.fail2ban.org/wiki/index.php/Main_Page) to dynamically manage firewall rules to rate limit incoming connections and discourage brute force hacking attempts. sshguard (<http://www.sshguard.net/>) is similar to fail2ban but only used for ssh connections. Further rate limiting can be accomplished at the firewall level. As an example, you can restrict the number of connections used by a single IP address to your server using iptables. Only allow 4 ssh connections per client system: `iptables -A INPUT -p tcp --syn --dport 22 -m connlimit --connlimit-above 4 -j DROP` You can limit the number of connections per minute. The following example will drop incoming connections if an IP address makes more than 10 connection attempts to port 22 within 60 seconds: `iptables -A INPUT -p tcp -dport 22 -i eth0 -m state --state NEW -m recent --set iptables -A INPUT -p tcp -dport 22 -i eth0 -m state --state NEW -m recent --update --seconds 60 --hitcount 10 -j DROP`
- Use a file system integrity tool to monitor for unexpected file changes. Tools such as tripwire (<http://sourceforge.net/projects/tripwire/>) to monitor filesystem for changed files. Another file integrity checking tool to consider is AIDE (Advanced Intrusion Detect Environment) (<http://aide.sourceforge.net/>).
- Use filesystem scanning tools periodically to check for exploits. Available tools such as checkrootkit (<http://www.chkrootkit.org>), rkhunter (<http://rkhunter.sourceforge.net>) and others should be used to check for known exploits on a periodic basis and report their results.
- VOLTRON does not use apache or require it. If Apache is being used, we recommend using mod_security and mod_evasive modules.

System Monitoring

- Monitor system state and resources. Use a monitoring tool such as Xymon (<http://xymon.sourceforge.net>) or big brother (<http://www.bb4.org/features.html>) to remotely monitor the system resources and state. Set the monitoring tools to alert the system administrators if anomalous use of resources (e.g. connections, memory, etc) are detected. An administrator can also use unix commands such as netstat to look for open connections periodically.
- Watch system logs and get logs off the system. Use a utility such as logwatch (<http://sourceforge.net/projects/logwatch/files/>) or logcheck (<http://logcheck.org>) to get daily summary of system activity via email. For Linux distributions that use systemd, use journalwatch (<http://git.the-compiler.org/journalwatch/>) to accomplish the same task. Additionally, use a remote syslog server to collect logs from all VOLTRON systems in the field at a centralized location for analysis. A tool such as splunk is ideal for this task and comes with many built-in analysis applications. Another benefit of sending logs remotely off the platform is the ability to inspect the logs even when the platform may be compromised.
- An active intrusion sensor such as PSAD (<http://cipherdyne.org/psad/>) can be used to look for intrusions as well.

Security Testing

Every security control discussed in the previous sections must be tested to determine correct operation and impact. For example, if we inserted a firewall rule to ban connections from an IP address such as 10.10.10.2, then we need to test that the connections actually fail.

In addition to functional correctness testing, common security testing tools such as Nessus (<http://www.tenable.com/products/nessus>) and nmap (<http://nmap.org>) should be used to perform cyber security testing.

Conclusion

No system is 100% secure unless it is disconnected from the network and is in a physically secure location. VOLTRON team recommends a risk-based cyber security approach that considers each risk, and the impact of an exploit. Mitigating technologies can then be used to mitigate the most impactful risks first. VOLTRON is built with security in mind from the ground up. But it is only as secure as the operating system that it runs on top of. This document is intended to help VOLTRON users to secure the underlying Linux operating system to further improve the robustness of the VOLTRON platform. Any security questions should be directed to voltrtron@pnnl.gov.

Platform External Address Configuration

In the configuration file located in `$VOLTRON_HOME/config` add `vip-address=tcp://ip:port` for each address you want to listen on

```
Example
vip-address=tcp://127.0.0.102:8182
vip-address=tcp://127.0.0.103:8083
vip-address=tcp://127.0.0.103:8183
```

Note: The config file is generated after running the `vcfg` command. The `vip-address` is for the local platform, NOT the remote platform.

3.4.3 Walkthroughs

How to authenticate an agent to communicate with VOLTRON platform:

An administrator can allow an agent to communicate with VOLTRON platform by creating an authentication record for that agent. An authentication record is created by using `voltrtron-ctl auth add` command and entering values to asked arguments.

```
voltrtron-ctl auth add

domain []:
address []:
user_id []:
capabilities (delimit multiple entries with comma) []:
roles (delimit multiple entries with comma) []:
groups (delimit multiple entries with comma) []:
mechanism [CURVE]:
credentials []:
comments []:
enabled [True]:
```

The listed fields can also be specified on the command line:

```
voltttron-ctl auth add --user_id bob --credentials ABCD...
```

If any field is specified on the command line, then the interactive menu will not be used.

The simplest way of creating an authentication record is by entering the `user_id` and credential values. `User_id` is an arbitrary string for VOLTTRON to identify the agent. `Credential` is the encoded public key string for the agent. Create a public/private key pair for the agent and enter encoded public key for credential parameter.

```
voltttron-ctl auth add

domain []:
address []:
user_id []: my-test-agent
capabilities (delimit multiple entries with comma) []:
roles (delimit multiple entries with comma) []:
groups (delimit multiple entries with comma) []:
mechanism [CURVE]:
credentials []: encoded-public-key-for-my-test-agent
comments []:
enabled [True]:
```

In next sections, we will discuss each parameter, its purpose and what all values it can take.

Domain:

Domain is the name assigned to locally bound address. Domain parameter is currently not being used in VOLTTRON and is placeholder for future implementation.

Address:

By specifying address, administrator can allow an agent to connect with VOLTTRON only if that agent is running on that address. Address parameter can take a string representing an IP addresses. It can also take a regular expression representing a range of IP addresses.

```
address []: 192.168.111.1
address []: /192.168.*/
```

User_id:

`User_id` can be any arbitrary string that is used to identify the agent by the platform. If a regular expression is used for address or credential to combine agents in an authentication record then all those agents will be identified by this `user_id`. It is primarily used for identifying agents during logging.

Capabilities:

Capability is an arbitrary string used by an agent to describe its exported RPC method. It is used to limit the access to that RPC method to only those agents who have that capability listed in their authentication record.

If administrator wants to authorize an agent to access an exported RPC method with capability of another agent, he/she can list that capability string in this parameter. Capability parameter takes a string or an array of strings listing all the

capabilities this agent is authorized to access. Listing capabilities here will allow this agent to access corresponding exported RPC methods of other agents.

For example, if there is an AgentA with capability enables exported RPC method and AgentB needs to access that method then AgentA's code and AgentB's authentication record would be as follow:

AgentA's capability enabled exported RPC method:

```
@RPC.export
@RPC.allow('can_call_bar')
def bar(self):
    return 'If you can see this, then you have the required capabilities'
```

AgentB's authentication record to access bar method:

```
volttron-ctl auth add

domain []:
address []:
user_id []: agent-b
capabilities (delimit multiple entries with comma) []: can_call_bar
roles (delimit multiple entries with comma) []:
groups (delimit multiple entries with comma) []:
mechanism [NULL]: CURVE
credentials []: encoded-public-key-for-agent-b
comments []:
enabled [True]:
```

Similarly, capability parameter can take an array of string:

```
capabilities (delimit multiple entries with comma) []: can_call_bar
capabilities (delimit multiple entries with comma) []: can_call_method1, can_call_
↪method2
```

Roles:

A role is a name for a set of capabilities. Roles can be used to grant an agent multiple capabilities without listing each capability in the in the agent's authorization entry. Capabilities can be fully utilized without roles. Roles are purely for organizing sets of capabilities.

Roles can be viewed and edited with the following commands:

- `volttron-ctl auth add-role`
- `volttron-ctl auth list-roles`
- `volttron-ctl auth remove-role`
- `volttron-ctl auth updated-role`

For example, suppose agents protect certain methods with the following capabilities: `READ_BUILDING_A_TEMP`, `SET_BUILDING_A_TEMP`, `READ_BUILDING_B_TEMP`, and `SET_BUILDING_B_TEMP`.

These capabilities can be organized into various roles:

```
volttron-ctl auth add-role TEMP_READER READ_BUILDING_A_TEMP READ_BUILDING_B_TEMP
volttron-ctl auth add-role BUILDING_A_ADMIN READ_BUILDING_A_TEMP SET_BUILDING_A_TEMP
volttron-ctl auth add-role BUILDING_B_ADMIN READ_BUILDING_B_TEMP SET_BUILDING_B_TEMP
```

To view these roles run `volttron-ctl auth list-roles`:

ROLE	CAPABILITIES
BUILDING_A_ADMIN	['READ_BUILDING_A_TEMP', 'SET_BUILDING_A_TEMP']
BUILDING_B_ADMIN	['READ_BUILDING_B_TEMP', 'SET_BUILDING_B_TEMP']
TEMP_READER	['READ_BUILDING_A_TEMP', 'READ_BUILDING_B_TEMP']

With this configuration, adding the BUILDING_A_ADMIN role to an agent's authorization entry implicitly grants that agent the READ_BUILDING_A_TEMP and SET_BUILDING_A_TEMP capabilities.

To add a new capabilities to an existing role:

```
voltttron-ctl auth update-role BUILDING_A_ADMIN CLEAR_ALARM TRIGGER_ALARM
```

To remove a capability from a role:

```
voltttron-ctl auth update-role BUILDING_A_ADMIN TRIGGER_ALARM --remove
```

Groups:

Groups provide one more layer of *grouping*. A group is a named set of roles. Like roles, groups are optional and are meant to help with organization.

Groups can be viewed and edited with the following commands:

- voltttron-ctl auth add-group
- voltttron-ctl auth list-groups
- voltttron-ctl auth remove-group
- voltttron-ctl auth updated-group

These commands behave the same as the *role* commands. For example, to further organize the capabilities in the previous section, one could create an ALL_BUILDING_ADMIN group:

```
voltttron-ctl auth add-group ALL_BUILDING_ADMIN BUILDING_A_ADMIN BUILDING_B_ADMIN
```

With this configuration, agents in the ALL_BUILDING_ADMIN group would implicitly have the BUILDING_A_ADMIN and BUILDING_B_ADMIN roles. This means such agents would implicitly be granted the following capabilities: READ_BUILDING_A_TEMP, SET_BUILDING_A_TEMP, READ_BUILDING_B_TEMP, and SET_BUILDING_B_TEMP.

Mechanism:

Mechanism is the authentication method by which the agent will communicate with VOLTTRON platform. Currently VOLTTRON uses only CURVE mechanism to authenticate agents.

Credentials:

The credentials field must be an CURVE encoded public key (see *voltttron.platform.vip.socket.encode_key* for method to encode public key).

```
credentials []: encoded-public-key-for-agent
```

Comments:

Comments is arbitrary string to associate with authentication record

Enabled:

TRUE or FALSE value to enable or disable the authentication record. Record will only be used if this value is True

Deployment Walkthrough

This page is meant as an overview of setting up a VOLTRON deployment which consists of one or more platforms collecting data and being managed by another platform running the VOLTRON Central agent. High level instructions are included but for more details on each step, please follow links to that section of the wiki.

Assumptions:

- “Data Collector” is the box that has the drivers and is collecting data it needs to forward.
- “Voltrtron Central/VC” is the box that has the historian which will save data to the database.
- VOLTRON_HOME is assumed to the default on both boxes which is: /home/<user>/voltrtron

Notes/Tips:

- Aside from installing the required packages with apt-get, sudo is not required and should not be used. VOLTRON is designed to be run as a non-root user and running with sudo is not supported.
- The convenience *scripts* have been developed to simplify many of the repetitive multi-step processes. For instance, `scripts/core/make-listener` can be modified for any agent and make it one command to stop, remove, build, install, configure, tag, start, and (optionally) enable an agent for autostart.
- These instructions assume default directories are used (for instance, /home/<user>/voltrtron for the project directory and /home/<user>/voltrtron for the VOLTRON Home directory).
- Creating a separate `config` directory for agent configuration files used in the deployment can prevent them from being committed back to the repository.
- Double check firewall rules/policies when setting up a multi-node deployment to ensure that platforms can communicate

On all machines:

On all machines in the deployment, setup the platform, setup encryption, authentication, and authorization. Also, build the basic agents for the deployment. All platforms will need a PlatformAgent and a Historian. Using *scripts* will help simplify this project.

Install required packages

- `sudo apt-get install build-essential python-dev openssl libssl-dev libevent-dev git`

Build the project

- Clone the repository and build using `python bootstrap.py`

Configuring Platform

On VC:

- Run *voltron-cfg*
- Setup as VOLTRON Central.
- Set appropriate ip, port, etc for this machine
- Pick to install a platform historian (defaults to sqlite)
- Start up the platform and find the line with the server public key “cat voltron.log|grep “public key””:

```
2016-05-19 08:42:58,062 () voltron.platform.main INFO: public key: <KEY>
```

On the data collector:

Setup drivers

For a simple case, follow instructions to install a *Fake Driver* for testing purposes. For an actual deployment against real devices see the following:

- Create a Master Driver Agent to coordinate drivers for the devices controlled by this platform.
- For *MODBUS* devices, create config files and point configuration files.
- For BACnet devices, create a *Proxy Agent* for *BACnet drivers* to communicate through

Now that data is being published to the bus, a *Forward Historian* can be configured to send this data to the VC instance for storage.

- Use: `voltron-ctl keypair` to generate a keypair
- `cat VOLTTRON_HOME/keypair` to get the public and secret keys
- Create a config directory in the main project directory
- Setup a Forward Historian
 - `cp services/core/ForwardHistorian/config config/forwarder.config`
 - Edit `forwarder.config` using the VC’s VIP address, the public server key, and the keypair
 - “destination-vip”: “`tcp://<VC_IP>:<VC_PORT>?serverkey=<server_key>&secretkey=<secret_key>&publickey=<public_`
 - For ease of use, you can create a script to install the historian:
 - `cp scripts/core/make-listener ./make-forwarder`, then edit to look like:

```
make-forwarder:: export SOURCE=services/core/ForwardHistorian export CONFIG=config/forwarder.config export TAG=forward
```

```
./scripts/core/make-agent.sh enable
```

- Execute that script and the forward historian should be installed

To check that things are working: Start a listener agent on VC, you should see data from the data collector appear

In the log for VC, check for credentials success for the ip of data collector.

Registering the collection platform

- In a browser, go to the url for your VC instance.
- Click on Register Platforms
- Enter a name for the collection platform and the ip configured `http://<ip>:<discovery port>`
- Open the tree upper left of the UI and find your platform.

Troubleshooting:

- Check firewall rules registering VC on VC: `ipc://@/home/voltron/.voltron/run/vip.socket` Change password by putting pw hash in config file Add remote ip address to config file

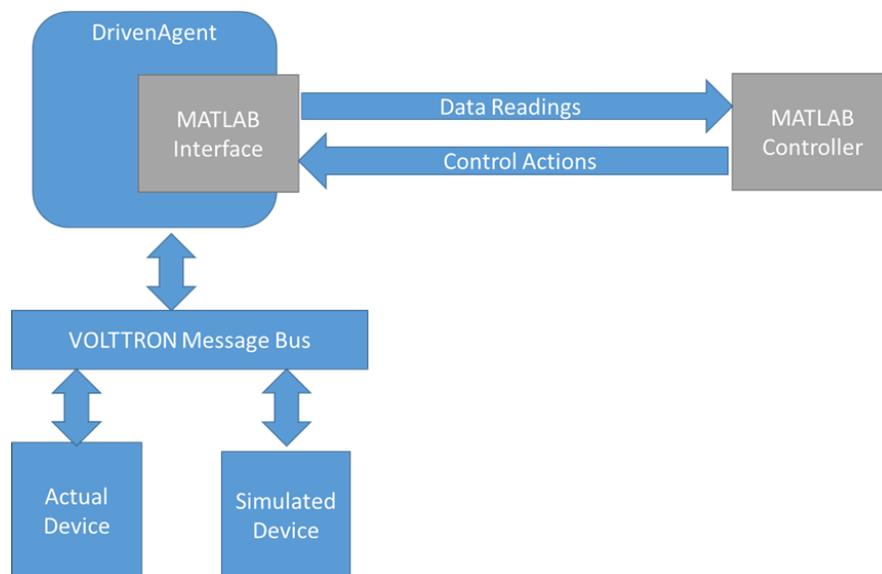
MatLab Integration

Overview:

Matlab-VOLTRON integration allows Matlab applications to receive data from devices and send control commands to change points on those devices.

DrivenMatlabAgent in VOLTRON allows this interaction by using ZeroMQ sockets to communicate with the Matlab application.

Data Flow Architecture:



Installation steps for system running Matlab:

1. Install python. Suggested 3.4. Other supported versions are 2.7, 3.3.
2. Install pyzmq (tested with version 15.2.0) Follow steps at: <https://github.com/zeromq/pyzmq>

3. Install Matlab (tested with R2015b)
4. Start Matlab and set the python path. In the Matlab command window set the python path with *pyversion*:

```
>> pyversion python.exe
```

5. To test that the python path has been set correctly type following in the Matlab command window. Matlab should print the python path with version information.

```
>> pyversion
```

6. To test that the pyzmq library is installed correctly and is accessible from python inside Matlab, type the following in Matlab command window and it should show pyzmq version installed.

```
>> py.zmq.pyzmq_version()
```

7. Copy *example.m* from *voltrtron/examples/ExampleMatlabApplication/matlab* to your desired folder.

Run and test Matlab VOLTTRON Integration:

Assumptions

- Device driver agent is already developed

Installation:

1. Install VOLTTRON on a VM or different system than the one running Matlab.

Follow link: <http://voltrtron.readthedocs.io/en/develop/install.html>

2. Add subtree *voltrtron-applications* under *voltrtron/applications* by using the following command:

```
git subtree add --prefix applications https://github.com/VOLTTRON/voltrtron-
↳applications.git develop --squash
```

Configuration

1. Copy example configuration file *applications/pnnl/DrivenMatlabAgent/config_waterheater* to *voltrtron/config*.
2. Change *config_url* and *data_url* in the new config file to the ipaddress of machine running Matlab. Keep the same port numbers.
3. Change *campus*, *building* and *unit (device) name* in the config file.
4. Open *example.m* and change following line:

```
matlab_result = '{"commands":{"Zone1":["temperature",27],"Zone2":["temperature",
↳28]}}';
```

Change it to include correct device name and point names in the format:

```
'{"commands":{"device1":["point1",value1],"device2":["point2",value2]}}';
```

Steps to test integration:

1. Start VOLTRON
2. Run Actuator
3. Run device driver agent
4. Run DrivenMatlabAgent with the new config file
5. Run example.m in Matlab

Now whenever the device driver publishes the state of devices listed in the config file of DrivenMatlabAgent, DrivenMatlabAgent will send it to Matlab application and receive commands to send to devices.

Resources

http://www.mathworks.com/help/matlab/getting-started_buik_wp-3.html

Forward Historian Deployment

This guide describes a simple setup where one Volttron instance collects data from a fake devices and sends to another instance . Lets consider the following example.

We are going to create two VOLTRON instances and send data from one VOLTRON instance running a fake driver(subscribing values from a fake device)and sending the values to the second VOLTRON instance.

VOLTRON instance 1 forwards data to VOLTRON instance 2**VOLTRON instance 1**

- Vctl shutdown –platform (if the platform is already working)
- Volttron-cfg (this helps in configuring the volttron instance http://volttron.readthedocs.io/en/releases-4.1/core_services/control/VOLTRON-Config.html)
 - Specify the IP of the machine : `tcp://130.20.*.*:22916`
 - Specify the port you want to use
 - Specify if you want to run VC(Volttron Central) here or this instance would be controlled by a VC and the IP and port of the VC
 - * Then install agents like Master driver Agent with fake driver agent for the instance.
 - * Install a listener agent so see the topics that are coming from the diver agent
 - * Then run the volttron instance by : `volttron -vv -l volttron.log&`
- Volttron authentication : We need to add the IP of the instance 2 in the auth.config file of the VOLTRON agent. This is done as follow :
 - Vctl auth-add
 - We specify the IP of the instance 2 and the credentials of the agent (<http://volttron.readthedocs.io/en/releases-4.1/devguides/walkthroughs/Agent-Authentication-Walkthrough.html?highlight=auth-add>)
 - For specifying authentication for all the agents , we specify `./*` for credentials as shown in http://volttron.readthedocs.io/en/releases-4.1/devguides/agent_development/index.html

- This should enable authentication for all the volttron-instance based on the IP you specify here

For this documentation, the topics from the driver agent will be send to the instance 2

- We use the existing agent called the Forward Historian for this purpose which is available in service/core in the VOLTTRON directory.
- In the config file under the ForwardHistorian directory , we modify the following field : - Destination-vip : the IP of the volttron instance to which we have to forward the data to along with the port number .

Example : “tcp://130.20.*.*:22916”

- Destination-serverkey: The server key of the VOLTTRON instance to which we need to forward the data to.

This can be obtained at the VOLTTRON instance by typing vctl auth serverkey

- Service_topic_list: specify the topics you want to forward specifically instead of all the values.
- Once the above values are set, your forwarder is all set .
- You can create a script file for the same and execute the agent.

VOLTTRON instance 2

- Vctl shutdown –platform (if the platform is already working)
- Volttron-cfg (this helps in configuring the volttron instance) http://volttron.readthedocs.io/en/releases-4.1/core_services/control/VOLTTRON-Config.html - Specify the IP of the machine : tcp://130.20.*.*:22916 - Specify the port you want to use. - Install the listener agent (this will show the connection from instance 1 if its successful and then show all the topics from instance 1.
- Volttron authentication : We need to add the IP of the instance 1 in the auth.config file of the VOLTTRON agent .This is done as follow : - Vctl auth-add - We specify the IP of the instance 1 and the credentials of the agent
<http://volttron.readthedocs.io/en/releases-4.1/devguides/walkthroughs/Agent-Authentication-Walkthrough.html?highlight=auth-add>
 - For specifying authentication for all the agents , we specify ./ */ for credentials as shown in http://volttron.readthedocs.io/en/releases-4.1/devguides/agent_development/index.html
 - This should enable authentication for all the volttron-instance based on the IP you specify here

Listener Agent

- Run the listener agent on this instance to see the values being forwarded from instance 1.

Once the above setup is done, you should be able to see the values from instance 1 on the listener agent of instance 2.

Forward Historian Walkthrough

This guide describes a simple setup where one VOLTTRON instance collects data from a fake devices and sends to another instance . Lets consider the following example.

We are going to create two VOLTTRON instances and send data from one VOLTTRON instance running a fake driver(subscribing values from a fake device)and sending the values to the second VOLTTRON instance.

VOLTRON instance 1 forwards data to VOLTRON instance 2

VOLTRON INSTANCE 1

- `voltttron-ctl shutdown --platform` (If VOLTRON is already running it must be shut down before running `voltttron-cfg`).
- `voltttron-cfg` - this helps in configuring the VOLTRON instance (*VOLTRON Config*).
 - Specify the IP of the machine : `tcp://127.0.0.1:22916`.
 - Specify the port you want to use.
 - Specify if you want to run VC (VOLTRON Central) here or this this instance would be controlled by a VC and the IP and port of the VC.
- Then start the VOLTRON instance by : `voltttron -vv & > voltttron.log&`.
- Then install agents like Master driver Agent with fake driver agent for the instance.
- Install a listener agent so see the topics that are coming from the diver agent.
- VOLTRON authentication : We need to add the IP of the instance 1 in the `auth.config` file of the VOLTRON agent .This is done as follow :
 - `voltttron-ctl auth-add`
 - We specify the IP of the instance 1 and the credentials of the agent. (*Agent authentication walkthrough*)
 - For specifying authentication for all the agents , we specify `./ */` for credentials as shown in *Agent Development*.
 - This should enable authentication for all the VOLTRON instances based on the IP you specify here .

For this documentation, the topics from the driver agent will be sent to the instance 2

- We use the existing agent called the Forward Historian for this purpose which is available in `service/core` in the VOLTRON directory.
- In the config file under the ForwardHistorian directory , we modify the following field:
 - `destination-vip` : the IP of the VOLTRON instance to which we have to forward the data to along with the port number . Example : `tcp://130.20.*.*:22916`.
 - `destination-serverkey`: The server key of the VOLTRON instance to which we need to forward the data to. This can be obtained at the VOLTRON instance by typing `voltttron-ctl auth serverkey`.
 - `service_topic_list`: specify the topics you want to forward specifically instead of all the values.
- Once the above values are set, your forwarder is all set .
- You can create a script file for the same and execute the agent.

VOLTRON INSTANCE 2

- `voltttron-ctl shutdown --platform` (If VOLTRON is already running it must be shut down before running `voltttron-cfg`).
- `voltttron-cfg` - this helps in configuring the VOLTRON instance. (*VOLTRON Config*) - Specify the IP of the machine : `tcp://127.0.0.1:22916`. - Specify the port you want to use. - Install the listener agent (this will show the connection from instance 1 if its successful and then show all the topics from instance 1.

- Then start the VOLTTRON instance by : `voltttron -vv & > voltttron.log&`.
- VOLTTRON authentication : We need to add the IP of the instance 1 in the `auth.config` file of the VOLTTRON agent .This is done as follow :
 - `voltttron-ctl auth-add`
 - We specify the IP of the instance 1 and the credentials of the agent.(*Agent authentication walkthrough*)
 - For specifying authentication for all the agents , we specify `./.*` for credentials as shown in *Agent Development*.
 - This should enable authentication for all the VOLTTRON instances based on the IP you specify here .

LISTENER AGENT

- Run the listener agent on this instance to see the values being forwarded from instance 1.

Once the above setup is done, you should be able to see the values from instance 1 on the listener agent of instance 2.

Multi-Platform Connection Walkthrough

Multi-Platform message bus communication alleviates the need for an agent in one platform to connect to another platform directly in order for it to send/receive messages from the other platform. With multi-platform communication, connections to external platforms will be maintained by the platforms itself and agents do not have the burden to manage the connections directly. This guide will show how to connect three VOLTTRON instances with a fake driver running on VOLTTRON instance 1 publishing to topic with prefix="devices" and listener agents running on other 2 VOLTTRON instances subscribed to topic "devices".

- *Getting Started*
- *Multi-Platform Configuration*
- *Configuration and Authentication in Setup Mode*
- *Setup Configuration and Authentication Manually*
- *Start Master driver on VOLTTRON instance 1*
- *Start Listener agents on VOLTTRON instance 2 and 3*
- *Stopping All the Platforms*

Getting Started

Modify the `subscribe` annotate method parameters in the listener agent (`examples/ListenerAgent/listener/agent.py` in the VOLTTRON root directory) to include `all_platforms=True` parameter to receive messages from external platforms.

```
@PubSub.subscribe('pubsub', '')
```

to

```
@PubSub.subscribe('pubsub', 'devices', all_platforms=True)
```

or add below line in the `onstart` method

```
self.vip.pubsub.subscribe('pubsub', 'devices', self.on_match, all_platforms=True)
```

Note: If using the onstart method remove the `@PubSub.subscribe('pubsub', '')` from the top of the method.

After *building VOLTRON*, open three shells with the current directory the root of the VOLTRON repository. Then activate the VOLTRON environment and export the `VOLTRON_HOME` variable. The home variable needs to be different for each instance.

```
$ source env/bin/activate
$ export VOLTRON_HOME=~/.volttron1
```

Run `volttron-cfg` in all the three shells. This command will ask how the instance should be set up. Many of the options have defaults and that will be sufficient. Enter a different VIP address for each platform. Configure fake master driver in the first shell and listener agent in second and third shell.

```

/bin/bash 92x52
(volttron)volttron@volttron-VirtualBox ~/git/volttron-2 $ source env/bin/activate
(volttron)volttron@volttron-VirtualBox ~/git/volttron-2 $ export VOLTRON_HOME=~/.volttron1
(volttron)volttron@volttron-VirtualBox ~/git/volttron-2 $ volttron-cfg

Your VOLTRON_HOME currently set to: /home/volttron/.volttron1
Is this the volttron you are attempting to setup? [Y]:
What is the external instance ipv4 address? [tcp://127.0.0.1]:
What is the instance port for the vip address? [22916]:
Is this instance a volttron central? [N]:
Will this instance be controlled by volttron central? [Y]: n
Would you like to install a platform historian? [N]:
Would you like to install a master driver? [N]: y
Configuring /home/volttron/git/volttron-2/services/core/MasterDriverAgent
Install a fake device on the master driver? [N]: y
Should agent autostart? [N]: n
Would you like to install a listener agent? [N]: n
Finished configuration

You can now start the volttron instance.

If you need to change the instance configuration you can edit
the config file at /home/volttron/.volttron1/config

(volttron)volttron@volttron-VirtualBox ~/git/volttron-2 $

/bin/bash 55x52
(volttron)volttron@volttron-VirtualBox ~/git/volttron-2
~/bin/activate
(volttron)volttron@volttron-VirtualBox ~/git/volttron-2
~/bin/activate
(volttron)volttron@volttron-VirtualBox ~/git/volttron-2
$ export VOLTRON_HOME=~/.volttron2
(volttron)volttron@volttron-VirtualBox ~/git/volttron-2
$ volttron-cfg

Your VOLTRON_HOME currently set to: /home/volttron/.volttron2
Is this the volttron you are attempting to setup? [Y]:
What is the external instance ipv4 address? [tcp://127.0.0.1]: tcp://127.0.0.3
What is the instance port for the vip address? [22916]:
Is this instance a volttron central? [N]:
Will this instance be controlled by volttron central? [Y]: n
Would you like to install a platform historian? [N]:
Would you like to install a master driver? [N]:
Would you like to install a listener agent? [N]: y
Configuring examples/ListenerAgent
Should agent autostart? [N]:
Finished configuration

You can now start the volttron instance.

If you need to change the instance configuration you can edit
the config file at /home/volttron/.volttron2/config

(volttron)volttron@volttron-VirtualBox ~/git/volttron-2
$

/bin/bash 84x52
(volttron)volttron@volttron-VirtualBox ~/git/volttron-2 $ source env/bin/activate
(volttron)volttron@volttron-VirtualBox ~/git/volttron-2 $ source env/bin/activate
(volttron)volttron@volttron-VirtualBox ~/git/volttron-2 $ export VOLTRON_HOME=~/.volttron3
(volttron)volttron@volttron-VirtualBox ~/git/volttron-2 $ volttron-cfg

Your VOLTRON_HOME currently set to: /home/volttron/.volttron3
Is this the volttron you are attempting to setup? [Y]:
What is the external instance ipv4 address? [tcp://127.0.0.1]: tcp://127.0.0.3
What is the instance port for the vip address? [22916]:
Is this instance a volttron central? [N]:
Will this instance be controlled by volttron central? [Y]: n
Would you like to install a platform historian? [N]:
Would you like to install a master driver? [N]:
Would you like to install a listener agent? [N]: y
Configuring examples/ListenerAgent
Should agent autostart? [N]:
Finished configuration

You can now start the volttron instance.

If you need to change the instance configuration you can edit
the config file at /home/volttron/.volttron3/config

(volttron)volttron@volttron-VirtualBox ~/git/volttron-2 $

```

Multi-Platform Configuration

For each instance, specify the instance name in platform config file under it's `VOLTRON_HOME` directory. If the platform supports web server, add the `bind-web-address` as well.

Here is an example,

Path of the config: `$VOLTRON_HOME/config`

```
[volttron]
vip-address = tcp://127.0.0.1:22916
instance-name = "platform1"
bind-web-address = http://127.0.0.1:8080
```

Instance name and bind web address entries added into each VOLTRON platform's config file is shown below.

```

/bin/bash 83x52
(volttron)volttron@volttron-VirtualBox ~/git/volttron-2 $ cat ~/.volttron1/config
[volttron]
vip-address = tcp://127.0.0.1:22916
instance-name = "platform1"
bind-web-address = http://127.0.0.1:8080
(volttron)volttron@volttron-VirtualBox ~/git/volttron-2 $

/bin/bash 64x52
(volttron)volttron@volttron-VirtualBox ~/git/volttron-2
~/bin/activate
(volttron)volttron@volttron-VirtualBox ~/git/volttron-2
~/bin/activate
(volttron)volttron@volttron-VirtualBox ~/git/volttron-2
$ cat ~/.volttron2/config
[volttron]
vip-address = tcp://127.0.0.2:22916
instance-name = "platform2"
bind-web-address = http://127.0.0.2:8080
(volttron)volttron@volttron-VirtualBox ~/git/volttron-2
$

/bin/bash 84x52
(volttron)volttron@volttron-VirtualBox ~/git/volttron-2 $ cat ~/.volttron3/config
[volttron]
vip-address = tcp://127.0.0.3:22916
instance-name = "platform3"
bind-web-address = http://127.0.0.3:8080
(volttron)volttron@volttron-VirtualBox ~/git/volttron-2 $

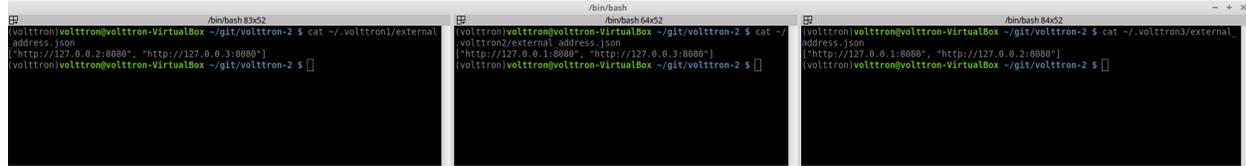
```

Next, each instance needs to know the VIP address, platform name and server keys of the remote platforms that it is connecting to. In addition, each platform has to authenticate or accept the connecting instances' public keys. We can do this step either by running VOLTRON in setup mode or configure the information manually.

Configuration and Authentication in Setup Mode

Note: It is necessary for **each** platform to have a web server if running in setup mode

Add list of web addresses of remote platforms in \$VOLTRON_HOME/external_address.json



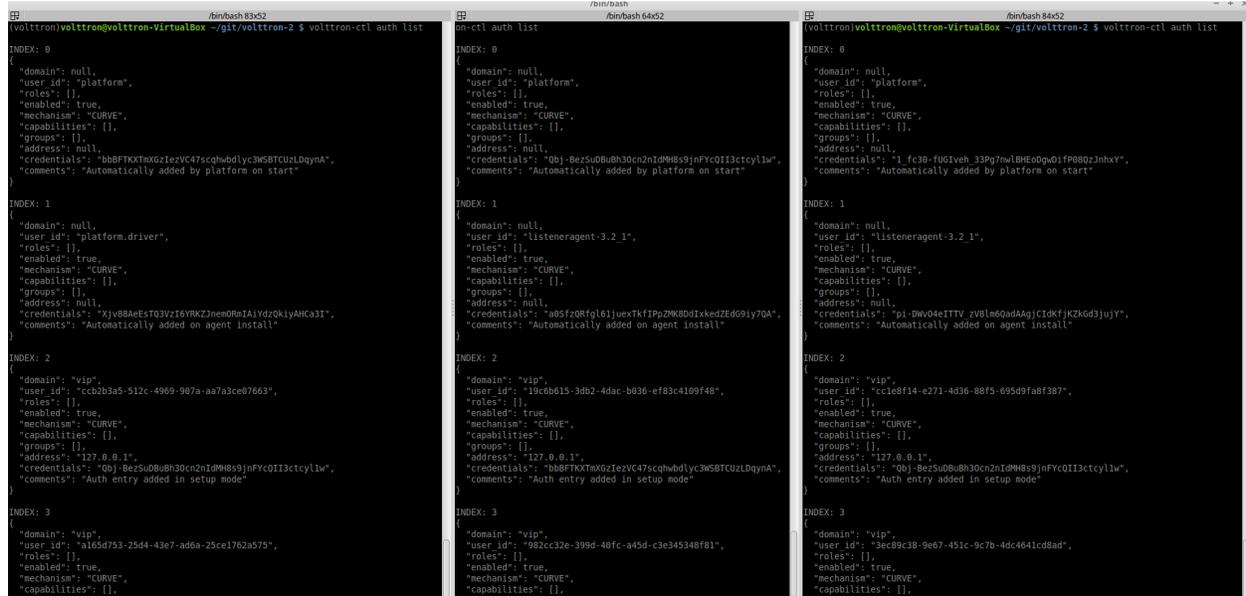
Start VOLTRON instances in setup mode in the three terminal windows. The “-l” option in the following command tells VOLTRON to log to a file. The file name should be different for each instance.

```
$ volttron -v -l 11.log --setup-mode&
```

Note: Don't forget the '&' on the end to put the process in the background.

A new auth entry is added for each new platform connection. This can be checked with below command in each terminal window.

```
$ volttron-ctl auth list
```



After all the connections are authenticated, we can start the instances in normal mode.

```
$ volttron-ctl shutdown --platform
$ volttron -v -l 11.log&
```

Setup Configuration and Authentication Manually

If you do not need web servers in your setup, then you will need to build the platform discovery config file manually. The config file should contain an entry containing VIP address, instance name and serverkey of each remote platform connection.

Name of the file: `external_platform_discovery.json`

Directory path: Each platform's `VOLTRON_HOME` directory.

For example, since VOLTRON instance 1 is connecting to VOLTRON instance 2 and 3, contents of `external_platform_discovery.json` will be

```
{
  "platform2": { "vip-address": "tcp://127.0.0.2:22916",
                "instance-name": "platform2",
                "serverkey": "YFyIqXy2H7gIKC1x6uPMdDOB_i9lzfAPB1IgbxfXLGc"},
  "platform3": { "vip-address": "tcp://127.0.0.3:22916",
                "instance-name": "platform3",
                "serverkey": "hzU2bnlacAhZSaI0rI8a6XK_bqLSpA0JRK4jq8ttZxw"}
}
```

We can obtain the serverkey of each platform using below command in each terminal window:

```
$ voltttron-ctl auth serverkey
```

Contents of `external_platform_discovery.json` of VOLTRON instance 1, 2, 3 is shown below.

```
(voltttron)voltttron@voltttron-VirtualBox ~/git/voltttron-2 $ cat ~/.voltttron1/external_platform_discovery.json
{"platform2":{"vip-address":"tcp://127.0.0.2:22916","instance-name":"platform2","serverkey":"Obj-BezSuDBuBh30cn2nIdMH8s9jnFycQII3ctcyllw"},
"platform3":{"vip-address":"tcp://127.0.0.3:22916","instance-name":"platform3","serverkey":"1_fc30-fUGIveh_33Pg7nwlBHEoDgwDifP08QzJnhxY"}}
(voltttron)voltttron@voltttron-VirtualBox ~/git/voltttron-2 $
(voltttron)voltttron@voltttron-VirtualBox ~/git/voltttron-2 $
(voltttron)voltttron@voltttron-VirtualBox ~/git/voltttron-2 $
(voltttron)voltttron@voltttron-VirtualBox ~/git/voltttron-2 $
(voltttron)voltttron@voltttron-VirtualBox ~/git/voltttron-2 $ cat ~/.voltttron2/external_platform_discovery.json
{"platform3":{"vip-address":"tcp://127.0.0.3:22916","instance-name":"platform3","serverkey":"1_fc30-fUGIveh_33Pg7nwlBHEoDgwDifP08QzJnhxY"},
"platform1":{"vip-address":"tcp://127.0.0.1:22916","instance-name":"platform1","serverkey":"bbBFTKXTmXGzIezVC47scqhwbdlyc3WSBTCuzLDqynA"}}
(voltttron)voltttron@voltttron-VirtualBox ~/git/voltttron-2 $
(voltttron)voltttron@voltttron-VirtualBox ~/git/voltttron-2 $
(voltttron)voltttron@voltttron-VirtualBox ~/git/voltttron-2 $
(voltttron)voltttron@voltttron-VirtualBox ~/git/voltttron-2 $
(voltttron)voltttron@voltttron-VirtualBox ~/git/voltttron-2 $ cat ~/.voltttron3/external_platform_discovery.json
{"platform2":{"vip-address":"tcp://127.0.0.2:22916","instance-name":"platform2","serverkey":"Obj-BezSuDBuBh30cn2nIdMH8s9jnFycQII3ctcyllw"},
"platform1":{"vip-address":"tcp://127.0.0.1:22916","instance-name":"platform1","serverkey":"bbBFTKXTmXGzIezVC47scqhwbdlyc3WSBTCuzLDqynA"}}
(voltttron)voltttron@voltttron-VirtualBox ~/git/voltttron-2 $
(voltttron)voltttron@voltttron-VirtualBox ~/git/voltttron-2 $
(voltttron)voltttron@voltttron-VirtualBox ~/git/voltttron-2 $ nano ~/.voltttron3/external_platform_discovery.json
(voltttron)voltttron@voltttron-VirtualBox ~/git/voltttron-2 $
```

After this, you will need to add the server keys of the connecting platforms using the `voltttron-ctl` utility. Type `voltttron-ctl auth add` command on the command prompt and simply hit Enter to select defaults on all fields except **credentials**. Here, we can either add serverkey of connecting platform or type `./*` to allow ALL connections.

Warning: `./*` allows ALL agent and platform connections without authentication.

```
$ voltttron-ctl auth add
domain []:
address []:
user_id []:
capabilities (delimit multiple entries with comma) []:
roles (delimit multiple entries with comma) []:
groups (delimit multiple entries with comma) []:
```

(continues on next page)

(continued from previous page)

```
mechanism [CURVE]:
credentials []: ./.*
comments []:
enabled [True]:
added entry domain=None, address=None, mechanism='CURVE', credentials=u'./.*', user_
↪id=None
```

For more information on authentication see *authentication*.

Once the initial configuration are setup, you can start all the VOLTRON instances in normal mode.

```
$ volttron -v -l 11.log&
```

Next step is to start agents in each platform to observe the multi-platform PubSub communication behavior.

Start Master driver on VOLTRON instance 1

If master driver is not configured to auto start when the instance starts up, we can start it explicitly with this command.

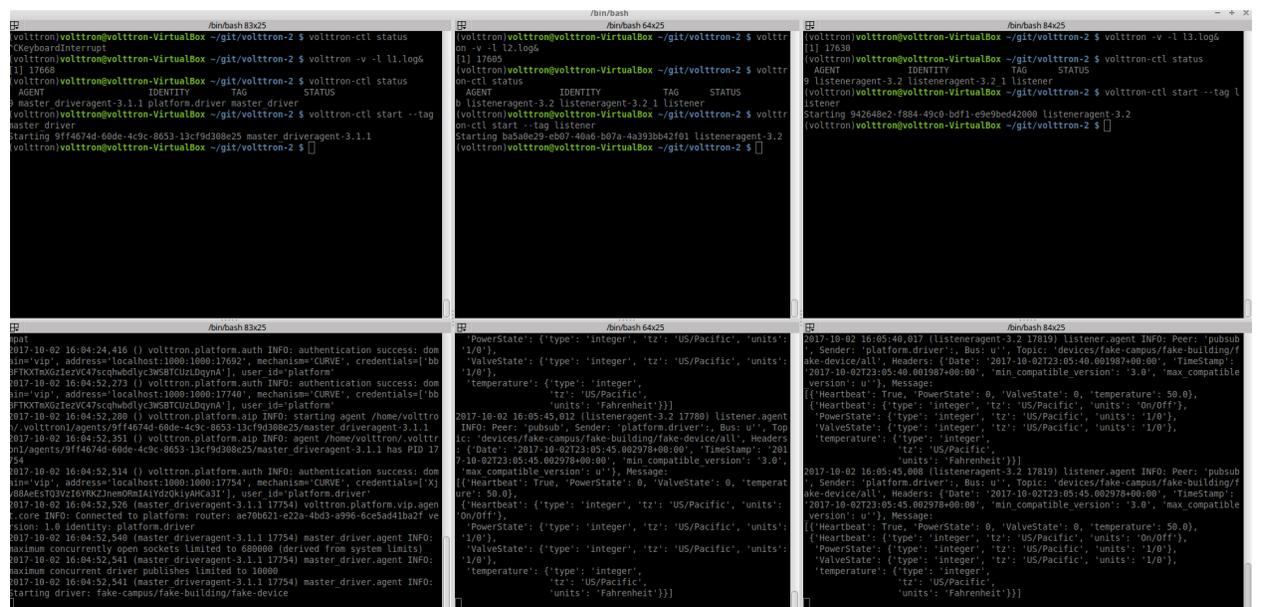
```
$ volttron-ctl start --tag master_driver
```

Start Listener agents on VOLTRON instance 2 and 3

If the listener agent is not configured to auto start when the instance starts up, we can start it explicitly with this command.

```
$ volttron-ctl start --tag listener
```

We should start seeing messages with prefix="devices" in the logs of VOLTRON instances 2 and 3.



Stopping All the Platforms

We can stop all the VOLTRON instances by executing below command in each terminal window.

```
$ voltttron-ctl shutdown --platform
```

Simple Web Agent Walkthrough

A simple web enabled agent that will hook up with a voltttron message bus and allow interaction between it via http. This example agent shows a simple file serving agent, a json-rpc based call, and a websocket based connection mechanism.

Starting VOLTRON Platform

Note: Activate the environment first *active the environment*

In order to start the simple web agent, we need to bind the VOLTRON instance to the a web server. We need to specify the address and the port for the web server. For example, if we want to bind the localhost:8080 as the web server we start the VOLTRON platform as follows:

```
voltttron -vv -l voltttron.log --bind-web-address http://127.0.0.1:8080 &
```

Once the platform is started, we are ready to run the Simple Web Agent.

Running Simple Web Agent

Note: The following assumes the shell is located at the VOLTRON_ROOT.

Copy the following into your shell (save it to a file for executing it again later).

```
python scripts/install-agent.py \
  --agent-source examples/SimpleWebAgent \
  --tag simpleWebAgent \
  --vip-identity webagent \
  --force \
  --start
```

This will create a web server on <http://localhost:8080>. The index.html file under simpleweb/webroot/simpleweb/ can be any html page which binds to the VOLTRON message bus .This provides a simple example of providing a web endpoint in VOLTRON.

Path based registration examples

- Files will need to be in webroot/simpleweb in order for them to be browsed from <http://localhost:8080/simpleweb/index.html>
- Filename is required as we don't currently autoredirect to any default pages

as shown in `self.vip.web.register_path("/simpleweb", os.path.join(WEBROOT))`

The following two examples show the way to call either a jsonrpc (default) endpoint and one that returns a different content-type. With the JSON-RPC example from voltrtron central we only allow post requests, however this is not required.

- Endpoint will be available at `http://localhost:8080/simple/text` `self.vip.web.register_endpoint("/simple/text", self.text)`
- Endpoint will be available at `http://localhost:8080/simple/jsonrpc` `self.vip.web.register_endpoint("/simpleweb/jsonrpc", self.rpcendpoint)`
- Text/html content type specified so the browser can act appropriately like `[("Content-Type", "text/html")]`
- The default response is application/json so our endpoint returns appropriately with a json based response.

Single Machine Deployment

The purpose of this demonstration is to show the process of setting up a simple VOLTTRON instance for use on a single machine.

Install and Build VOLTTRON

First, *install* and *build* VOLTTRON:

For a quick reference:

```
sudo apt-get update
sudo apt-get install build-essential python-dev openssl libssl-dev libevent-dev git
git clone https://github.com/VOLTTRON/volttron/
cd volttron
python2.7 bootstrap.py
```

Activating the VOLTTRON Environment

After the build is complete, activate the VOLTTRON environment.

```
source env/bin/activate
```

Configuring VOLTTRON

The `vcfg` command allows for an easy configuration of the VOLTTRON environment.

Note: To create a simple instance of VOLTTRON, leave the default response, or select yes (y) if prompted for a yes or no response [Y/N]. You must choose a username and password for the VOLTTRON Central admin account.

A set of example responses are included here (username is user, localhost is volttron-pc):

```

(volttron)user@volttron-pc:~/volttron$ vcfg

Your VOLTRON_HOME currently set to: /home/user/.volttron

Is this the volttron you are attempting to setup? [Y]:
What type of message bus (rmq/zmq)? [zmq]:
What is the vip address? [tcp://127.0.0.1]:
What is the port for the vip address? [22916]:
Is this instance web enabled? [N]: y
What is the protocol for this instance? [https]:
Web address set to: https://volttron-pc
What is the port for this instance? [8443]:
Would you like to generate a new web certificate? [Y]:
WARNING! CA certificate does not exist.
Create new root CA? [Y]:

Please enter the following details for web server certificate:
    Country: [US]:
    State: WA
    Location: Richland
    Organization: PNNL
    Organization Unit: VOLTRON
Created CA cert
Creating new web server certificate.
Is this an instance of volttron central? [N]: y
Configuring /home/user/volttron/services/core/VolttronCentral.
Enter volttron central admin user name: <your volttron central admin username here>
Enter volttron central admin password: <your volttron central admin password here>
Retype password: <retype your volttron central admin password here>
Installing volttron central.
Should the agent autostart? [N]: y
Will this instance be controlled by volttron central? [Y]: y
Configuring /home/user/volttron/services/core/VolttronCentralPlatform.
What is the name of this instance? [volttron1]:
Volttron central address set to https://volttron-pc:8443
Should the agent autostart? [N]: y
Would you like to install a platform historian? [N]: y
Configuring /home/user/volttron/services/core/SQLHistorian.
Should the agent autostart? [N]: y
Would you like to install a master driver? [N]: y
Configuring /home/user/volttron/services/core/MasterDriverAgent.
Would you like to install a fake device on the master driver? [N]: y
Should the agent autostart? [N]: y
Would you like to install a listener agent? [N]: y
Configuring examples/ListenerAgent.
Should the agent autostart? [N]: y
Finished configuration!

You can now start the volttron instance.

If you need to change the instance configuration you can edit
the config file is at /home/user/.volttron/config

(volttron)user@volttron-pc:~/volttron$

```

Once this is finished, run VOLTRON and test the new configuration.

Testing VOLTRON

To test that the configuration was successful, start an instance of VOLTRON in the background:

```
./start-voltron
```

Note: This command must be run from the root voltron directory.

Command Line

If the example `vcfg` responses were used, the listener, master_driver, platform_historian, vcp, and vc agents should have all started automatically. This can be checked using `vctl status`.

The output should look similar to this:

```
(voltron)user@voltron-pc:~/voltron$ vctl status
AGENT          IDENTITY          TAG          STATUS          HEALTH
↔HEALTH
8 listeneragent-3.2      listeneragent-3.2_1 listener      running [2810]  GOOD
0 master_driveragent-3.2 platform.driver    master_driver  running [2813]  GOOD
3 sqlhistorianagent-3.7.0 platform.historian platform_historian running [2811]  GOOD
2 vcplatformagent-4.8    platform.agent    vcp           running [2812]  GOOD
9 voltroncentralagent-5.0 voltron.central    vc            running [2808]  GOOD
```

You can further verify that the agents are functioning correctly with `tail -f voltron.log`

VOLTRON Central

Open a web browser and navigate to <https://voltron-pc:8443/vc/index.html>

There may be a message warning about a potential security risk. Check to see if the certificate that was created in `vcfg` is being used. The process below is for firefox.

Note: Chrome does not allow one to accept certificate errors. You will need to use a different browser. Firefox is recommended.



Warning: Potential Security Risk Ahead

Firefox detected a potential security threat and did not continue to `volttron-pc`. If you visit this site, attackers could try to steal information like your passwords, emails, or credit card details.

What can you do about it?

The issue is most likely with the website, and there is nothing you can do to resolve it.

If you are on a corporate network or using anti-virus software, you can reach out to the support teams for assistance. You can also notify the website's administrator about the problem.

[Learn more...](#)

Go Back (Recommended)

Advanced...

Report errors like this to help Mozilla identify and block malicious sites



Someone could be trying to impersonate the site and you should not continue.

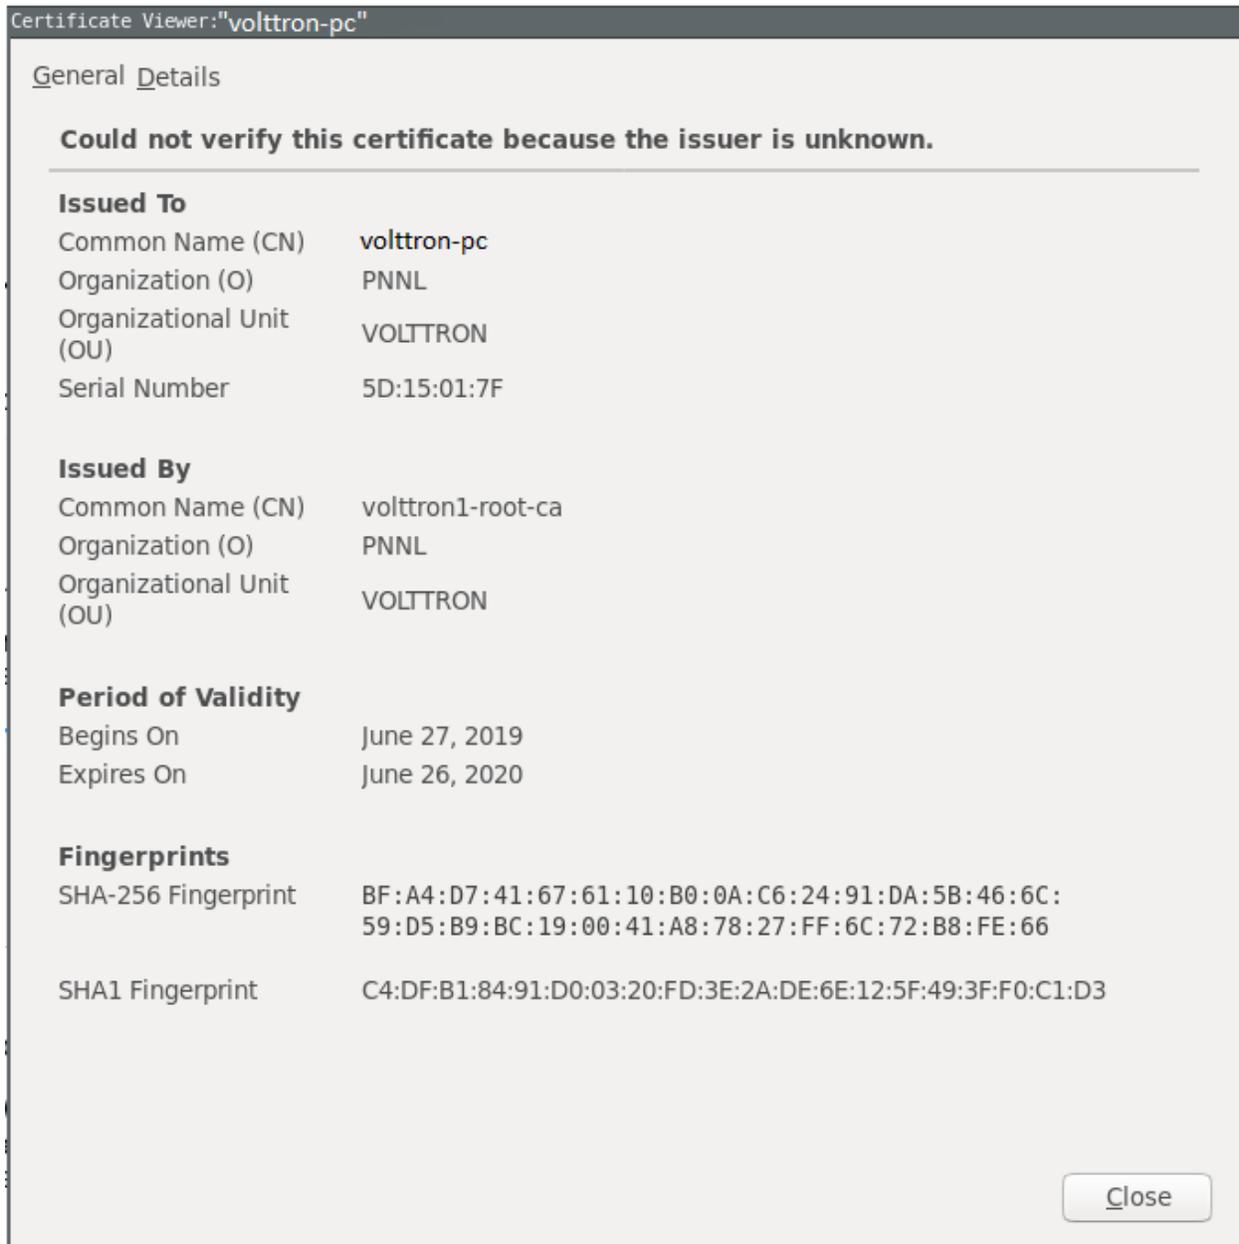
Websites prove their identity via certificates. Firefox does not trust `volttron-pc :8443` because its certificate issuer is unknown, the certificate is self-signed, or the server is not sending the correct intermediate certificates.

Error code: [SEC_ERROR_UNKNOWN_ISSUER](#)

[View Certificate](#)

Go Back (Recommended)

Accept the Risk and Continue



Someone could be trying to impersonate the site and you should not continue.

Websites prove their identity via certificates. Firefox does not trust `voltron-pc :8443` because its certificate issuer is unknown, the certificate is self-signed, or the server is not sending the correct intermediate certificates.

Error code: [SEC_ERROR_UNKNOWN_ISSUER](#)

[View Certificate](#)



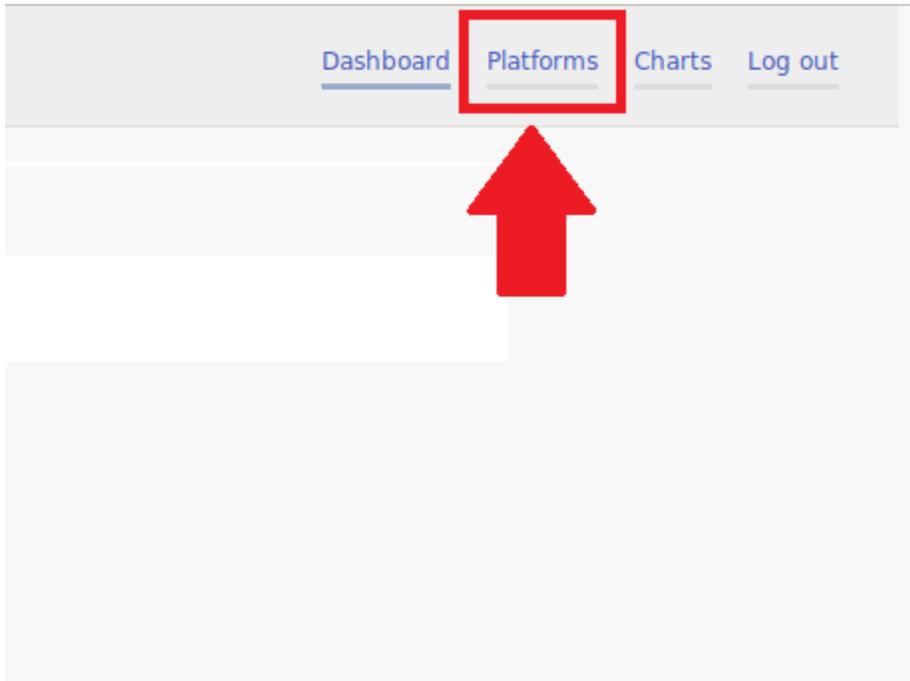
[Go Back \(Recommended\)](#)

[Accept the Risk and Continue](#)

Log in using the username and password you created during the `vetl` prompt.

The image shows a login page for VOLTRON™ Central BETA. The page has a light gray background. At the top, the text "VOLTRON™ Central" is displayed in a large, bold, dark blue font, with "BETA" in a smaller, gray font below it. Underneath, it says "Funded by DOE EERE BTO" in a smaller, dark blue font. In the center, there are three input fields: "Username", "Password", and "Log in". The "Username" and "Password" fields are white with a light gray border, and the "Log in" button is a solid gray color.

Once you have logged in, click on the Platforms tab in the upper right corner of the window.



Once in the Platforms screen, click on the name of the platform.



You will now see a list of agents. They should all be running.

Platforms / volttron1 (dm9sdHRyb24xLnBsYXRmb3JtLmFnZW50)

Agents

Name	Identity	UUID	Status	Action
listeneragent-3.2	listeneragent-3.2_1	ff8b93d0-b499-4d66-8304-94ac9fd5d9b4	Running (PID 3842)	<input type="button" value="Stop"/> <input type="button" value="Remove"/>
master_driveragent-3.2	platform.driver	30bf6384-3bde-4e28-8bf7-6695448f7b7d	Running (PID 3847)	<input type="button" value="Stop"/> <input type="button" value="Remove"/>
sqlhistorianagent-3.7.0	platform.historian	769d570d-e9b5-4263-8d07-31416ccec5c8	Running (PID 3845)	<input type="button" value="Stop"/> <input type="button" value="Remove"/>
vcplatformagent-4.8	platform.agent	f03b0ade-d161-4ddb-946b-30e129725584	Running (PID 3844)	<input type="button" value="Stop"/> <input type="button" value="Remove"/>
volttroncentralagent-5.0	volttron.central	63f67dd6-a27a-4064-949f-7c6649c43f73	Running (PID 3846)	<input type="button" value="Stop"/> <input type="button" value="Remove"/>

Install agents

No files selected.

For more information on VOLTRON Central, please see:

- [VOLTRON Central Management](#)
- [VOLTRON Central Demo](#)

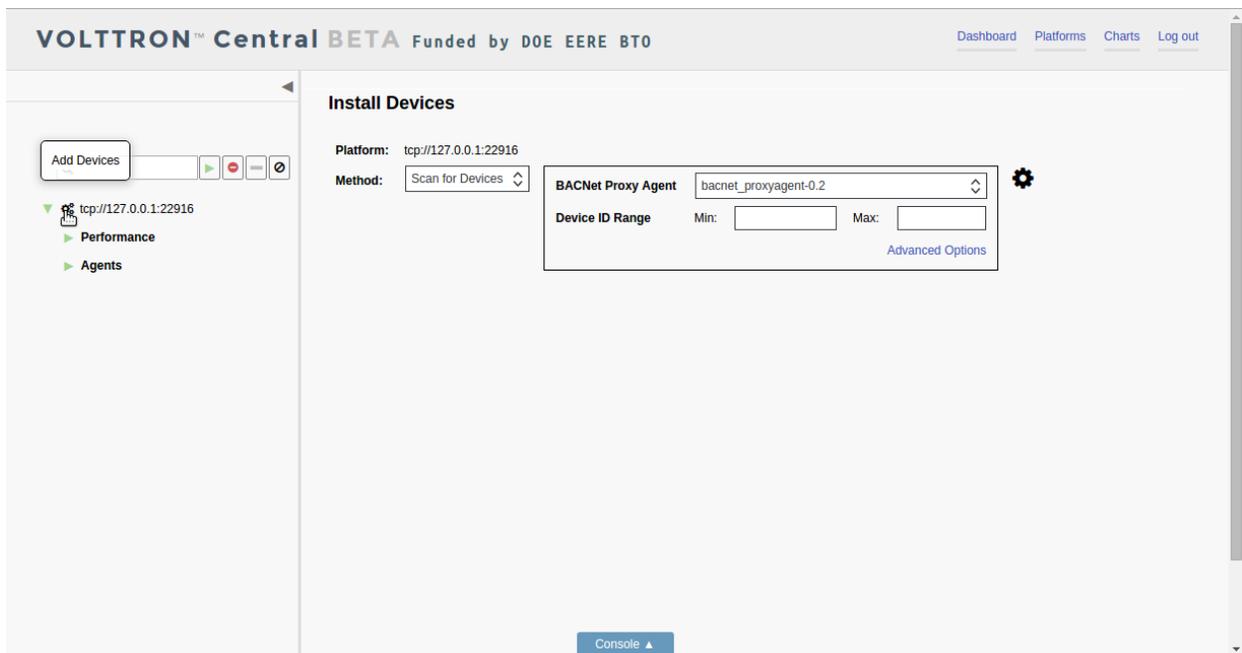
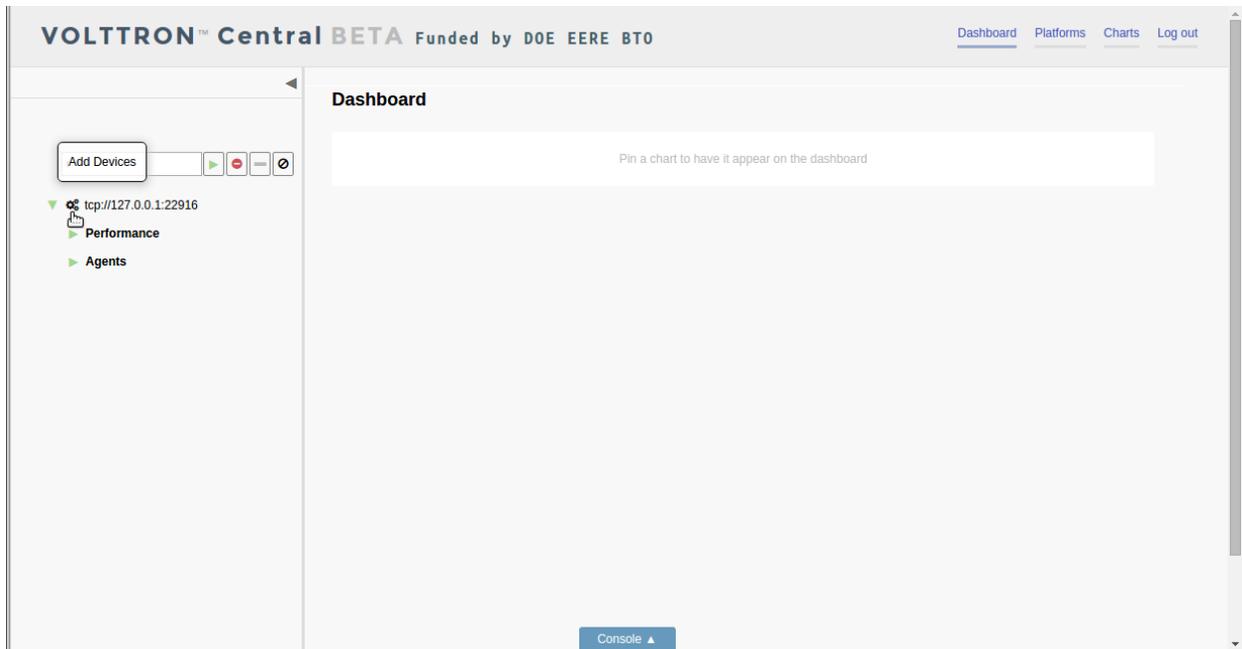
Device Configuration in VOLTRON Central

Devices in your network can be detected and configured through the VOLTRON Central UI. The current version of VOLTRON enables device detection and configuration for BACnet devices. The following sections describe the processes involved with performing scans to detect physical devices and get their points, and configuring them as virtual devices installed on VOLTRON instances.

- [Launching Device Configuration](#)
- [Scanning for Devices](#)
- [Scanning for Points](#)
- [Registry Configuration File](#)
- [Additional Attributes](#)
- [Quick Edit Features](#)
- [Keyboard Commands](#)
- [Registry Preview](#)
- [Registry Configuration Options](#)
- [Reloading Device Points](#)
- [Device Configuration Form](#)
- [Configuring Subdevices](#)
- [Reconfiguring Devices](#)
- [Exporting Registry Configuration Files](#)

Launching Device Configuration

To begin device configuration in VOLTRON Central, extend the side panel on the left and find the cogs button next to the platform instance you want to add a device to. Click the cogs button to launch the device configuration feature.



Currently the only method of adding devices is to conduct a scan to detect BACnet devices. A BACnet Proxy Agent must be running in order to do the scan. If more than one BACnet Proxy is installed on the platform, choose the one that will be used for the scan.

The scan can be conducted using default settings that will search for all physical devices on the network. However, optional settings can be used to focus on specific devices or change the duration of the scan. Entering a range of device IDs will limit the scan to return only devices with IDs in that range. Advanced options include the ability to specify the IP address of a device to detect as well as the ability to change the duration of the scan from the default of five seconds.

Scanning for Devices

To start the scan, click the large cog button to the right of the scan settings.

VOLTRON™ Central BETA Funded by DOE EERE BT0

Dashboard Platforms Charts Log out

Install Devices Find Devices

Platform: tcp://127.0.0.1:22916

Method: Scan for Devices

BACNet Proxy Agent: bacnet_proxyagent-0.2

Device ID Range: Min: Max:

Advanced Options

Console

Devices that are detected will appear in the space below the scan settings. Scanning can be repeated at any time by clicking the large cog button again.

VOLTRON™ Central BETA Funded by DOE EERE BT0

Dashboard Platforms Charts Log out

Install Devices

Platform: tcp://127.0.0.1:22916

Method: Scan for Devices

BACNet Proxy Agent: bacnet_proxyagent-0.2

Device ID Range: Min: Max:

Advanced Options

Address	Name	Description	Device ID	Vendor ID	Vendor	Type
▶ 10.0.2.4	Betelgeuse		70	15	Cornell University	bacnet
▶ 10.0.2.6	MacGyver		1000	19	TAC AB	bacnet

Console

Scanning for Points

Another scan can be performed on each physical device to retrieve its available points. This scan is initiated by clicking the triangle next to the device in the list. The first time the arrow is clicked, it initiates the scan. After the points are retrieved, the arrow becomes a hide-and-show toggle button and won't reinitiate scanning the device.

The screenshot shows the VOLTRON Central BETA interface. The main content area is titled "Install Devices". It features a "Platform" field set to "tcp://127.0.0.1:22916" and a "Method" dropdown set to "Scan for Devices". Below this, there are configuration options for the "BACNet Proxy Agent" (set to "bacnet_proxyagent-0.2") and a "Device ID Range" section with "Min" and "Max" input fields. A "Get Device Points" button is visible. A table displays the results of the scan:

Address	Name	Description	Device ID	Vendor ID	Vendor	Type
10.0.2.4	Betelgeuse		70	15	Cornell University	bacnet
10.0.2.6	MacGyver		1000	19	TAC AB	bacnet

The interface also includes a sidebar with a search bar and a tree view showing the current platform instance. A "Console" button is located at the bottom of the main content area.

After the points have been retrieved once, the only way to scan the same device for points again is to relaunch the device configuration process from the start by clicking on the small cogs button next to the platform instance in the panel tree.

Registry Configuration File

The registry configuration determines which points on the physical device will be associated with the virtual device that uses that particular registry configuration. The registry configuration determines which points' data will be published to the message bus and recorded by the historian, and it determines how the data will be presented.

When all the points on the device have been retrieved, the points are loaded into the registry configuration editor. There, the points can be modified and selected to go into the registry configuration file for a device.

Each row in the registry configuration editor represents a point, and each cell in the row represents an attribute of the point.

Only points that have been selected will be included in the registry configuration file. To select a point, check the box next to the point in the editor.

VOLTRON™ Central BETA Funded by DOE EERE BTO Dashboard Platforms Charts Log out

Search:

tcp://127.0.0.1:22916

Install Devices

Platform: tcp://127.0.0.1:22916

Method: Scan for Devices

BACNet Proxy Agent: bacnet_proxyagent-0.2

Device ID Range: Min: Max:

[Advanced Options](#)

Address	Name	Description	Device ID	Vendor ID	Vendor	Type
10.0.2.4	Betelgeuse		70	15	Cornell University	bacnet

Volttron Point Name	Writable	Units
<input type="checkbox"/> ReturnAirTemperature	FALSE	UNKNOWN UNITS
<input checked="" type="checkbox"/> ReturnAirHumidity	FALSE	UNKNOWN UNITS
<input type="checkbox"/> CoolingValveOutputCommand	FALSE	UNKNOWN UNITS
<input type="checkbox"/> MixedAirTemperature	FALSE	UNKNOWN UNITS
<input type="checkbox"/> OutdoorAirHumidity	FALSE	UNKNOWN UNITS
<input type="checkbox"/> PreheatTemperature	FALSE	UNKNOWN UNITS
<input type="checkbox"/> DischargeAirTemperature	FALSE	UNKNOWN UNITS
<input type="checkbox"/> DischargeAirStaticPressure	FALSE	UNKNOWN UNITS

[Console](#)

VOLTRON™ Central BETA Funded by DOE EERE BTO Dashboard Platforms Charts Log out

Search:

tcp://127.0.0.1:22916

Install Devices

Platform: tcp://127.0.0.1:22916

Method: Scan for Devices

BACNet Proxy Agent: bacnet_proxyagent-0.2

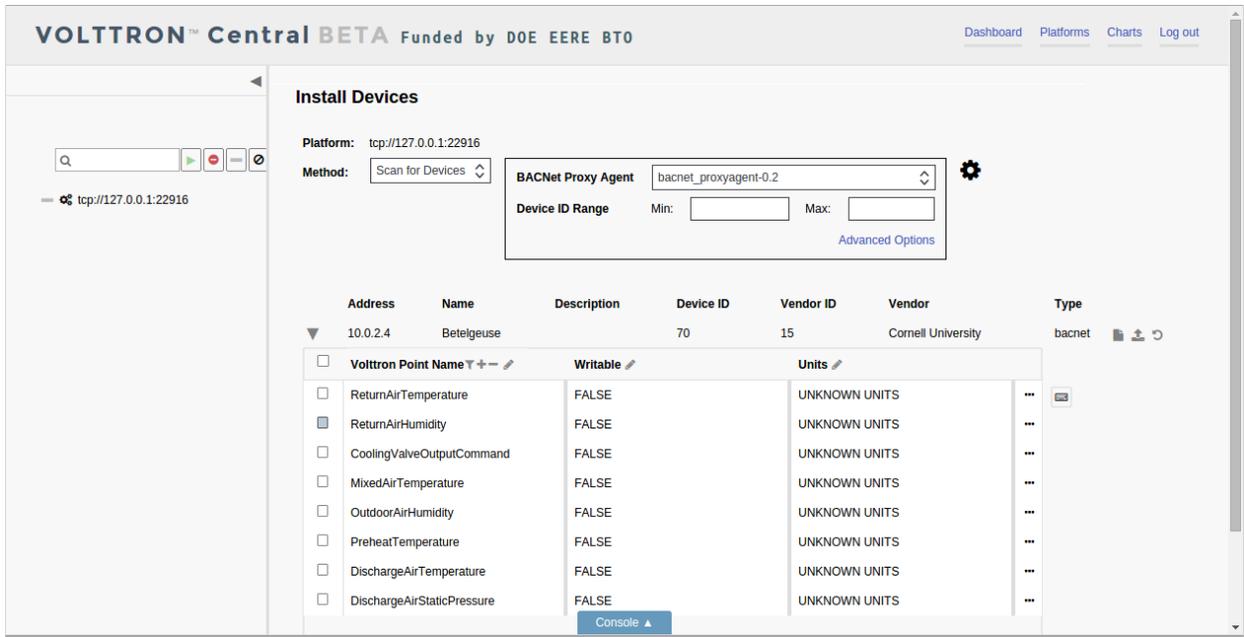
Device ID Range: Min: Max:

[Advanced Options](#)

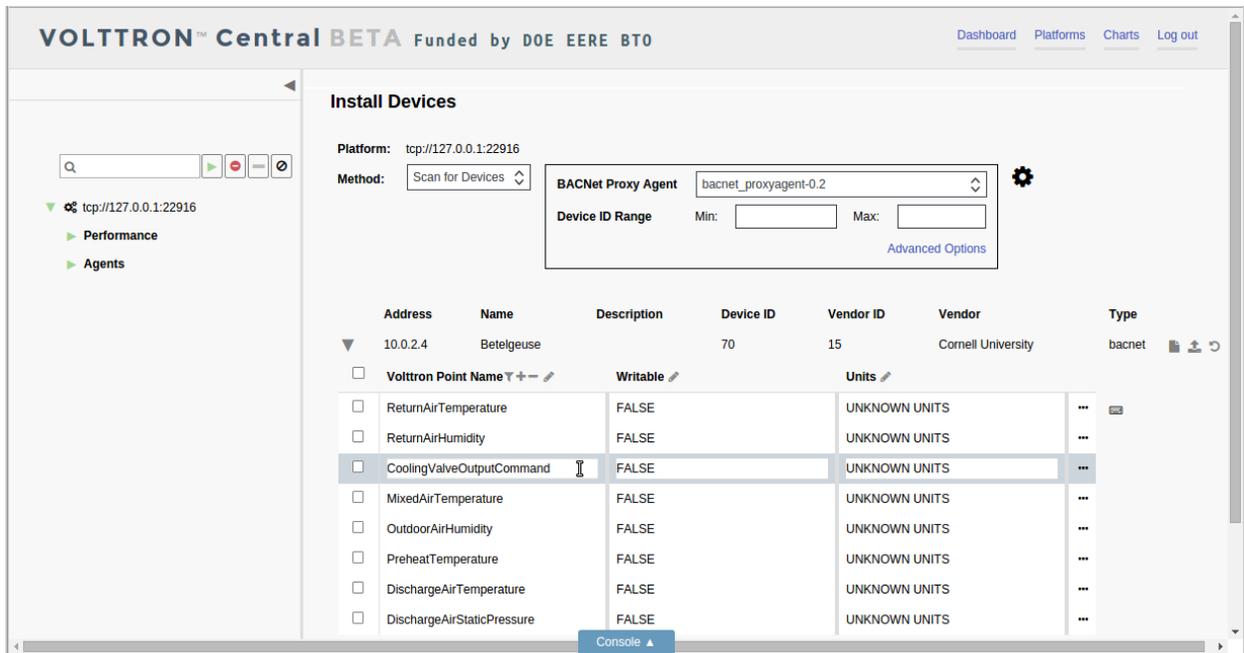
Address	Name	Description	Device ID	Vendor ID	Vendor	Type
10.0.2.4	Betelgeuse		70	15	Cornell University	bacnet

Volttron Point Name	Writable	Units
<input type="checkbox"/> ReturnAirTemperature	FALSE	UNKNOWN UNITS
<input checked="" type="checkbox"/> ReturnAirHumidity	FALSE	UNKNOWN UNITS
<input type="checkbox"/> CoolingValveOutputCommand	FALSE	UNKNOWN UNITS
<input type="checkbox"/> MixedAirTemperature	FALSE	UNKNOWN UNITS
<input type="checkbox"/> OutdoorAirHumidity	FALSE	UNKNOWN UNITS
<input type="checkbox"/> PreheatTemperature	FALSE	UNKNOWN UNITS
<input type="checkbox"/> DischargeAirTemperature	FALSE	UNKNOWN UNITS
<input type="checkbox"/> DischargeAirStaticPressure	FALSE	UNKNOWN UNITS

[Console](#)



Type directly in a cell to change an attribute value for a point.



Additional Attributes

The editor's default view shows the attributes that are most likely to be changed during configuration: the VOLTTRON point name, the writable setting, and the units. Other attributes are present but not shown in the default view. To see the entire set of attributes for a point, click the Edit Point button (the three dots) at the end of the point row.

VOLTRON™ Central BETA Funded by DOE EERE BT0

Dashboard Platforms Charts Log out

Install Devices

Platform: tcp://127.0.0.1:22916

Method: Scan for Devices

BACNet Proxy Agent: bacnet_proxyagent-0.2

Device ID Range: Min: [] Max: []

Advanced Options

Address	Name	Description	Device ID	Vendor ID	Vendor	Type
10.0.2.4	Beteigeuse		70	15	Cornell University	bacnet
<input type="checkbox"/> Volttron Point Name		<input type="checkbox"/> Writable	<input type="checkbox"/> Units			
<input type="checkbox"/>	ReturnAirTemperature	FALSE	UNKNOWN UNITS			
<input type="checkbox"/>	ReturnAirHumidity	FALSE	UNKNOWN UNITS			
<input type="checkbox"/>	CoolingValveOutputCommand	FALSE	UNKNOWN UNITS			
<input type="checkbox"/>	MixedWaterTemperature	FALSE	UNKNOWN UNITS			
<input type="checkbox"/>	OutdoorAirHumidity	FALSE	UNKNOWN UNITS			
<input type="checkbox"/>	PreheatTemperature	FALSE	UNKNOWN UNITS			
<input type="checkbox"/>	DischargeAirTemperature	FALSE	UNKNOWN UNITS			
<input type="checkbox"/>	DischargeAirStaticPressure	FALSE	UNKNOWN UNITS			

Console

In the window that opens, point attributes can be changed by typing in the fields and clicking the Apply button.

VOLTRON™ Central BETA Funded by DOE EERE BT0

Dashboard Platforms Charts Log out

Point

Point	Value	Show in Table
Index	3000124	<input type="checkbox"/>
Reference Point Name	ReturnAirHumidity	<input type="checkbox"/>
Volttron Point Name	ReturnAirHumidity	<input checked="" type="checkbox"/>
Unit Details	No limits.	<input type="checkbox"/>
BACnet Object Type	analogInput	<input type="checkbox"/>
Notes		<input type="checkbox"/>
Writable	FALSE	<input checked="" type="checkbox"/>
Units	UNKNOWN UNITS	<input checked="" type="checkbox"/>
Property	presentValue	<input type="checkbox"/>

Cancel Apply

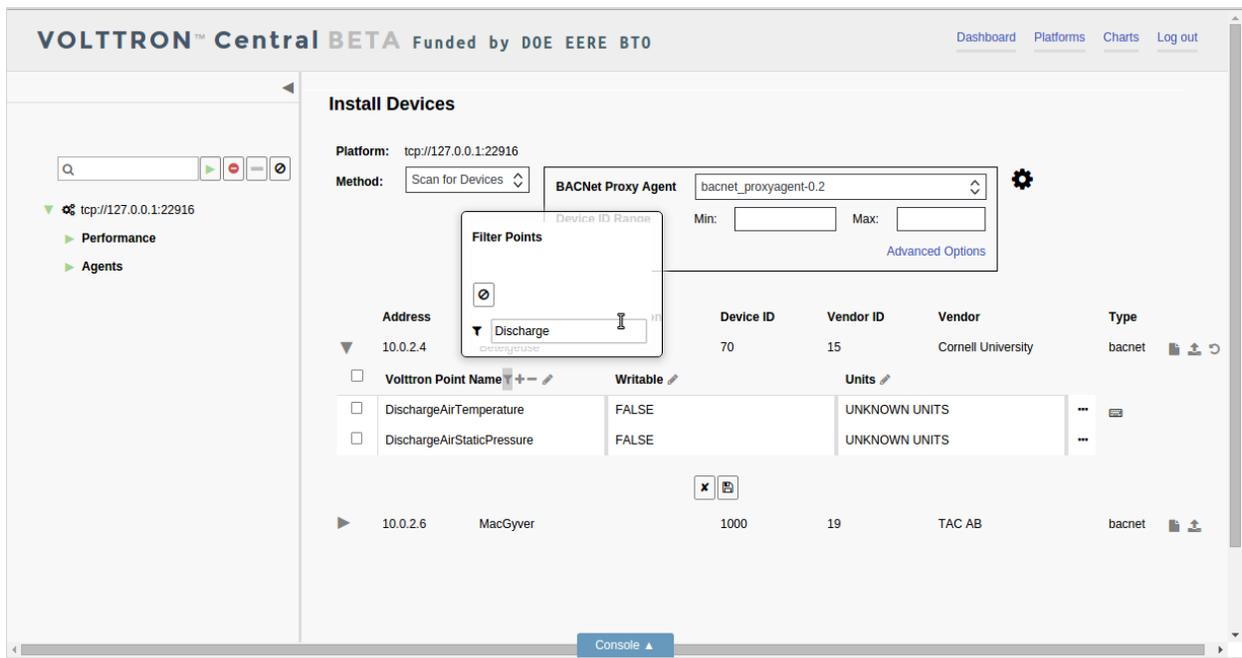
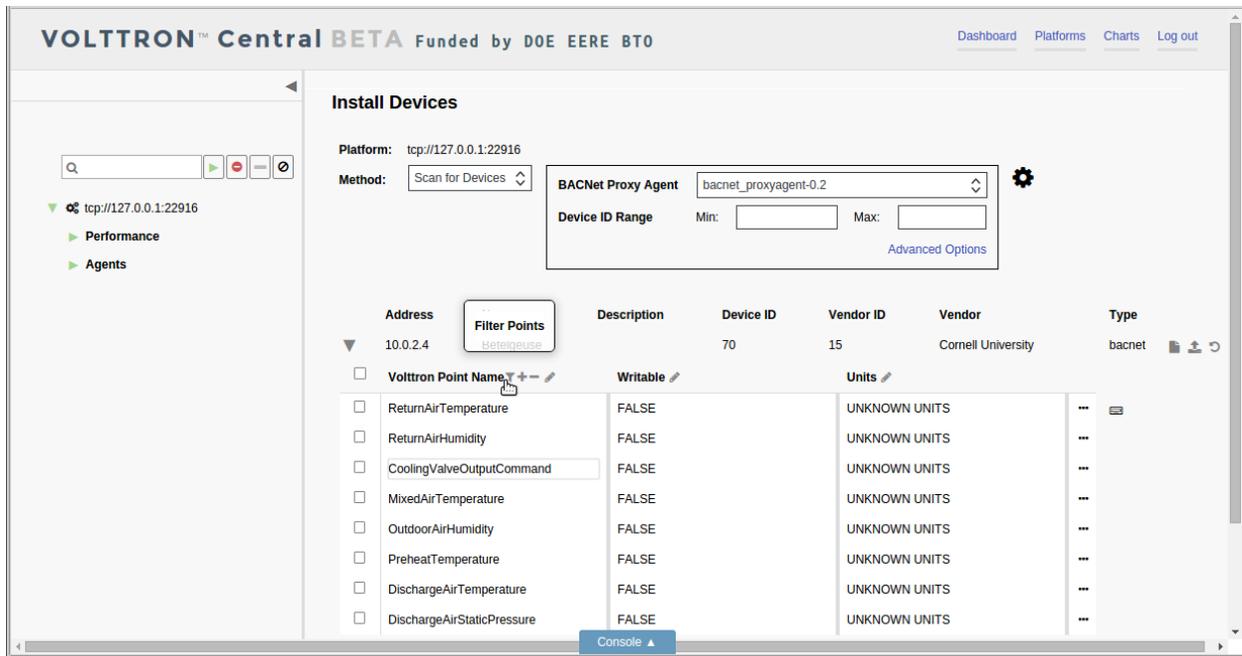
Console

Checking or unchecking the “Show in Table” box for an attribute will add or remove it as a column in the registry configuration editor.

Quick Edit Features

Several quick-edit features are available in the registry configuration editor.

The list of points can be filtered based on values in the first column by clicking the filter button in the first column’s header and entering a filter term.



The filter feature allows points to be edited, selected, or deselected more quickly by narrowing down potentially large lists of points. However, the filter doesn't select points, and if the registry configuration is saved while a filter is applied, any selected points not included in the filter will still be included in the registry file.

To clear the filter, click on the Clear Filter button in the filter popup.

VOLTRON™ Central BETA Funded by DOE EERE BTO

Dashboard Platforms Charts Log out

Install Devices

Platform: tcp://127.0.0.1:22916

Method: Scan for Devices

BACNet Proxy Agent: bacnet_proxyagent-0.2

Device ID Range: Min: Max:

Advanced Options

Filter Points

Clear Filter

Address	Device ID	Vendor ID	Vendor	Type
10.0.2.4	70	15	Cornell University	bacnet

Voltron Point Name	Writable	Units
<input type="checkbox"/> ReturnAirTemperature	FALSE	UNKNOWN UNITS
<input type="checkbox"/> ReturnAirHumidity	FALSE	UNKNOWN UNITS
<input type="checkbox"/> CoolingValveOutputCommand	FALSE	UNKNOWN UNITS
<input type="checkbox"/> MixedAirTemperature	FALSE	UNKNOWN UNITS
<input type="checkbox"/> OutdoorAirHumidity	FALSE	UNKNOWN UNITS
<input type="checkbox"/> PreheatTemperature	FALSE	UNKNOWN UNITS
<input type="checkbox"/> DischargeAirTemperature	FALSE	UNKNOWN UNITS
<input type="checkbox"/> DischargeAirStaticPressure	FALSE	UNKNOWN UNITS

Console

To add a new point to the points listed in the registry configuration editor, click on the Add Point button in the header of the first column.

VOLTRON™ Central BETA Funded by DOE EERE BTO

Dashboard Platforms Charts Log out

Install Devices

Platform: tcp://127.0.0.1:22916

Method: Scan for Devices

BACNet Proxy Agent: bacnet_proxyagent-0.2

Device ID Range: Min: Max:

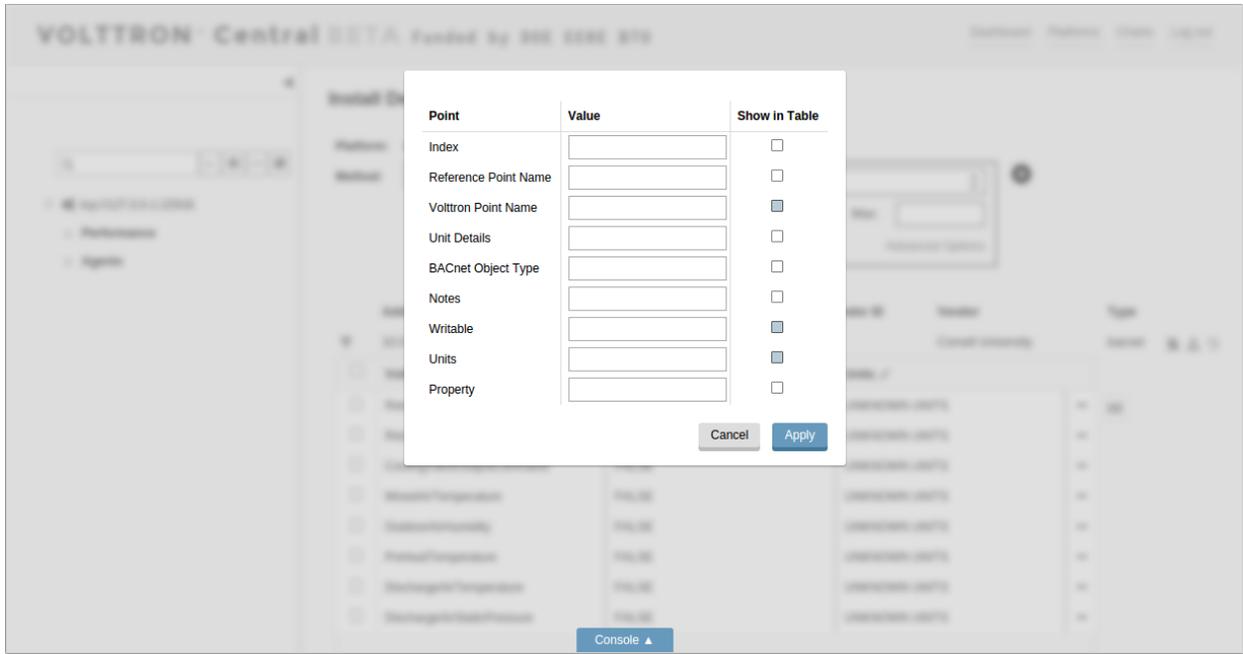
Advanced Options

Add New Point

Address	Description	Device ID	Vendor ID	Vendor	Type
10.0.2.4		70	15	Cornell University	bacnet

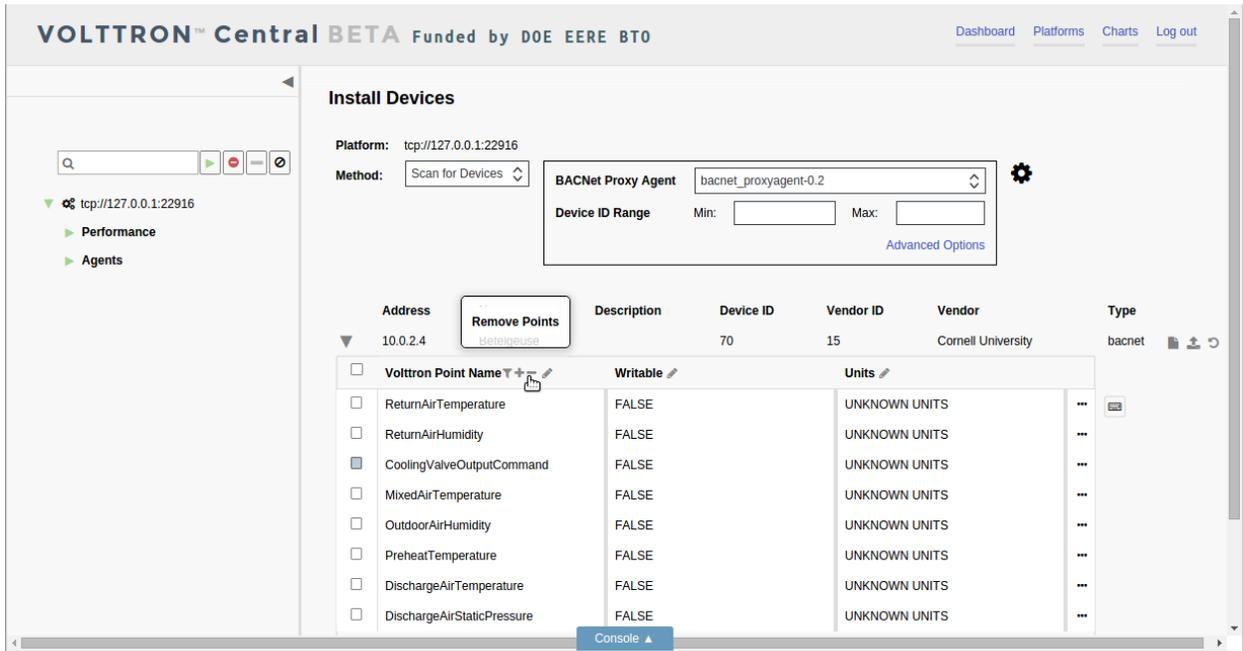
Voltron Point Name	Writable	Units
<input type="checkbox"/> ReturnAirTemperature	FALSE	UNKNOWN UNITS
<input type="checkbox"/> ReturnAirHumidity	FALSE	UNKNOWN UNITS
<input type="checkbox"/> CoolingValveOutputCommand	FALSE	UNKNOWN UNITS
<input type="checkbox"/> MixedAirTemperature	FALSE	UNKNOWN UNITS
<input type="checkbox"/> OutdoorAirHumidity	FALSE	UNKNOWN UNITS
<input type="checkbox"/> PreheatTemperature	FALSE	UNKNOWN UNITS
<input type="checkbox"/> DischargeAirTemperature	FALSE	UNKNOWN UNITS
<input type="checkbox"/> DischargeAirStaticPressure	FALSE	UNKNOWN UNITS

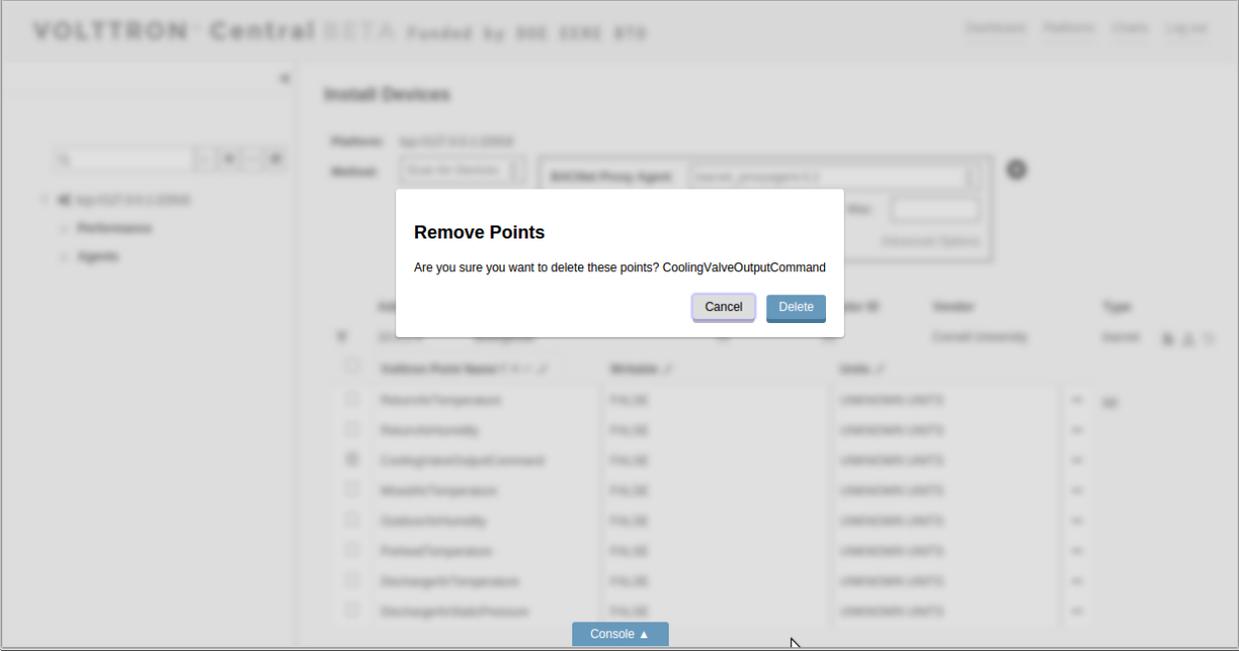
Console



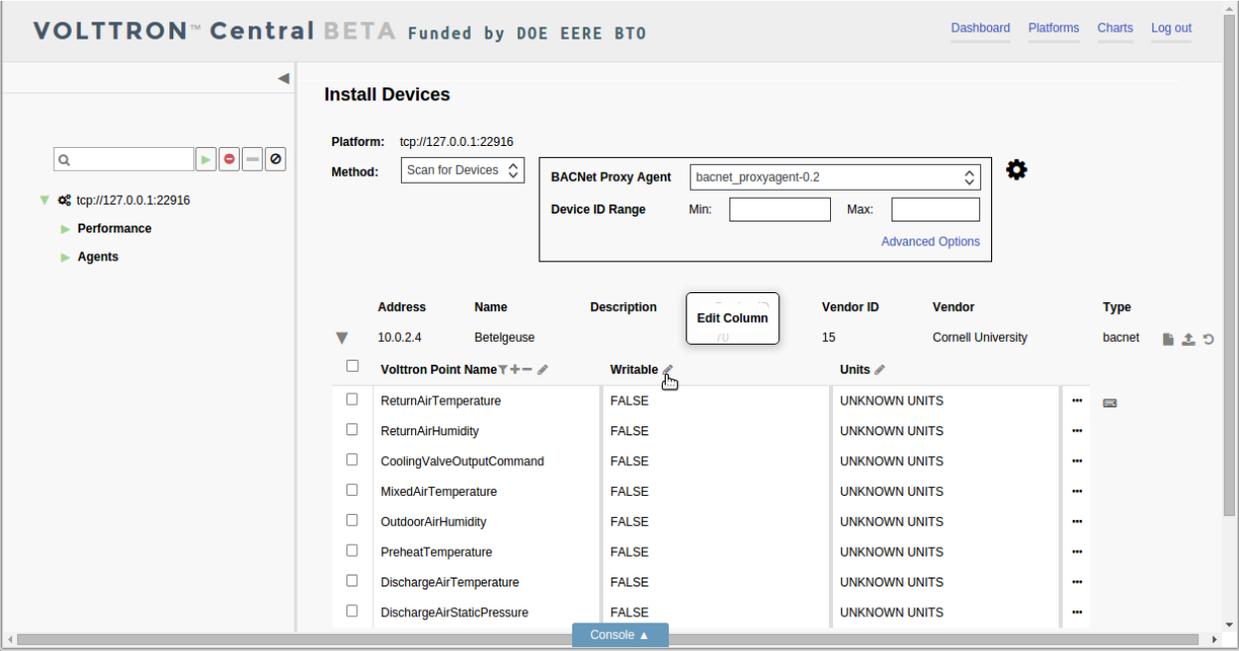
Provide attribute values, and click the Apply button to add the new point, which will be appended to the bottom of the list.

To remove points from the list, select the points and click the Remove Points button in the header of the first column.

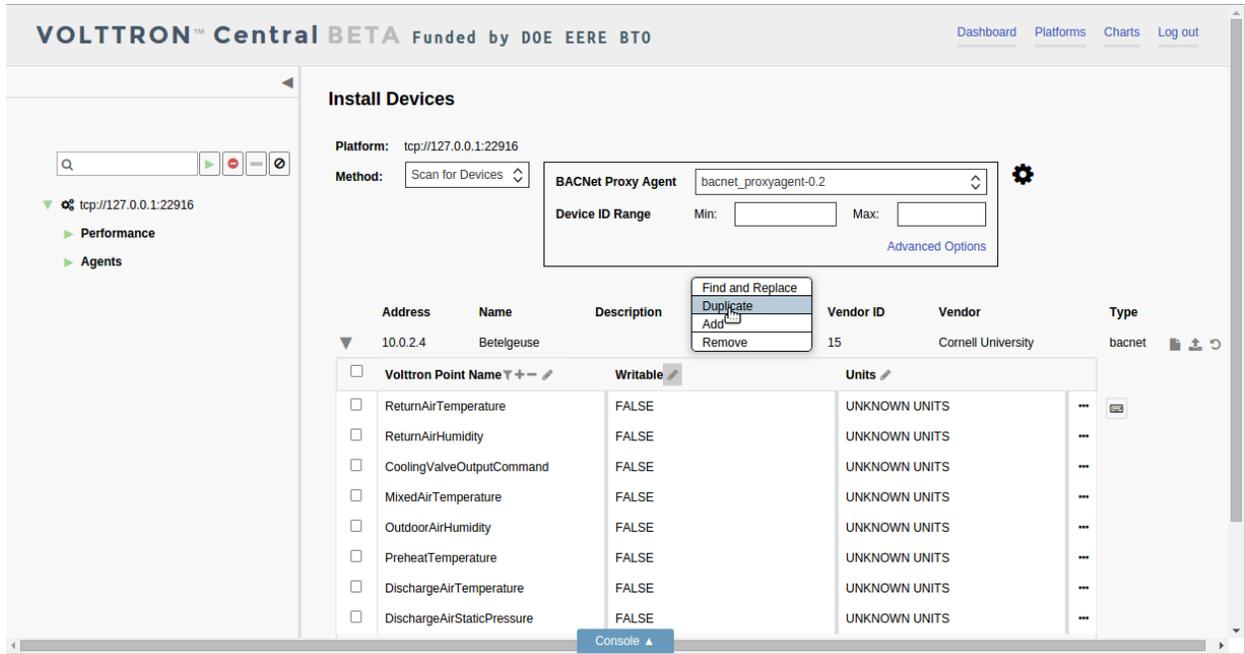




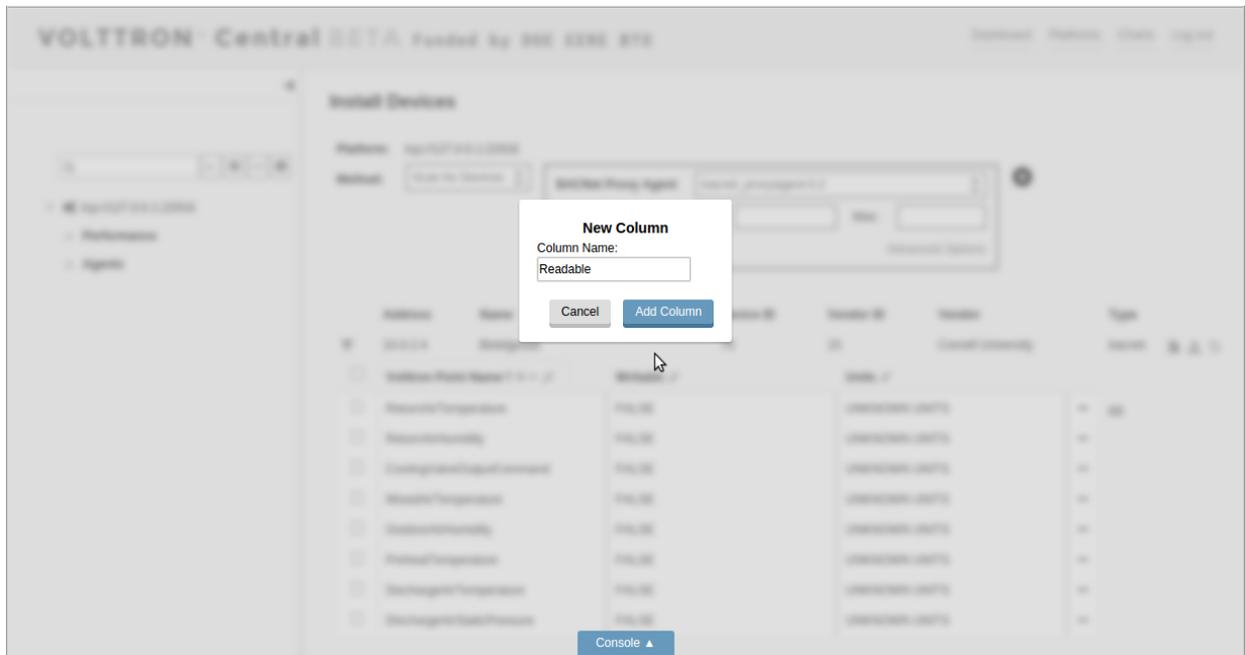
Each column has an Edit Column button in its header.



Click on the button to display a popup menu of operations to perform on the column. The options include inserting a blank new column, duplicating an existing column, removing a column, or searching for a value within a column.

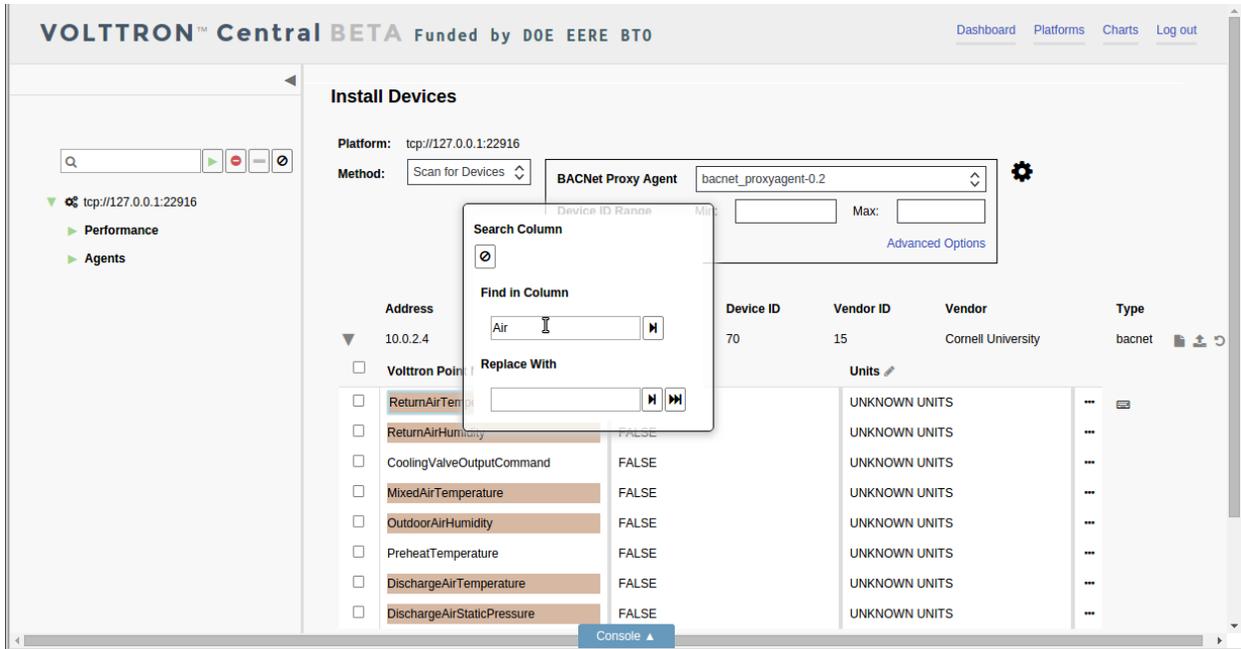


A duplicate or new column has to be given a unique name.



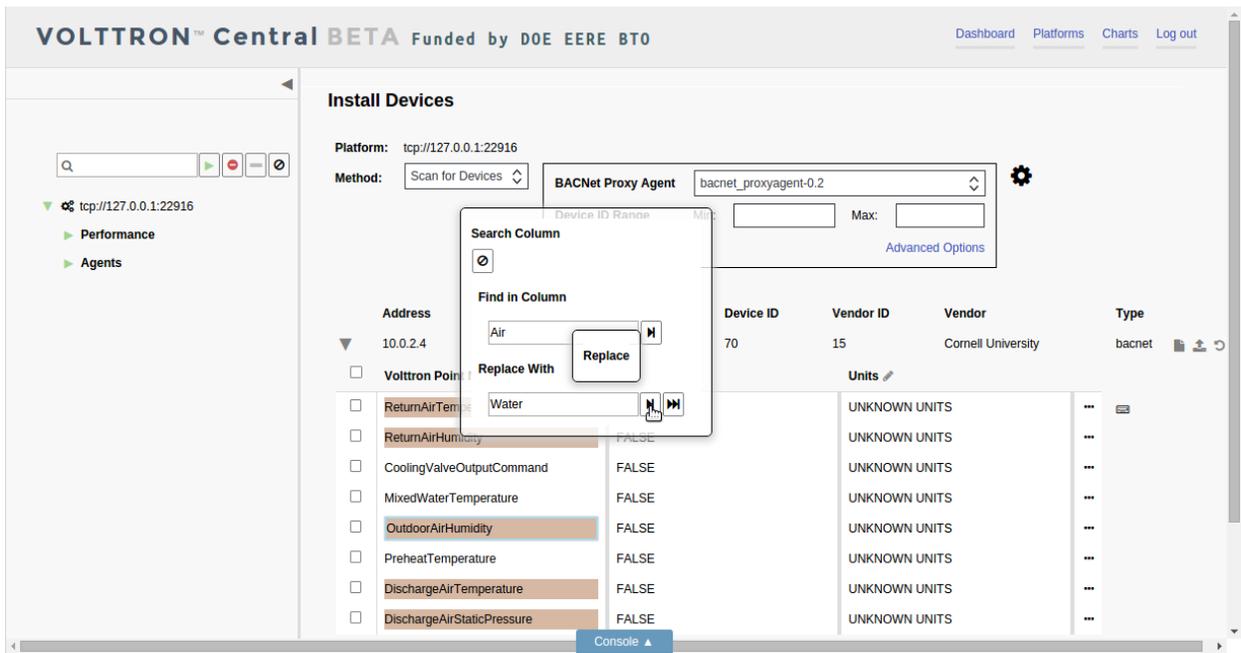
To search for values in a column, choose the Find and Replace option in the popup menu.

Type the term to search for, and click the Find Next button to highlight all the matched fields in the column.



Click the Find Next button again to advance the focus down the list of matched terms.

To quickly replace the matched term in the cell with focus, type a replacement term, and click on the Replace button.



To replace all the matched terms in the column, click on the Replace All button. Click the Clear Search button to end the search.

Keyboard Commands

Some keyboard commands are available to expedite the selection or de-selection of points. To initiate use of the keyboard commands, strike the Control key on the keyboard. For keyboard commands to be activated, the registry

configuration editor has to have focus, which comes from interacting with it. But the commands won't be activated if the cursor is in a typable field.

If the keyboard commands have been successfully activated, a faint highlight will appear over the first row in the registry configuration editor.

The screenshot shows the VOLTRON Central BETA web interface. The header includes the logo and navigation links: Dashboard, Platforms, Charts, and Log out. The main content area is titled 'Install Devices' and shows the following configuration:

- Platform: tcp://127.0.0.1:22916
- Method: Scan for Devices
- BACNet Proxy Agent: bacnet_proxyagent-0.2
- Device ID Range: Min: [] Max: []

Below the configuration, there is a table of devices. The first row is highlighted:

Address	Name	Description	Device ID	Vendor ID	Vendor	Type
10.0.2.4	Betelgeuse		70	15	Cornell University	bacnet

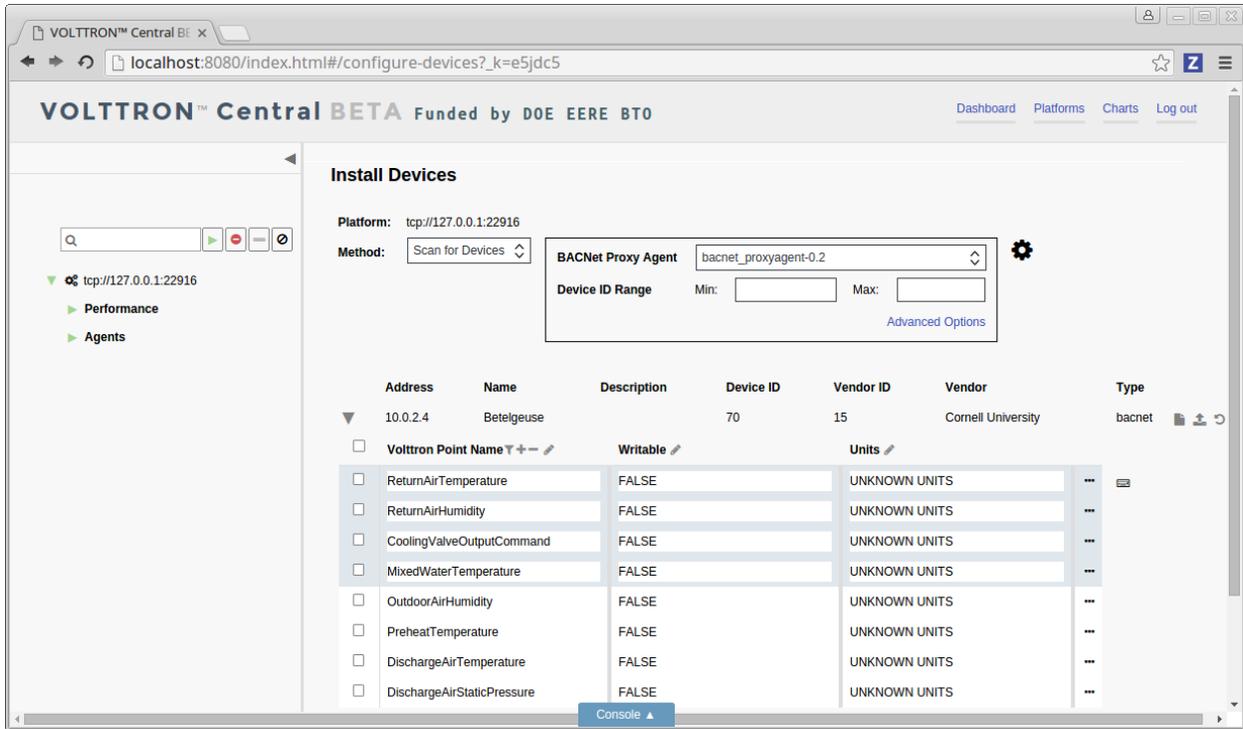
Below the table, there is a section for 'Voltron Point Name' and 'Writable' properties. The first row is highlighted:

Voltron Point Name	Writable	Units
ReturnAirTemperature	FALSE	UNKNOWN UNITS
ReturnAirHumidity	FALSE	UNKNOWN UNITS
CoolingValveOutputCommand	FALSE	UNKNOWN UNITS
MixedWaterTemperature	FALSE	UNKNOWN UNITS
OutdoorAirHumidity	FALSE	UNKNOWN UNITS
PreheatTemperature	FALSE	UNKNOWN UNITS
DischargeAirTemperature	FALSE	UNKNOWN UNITS
DischargeAirStaticPressure	FALSE	UNKNOWN UNITS

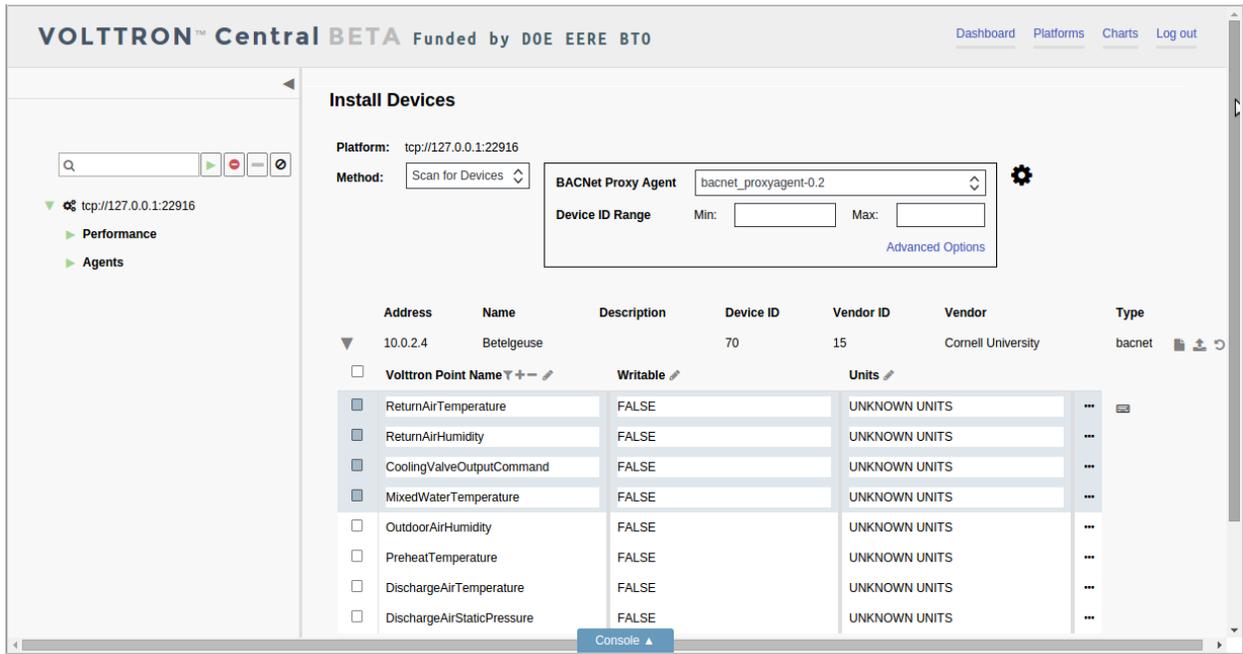
The interface also includes a search bar on the left and a 'Console' button at the bottom.

Keyboard commands are deactivated when the mouse cursor moves over the configuration editor. If unintentional deactivation occurs, strike the Control key again to reactivate the commands.

With keyboard commands activated, the highlighted row can be advanced up or down by striking the up or down arrow on the keyboard. A group of rows can be highlighted by striking the up or down arrow while holding down the Shift key.



To select the highlighted rows, strike the Enter key.



Striking the Enter key with rows highlighted will also deselect any rows that were already selected.

Click on the Keyboard Shortcuts button to show a popup list of the available keyboard commands.

VOLTRON™ Central BETA Funded by DOE EERE BT0 Dashboard Platforms Charts Log out

Search

- tcp://127.0.0.1:22916
 - Performance
 - Agents

Install Devices

Platform: tcp://127.0.0.1:22916

Method: Scan for Devices

BACNet Proxy Agent: bacnet_proxyagent-0.2

Device ID Range: Min: Max:

[Advanced Options](#)

Address	Name	Description	Device ID	Vendor ID	Vendor	Keyboard Shortcuts
10.0.2.4	Betelgeuse		70	15	Cornell University	
<input type="checkbox"/> Volttron Point Name <small>↑ + -</small>			Writable		Units	
<input type="checkbox"/>	ReturnAirTemperature	FALSE		UNKNOWN UNITS		...
<input type="checkbox"/>	ReturnAirHumidity	FALSE		UNKNOWN UNITS		...
<input type="checkbox"/>	CoolingValveOutputCommand	FALSE		UNKNOWN UNITS		...
<input type="checkbox"/>	MixedWaterTemperature	FALSE		UNKNOWN UNITS		...
<input type="checkbox"/>	OutdoorAirHumidity	FALSE		UNKNOWN UNITS		...
<input type="checkbox"/>	PreheatTemperature	FALSE		UNKNOWN UNITS		...
<input type="checkbox"/>	DischargeAirTemperature	FALSE		UNKNOWN UNITS		...
<input type="checkbox"/>	DischargeAirStaticPressure	FALSE		UNKNOWN UNITS		...

Console

VOLTRON™ Central BETA Funded by DOE EERE BT0 Dashboard Platforms Charts Log out

Search

- tcp://127.0.0.1:22916
 - Performance
 - Agents

Install Devices

Platform: tcp://127.0.0.1:22916

Method: Scan for Devices

BACNet Proxy Agent: bacnet_proxyagent-0.2

Device ID Range: Min: Max:

[Advanced Options](#)

Address	Name	Description	Device ID	Vendor ID	Vendor	Keyboard Shortcuts
10.0.2.4	Betelgeuse		70	15	Cornell University	
<input type="checkbox"/> Volttron Point Name <small>↑ + -</small>			Writable		Units	
<input type="checkbox"/>	ReturnAirTemperature	FALSE		UNKNOWN UNITS		...
<input type="checkbox"/>	ReturnAirHumidity	FALSE		UNKNOWN UNITS		...
<input type="checkbox"/>	CoolingValveOutputCommand	FALSE		UNKNOWN UNITS		...
<input type="checkbox"/>	MixedWaterTemperature	FALSE		UNKNOWN UNITS		...
<input type="checkbox"/>	OutdoorAirHumidity	FALSE		UNKNOWN UNITS		...
<input type="checkbox"/>	PreheatTemperature	FALSE		UNKNOWN UNITS		...
<input type="checkbox"/>	DischargeAirTemperature	FALSE		UNKNOWN UNITS		...
<input type="checkbox"/>	DischargeAirStaticPressure	FALSE		UNKNOWN UNITS		...

Console

Keyboard Shortcuts

Activate keyboard commands for device table that has focus.

Ctrl

ESC Deactivate keyboard commands.

Up Move keyboard selection up one row.

Down / Space Move keyboard selection down one row.

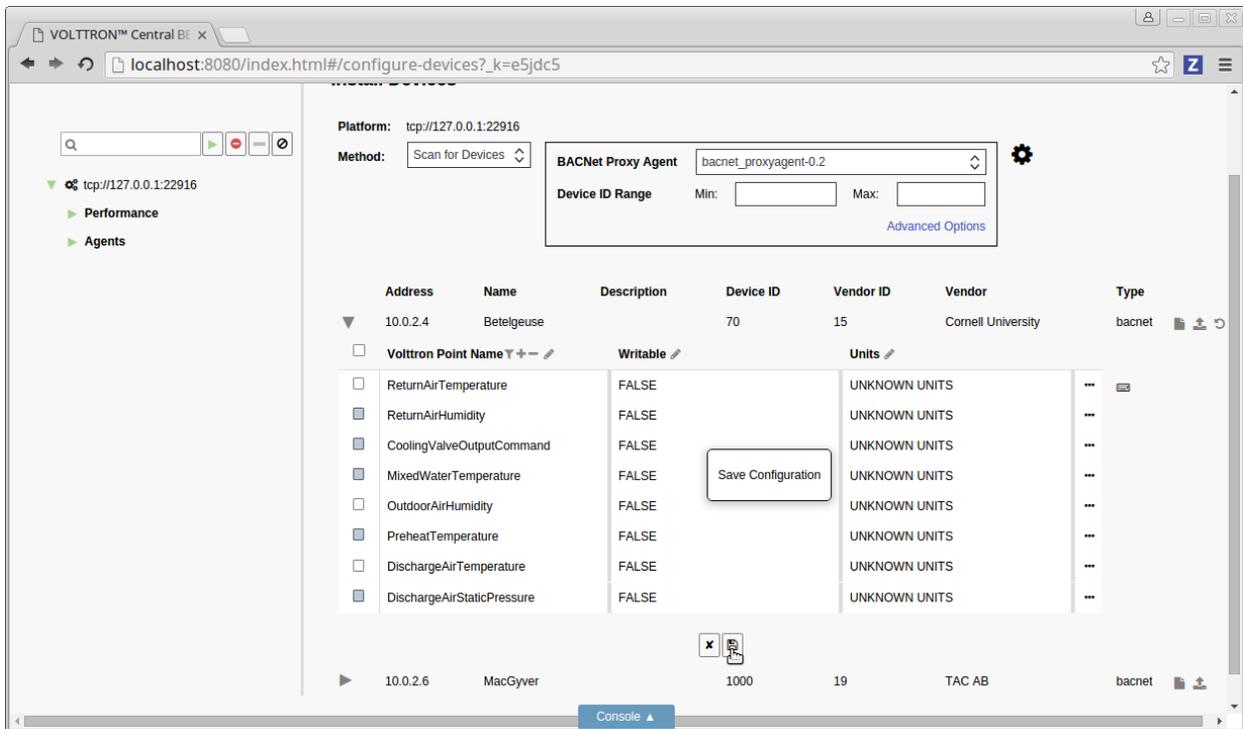
Shift+Up Extend keyboard selection up one row.

Shift+Down Extend keyboard selection down one row.

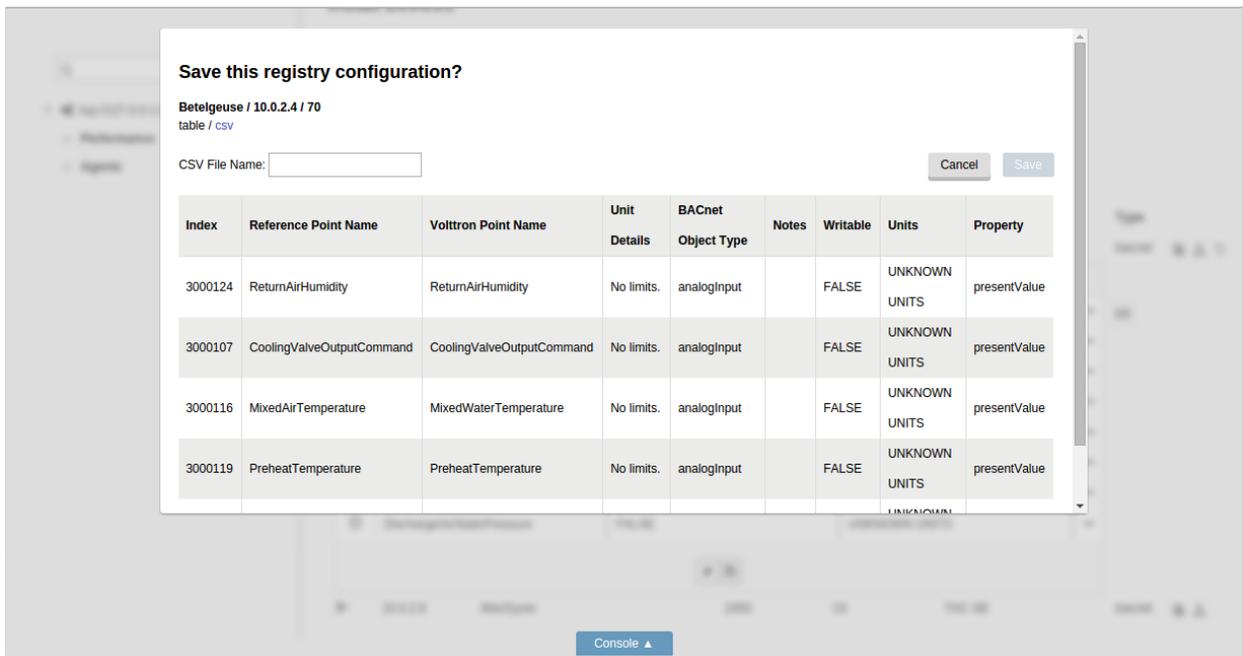
Enter Lock in keyboard selections.

Registry Preview

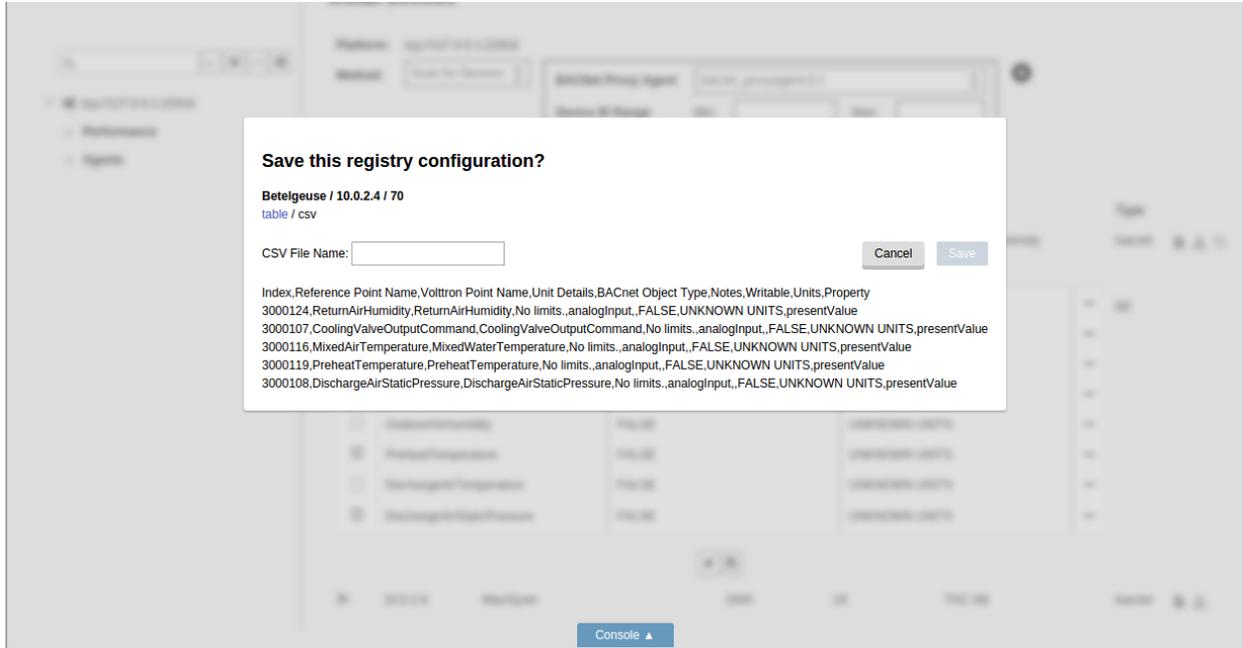
To save the registry configuration, click the Save button at the bottom of the registry configuration editor.



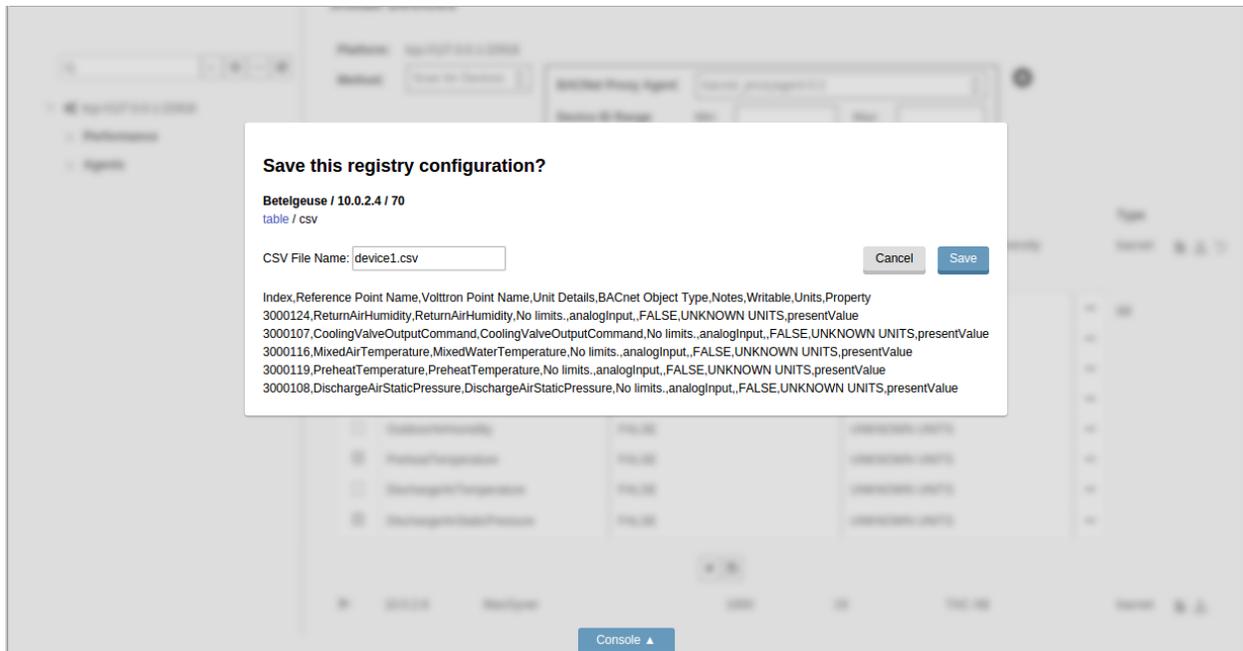
A preview will appear to let you confirm that the configuration is what you intended.

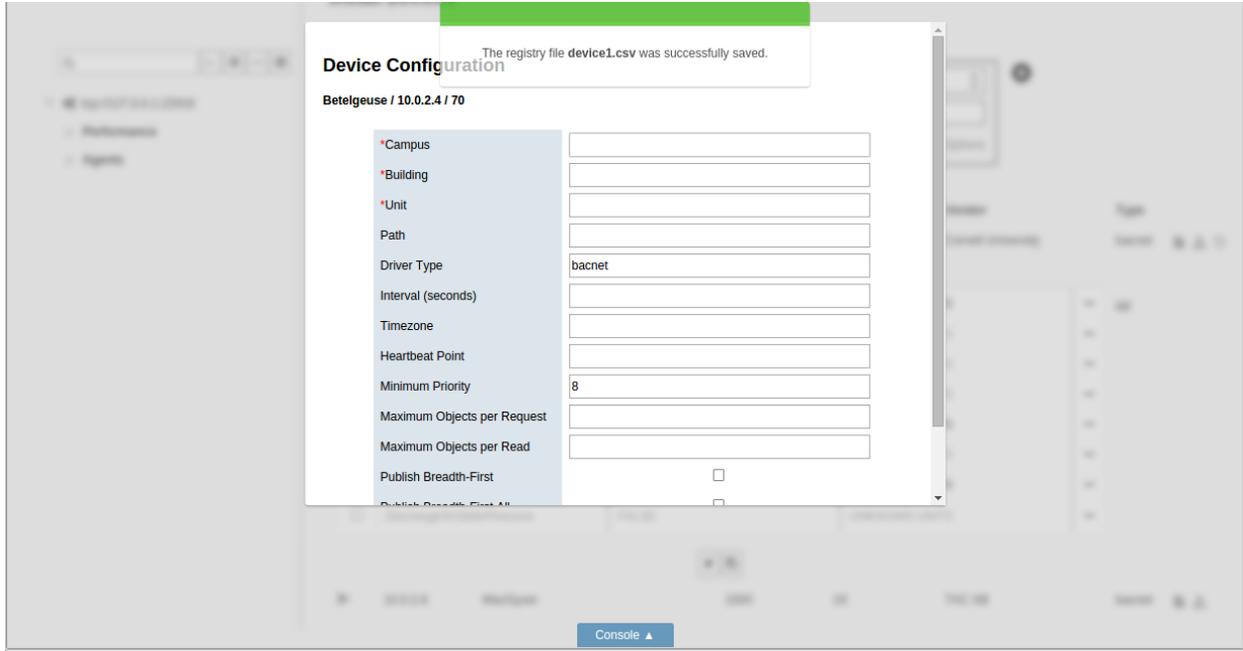


The configuration also can be inspected in the comma-separated format of the actual registry configuration file.



Provide a name for the registry configuration file, and click the Save button to save the file.

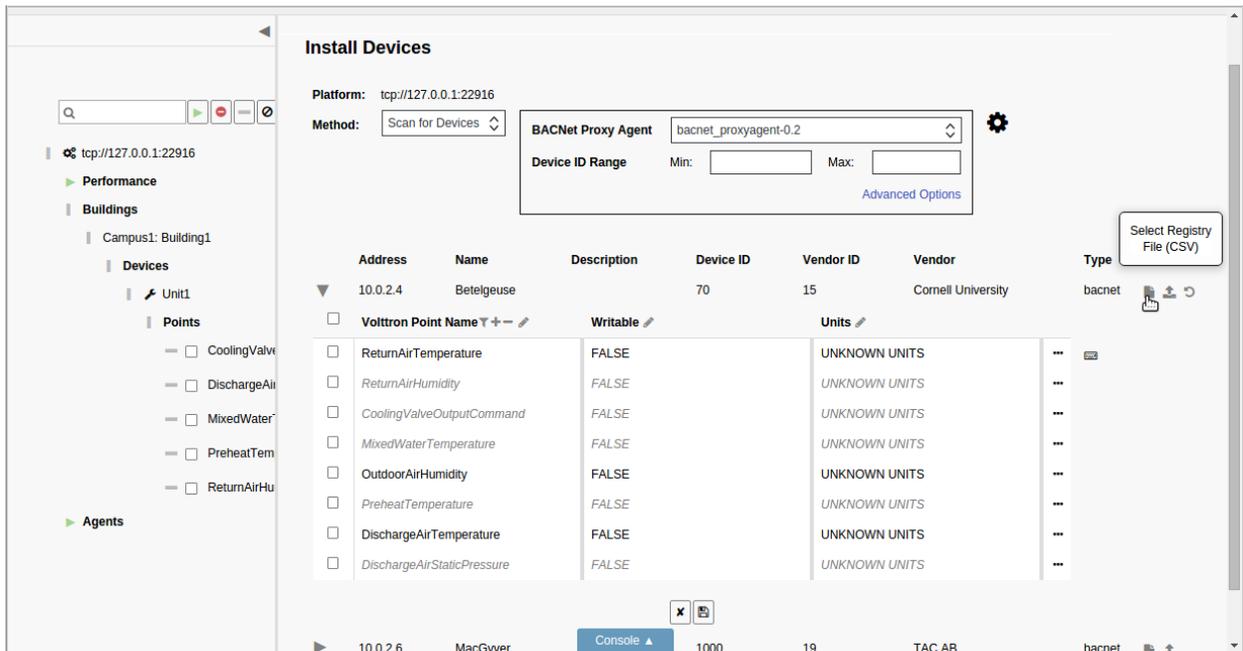




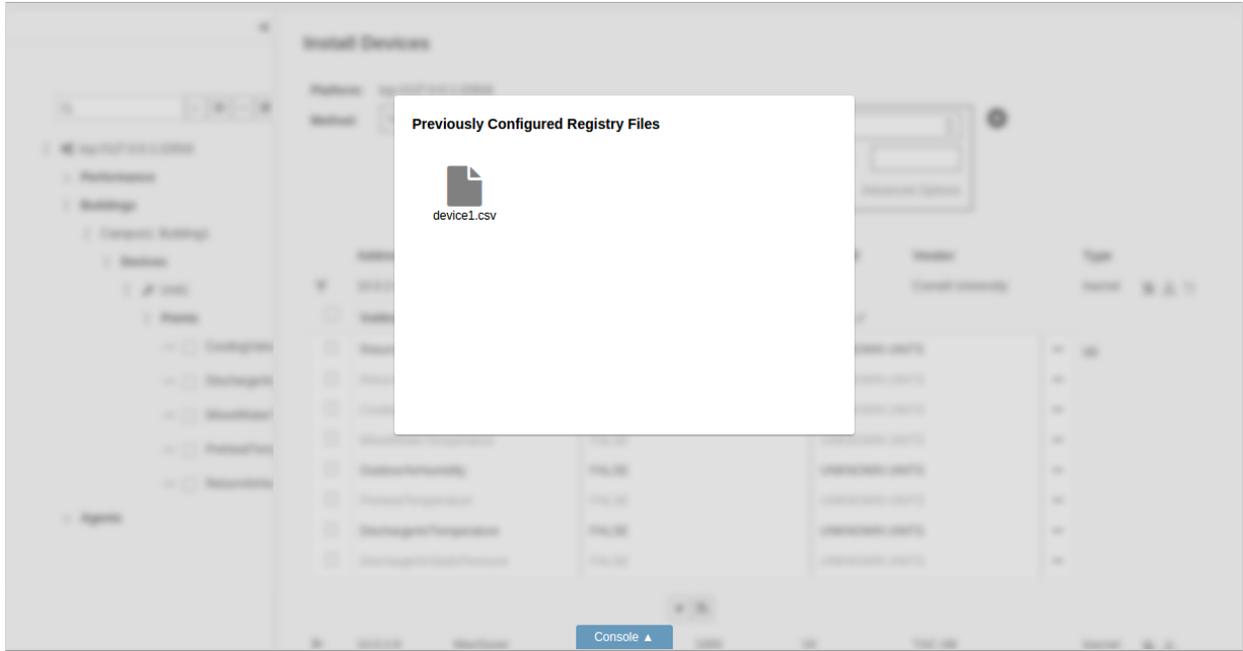
Registry Configuration Options

Different subsets of configured points can be saved from the same physical device and used to create separate registry files for multiple virtual devices and subdevices. Likewise, a single registry file can be reused by multiple virtual devices and subdevices.

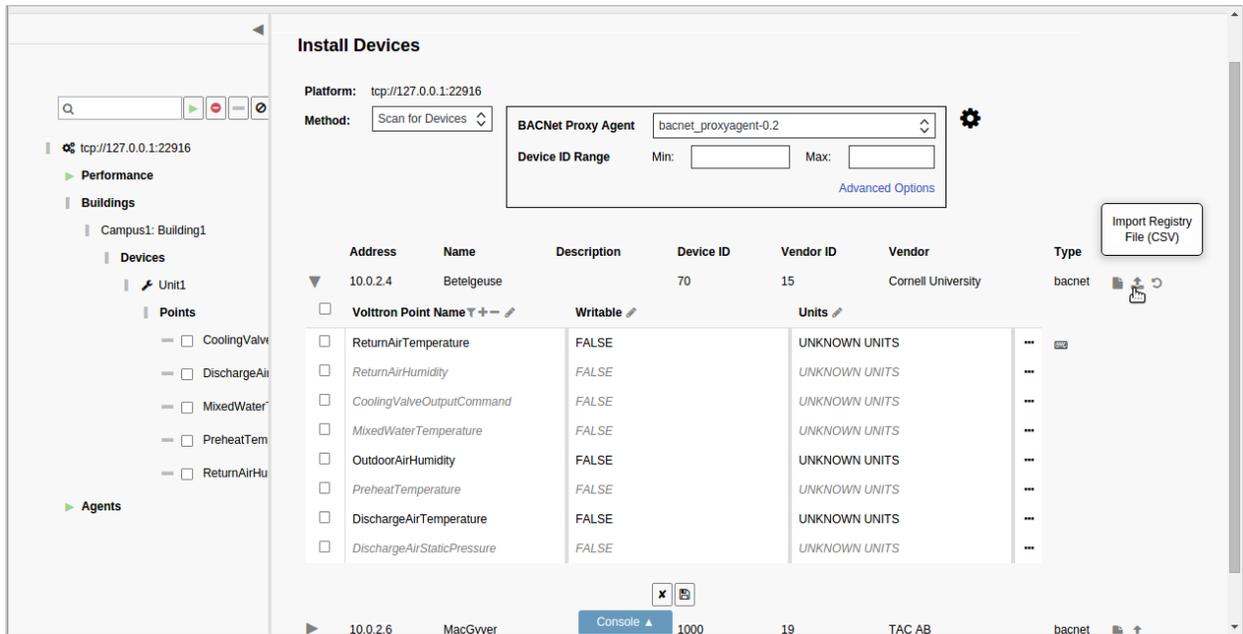
To reuse a previously saved registry file, click on the Select Registry File (CSV) button at the end of the physical device’s listing.



The Previously Configured Registry Files window will appear, and a file can be selected to load it into the registry configuration editor.

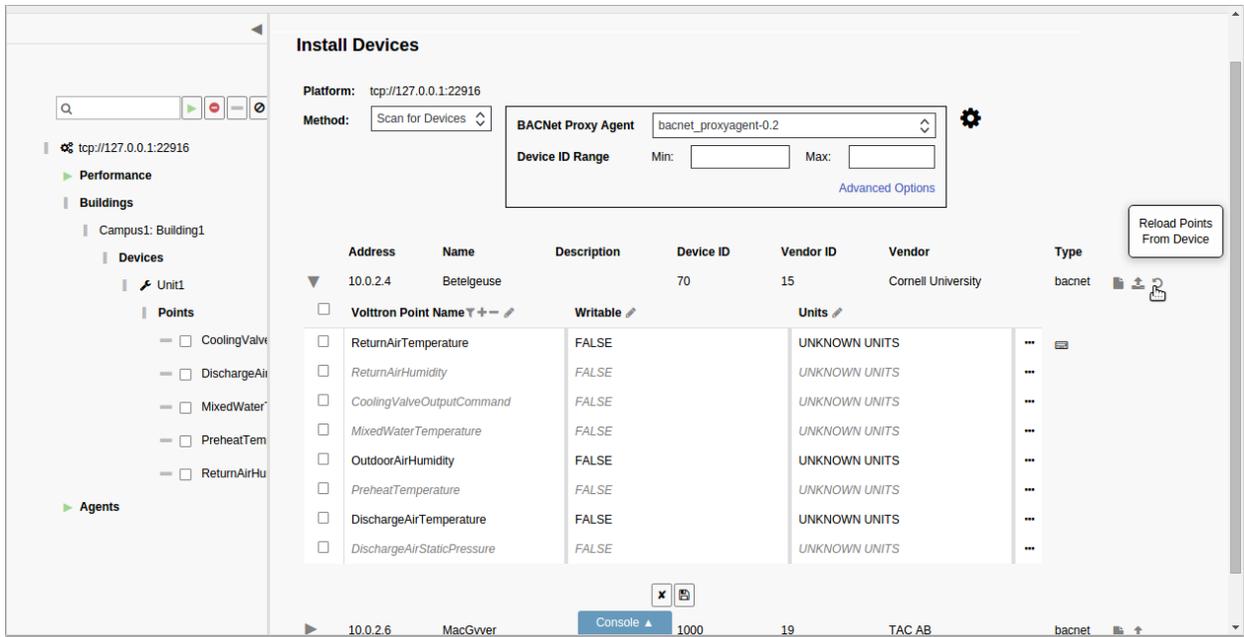


Another option is to import a registry configuration file from the computer running the VOLTRON Central web application, if one has been saved to local storage connected to the computer. To import a registry configuration file from local storage, click on the Import Registry File (CSV) button at the end of the physical device's listing, and use the file selector window to locate and load the file.



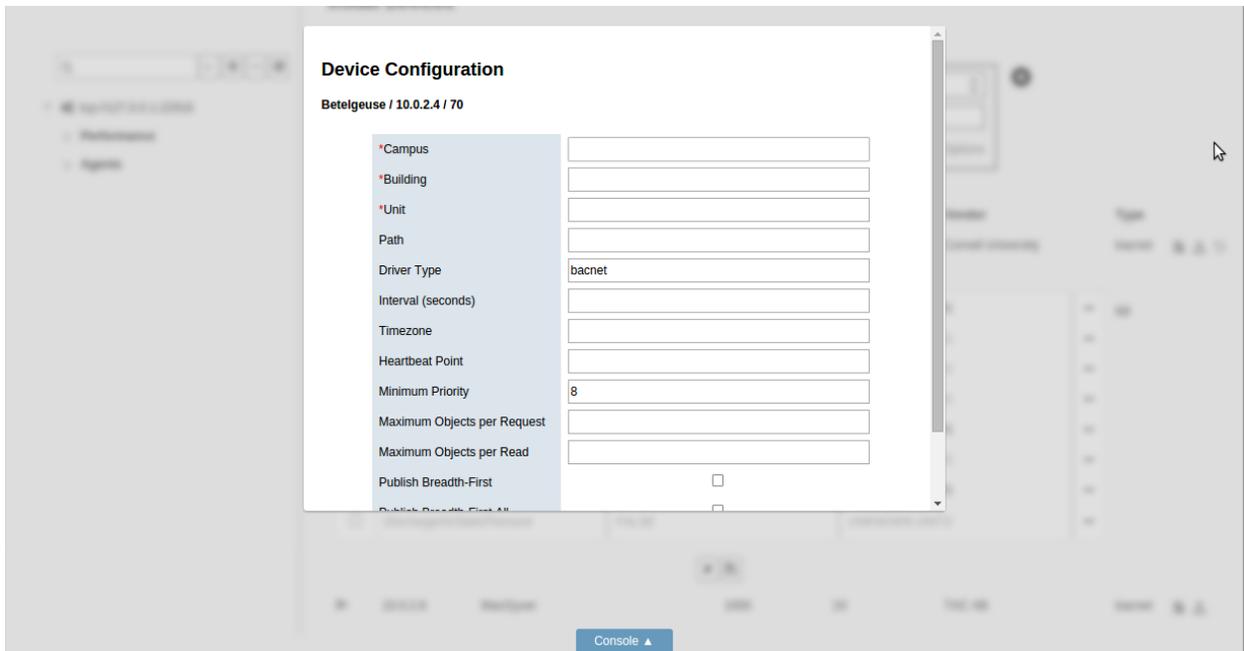
Reloading Device Points

Once a physical device has been scanned, the original points from the scan can be reloaded at any point during device configuration by clicking on the Reload Points From Device button at the end of the device's listing.

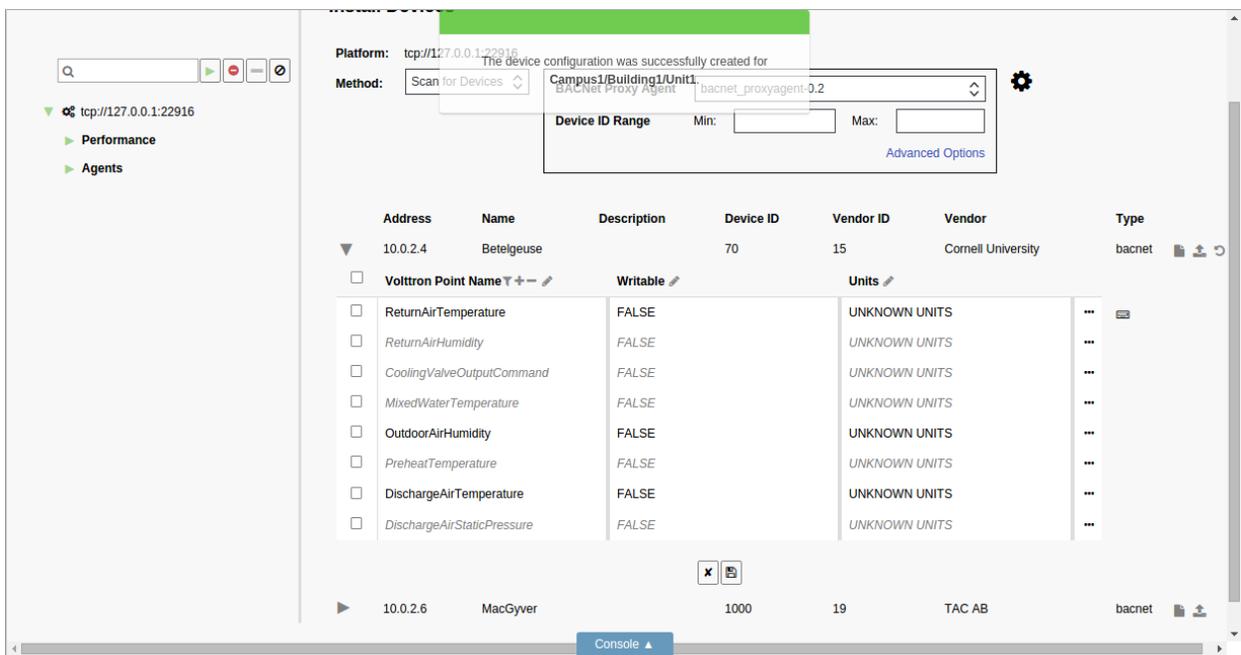
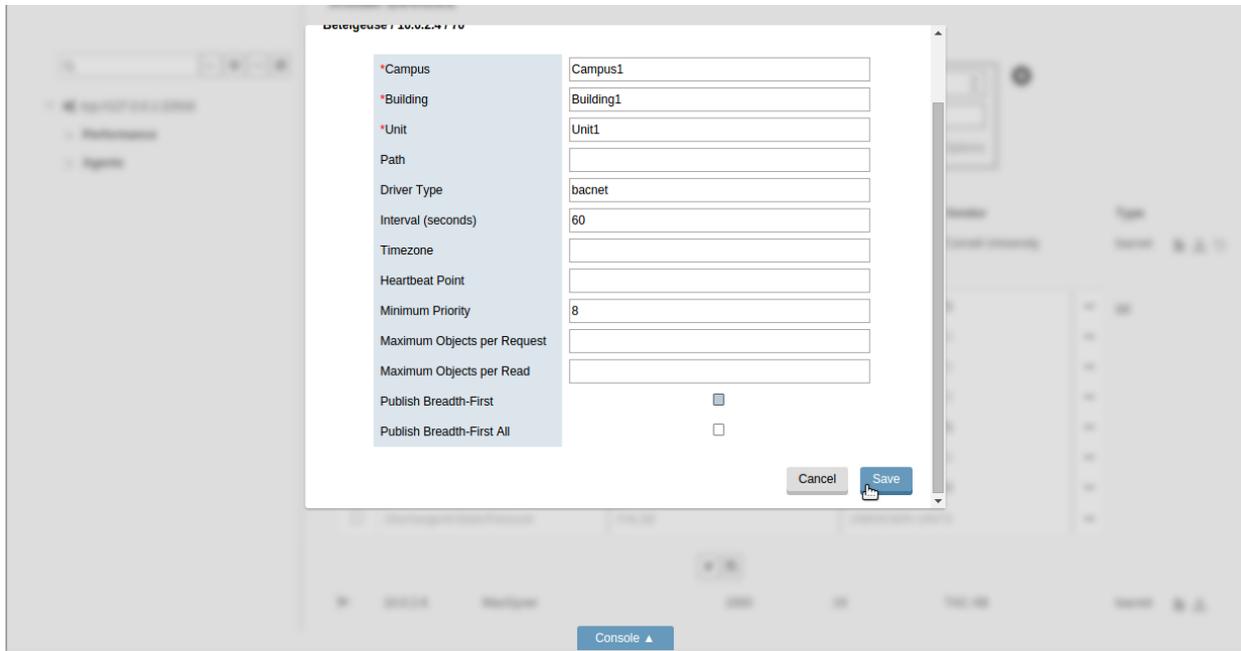


Device Configuration Form

After the registry configuration file has been saved, the device configuration form appears. Creating the device configuration results in the virtual device being installed in the platform and determines the device’s position in the side panel tree. It also contains some settings that determine how data is collected from the device.

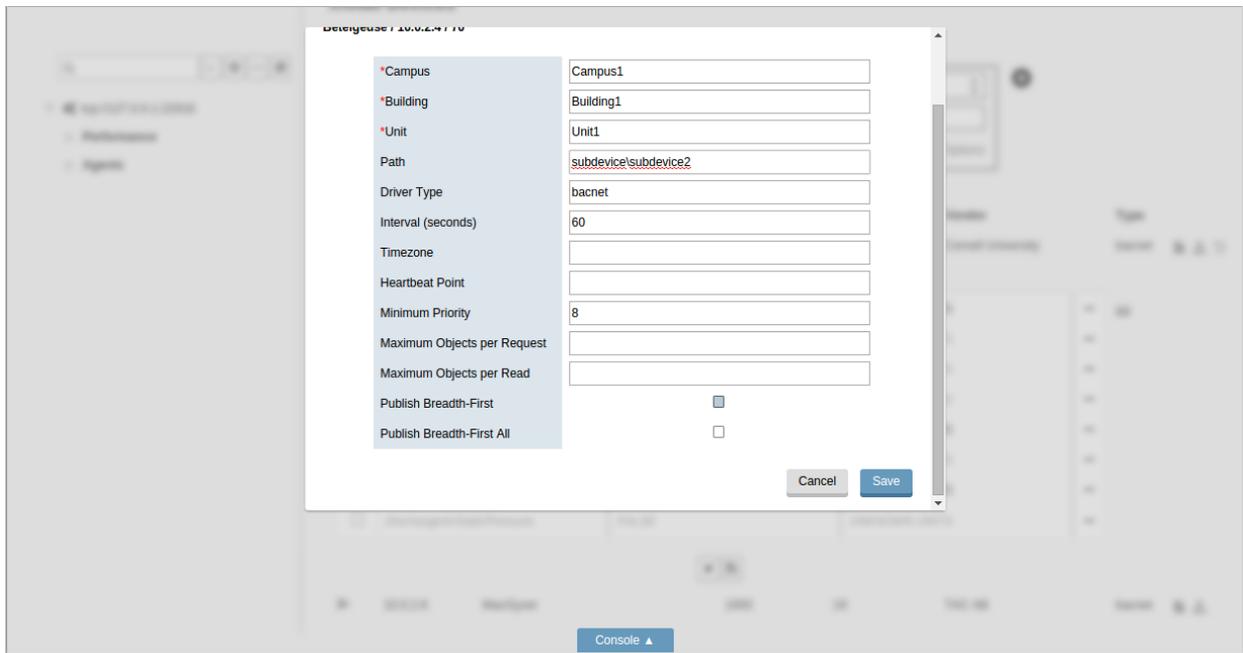
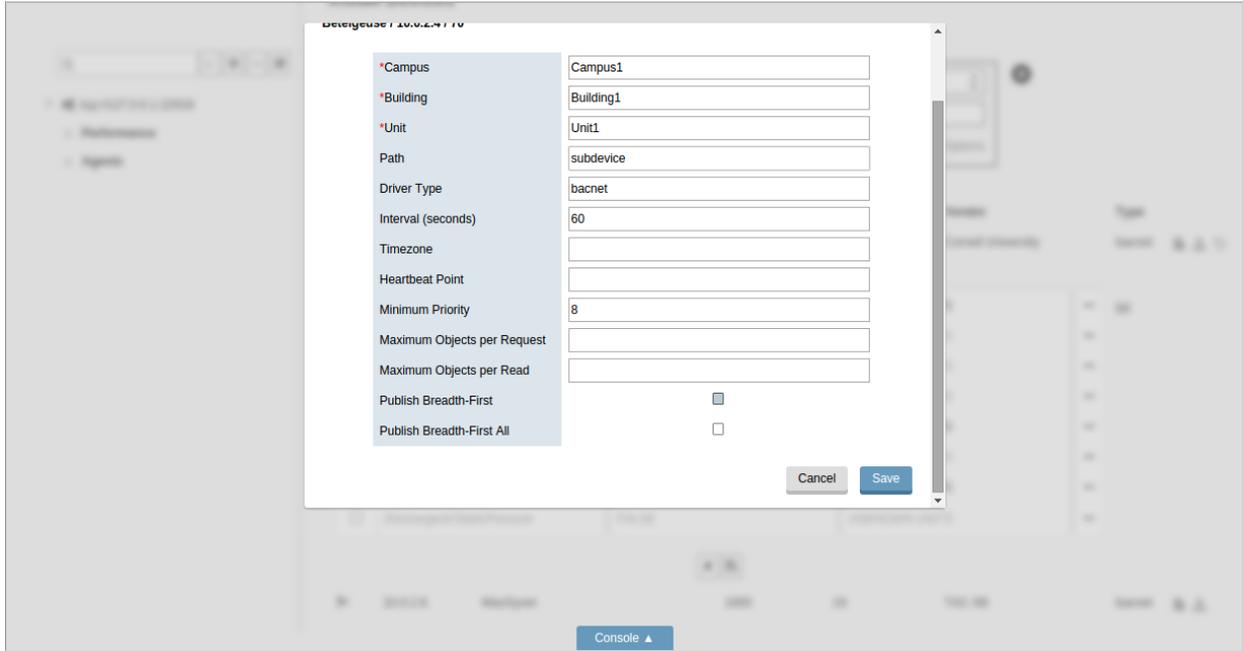


After the device configuration settings have been entered, click the Save button to save the configuration and add the device to the platform.

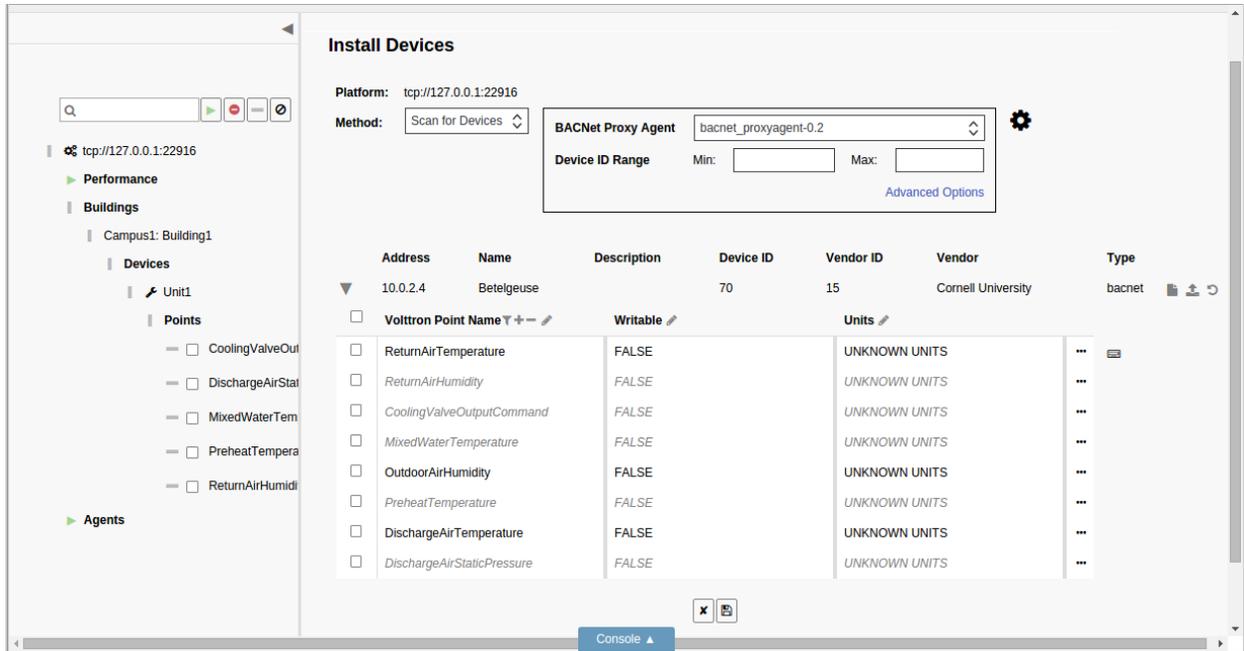


Configuring Subdevices

After a device has been configured, subdevices can be configured by pointing to their position in the Path attribute of the device configuration form. But a subdevice can't be configured until its parent device has been configured first.

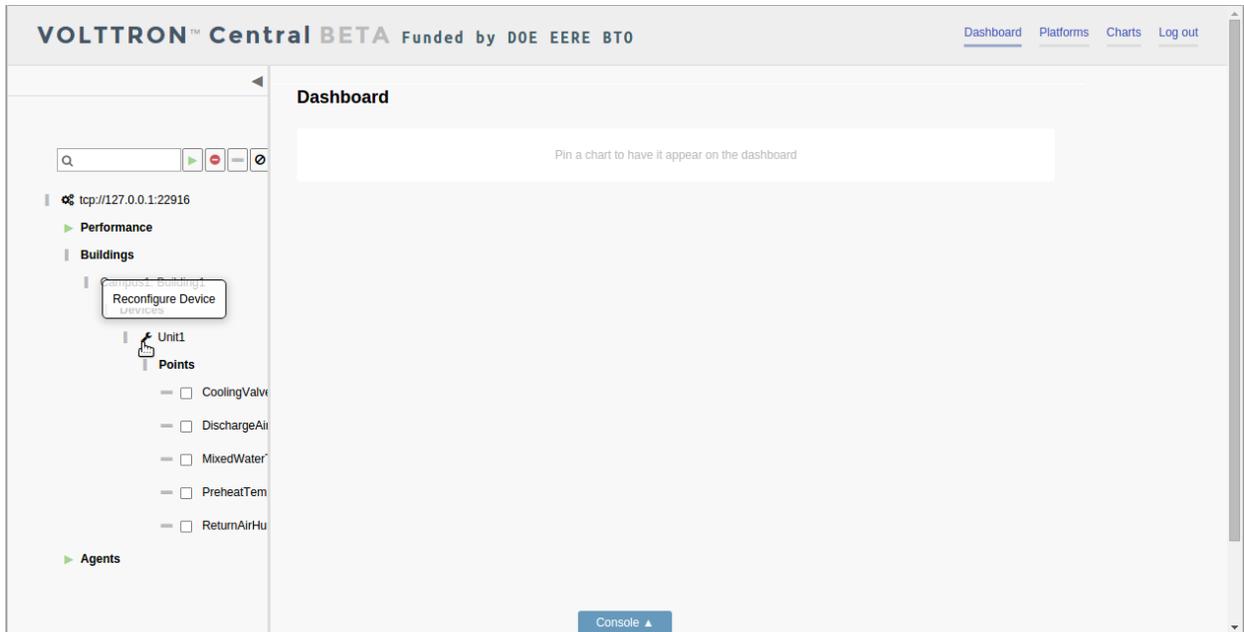


As devices are configured, they're inserted into position in the side panel tree, along with their configured points.

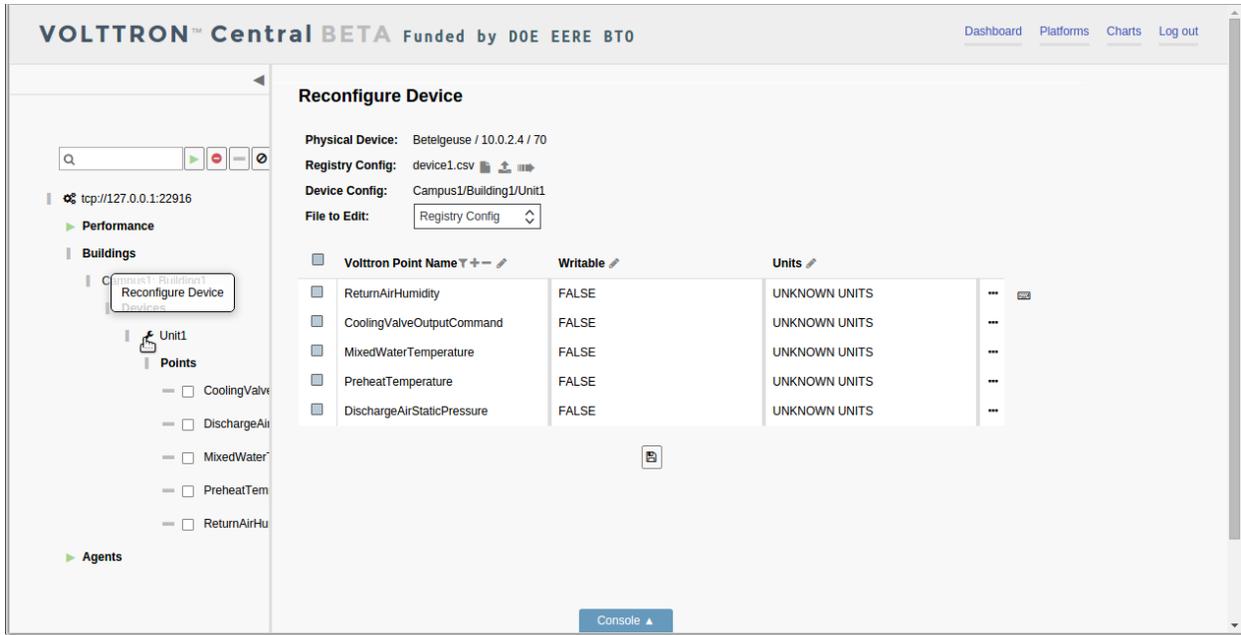


Reconfiguring Devices

A device that's been added to a VOLTRON instance can be reconfigured by changing its registry configuration or its device configuration. To launch reconfiguration, click on the wrench button next to the device in the side panel tree.



Reconfiguration reloads the registry configuration editor and the device configuration form for the virtual device. The editor and the form work the same way in reconfiguration as during initial device configuration.

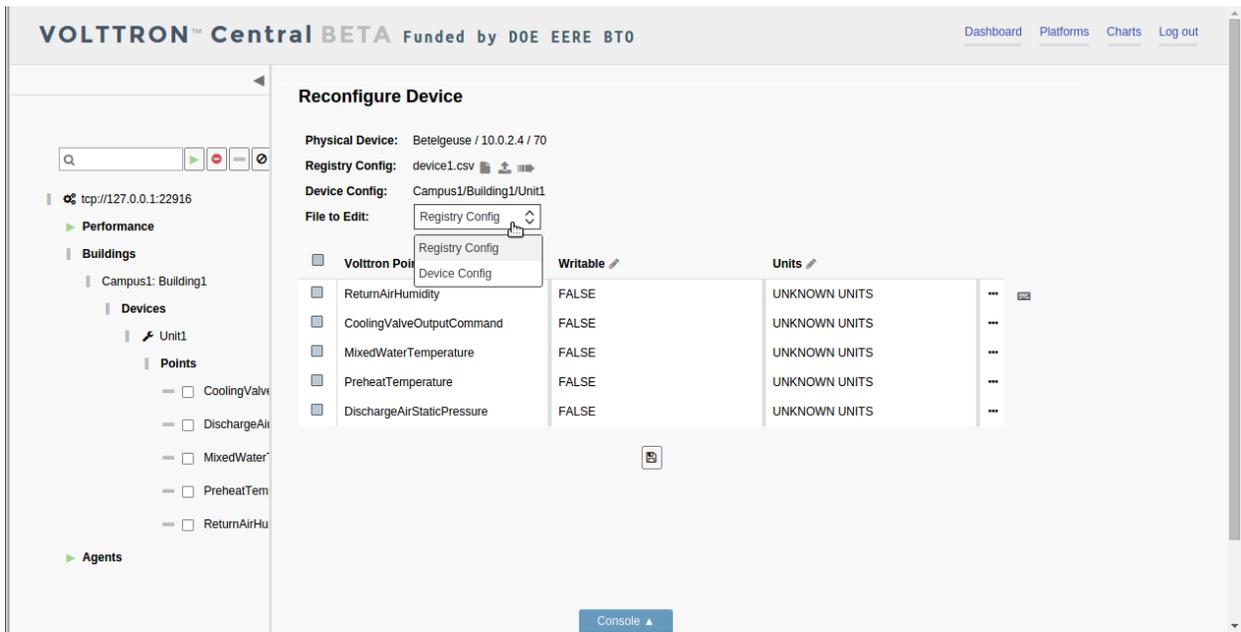


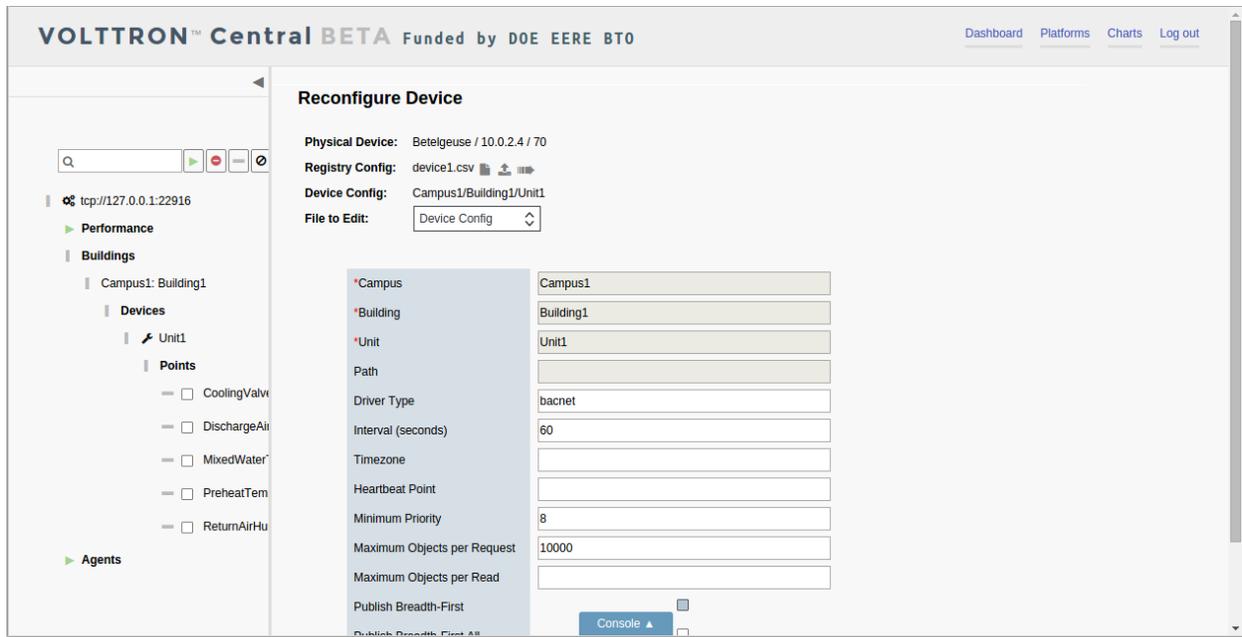
The reconfiguration view shows the name, address, and ID of the physical device that the virtual device was configured from. It also shows the name of the registry configuration file associated with the virtual device as well as its configured path.

A different registry configuration file can be associated with the device by clicking on the Select Registry File (CSV) button or the Import Registry File (CSV) button.

The registry configuration can be edited by making changes to the configuration in the editor and clicking the Save button.

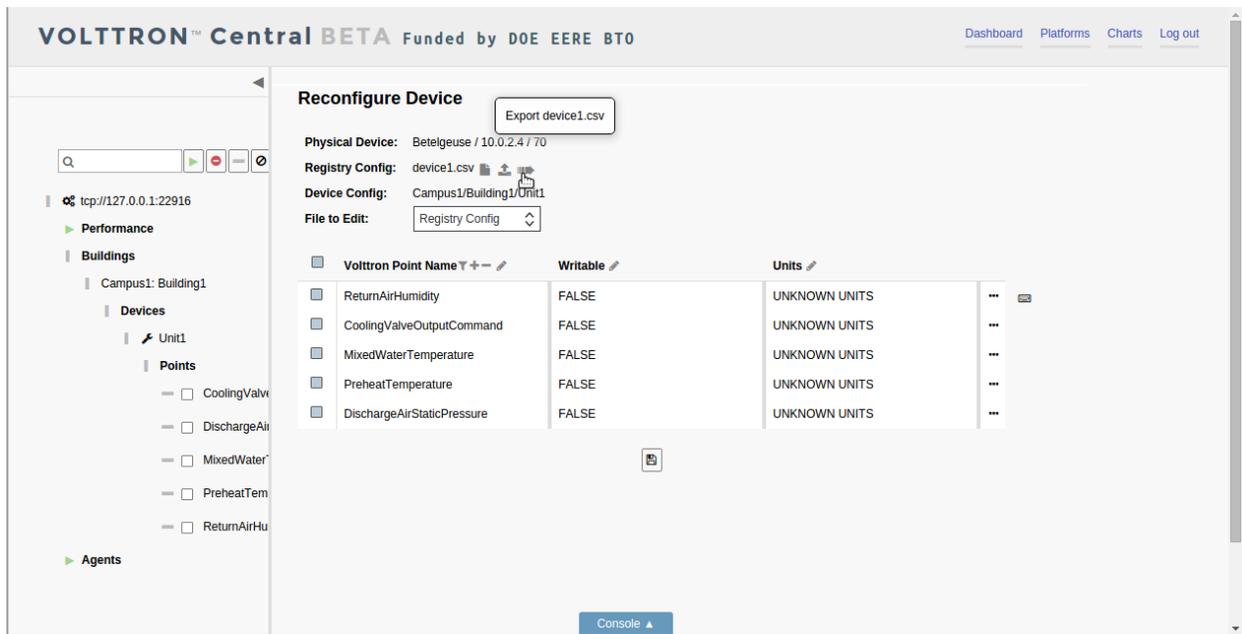
To make changes to the device configuration form, click on the File to Edit selector and choose Device Config.





Exporting Registry Configuration Files

The registry configuration file associated with a virtual device can be exported from the web browser to the computer's local storage by clicking on the File Export Button in the device reconfiguration view.



VOLTRON Central Demo

VOLTRON Central is a platform management web application that allows platforms to communicate and to be managed from a centralized server. This agent alleviates the need to ssh into independent nodes in order to manage

them. The demo will start up three different instances of VOLTTRON with three historians and different agents on each host. The following entries will help to navigate around the VOLTTRON Central interface.

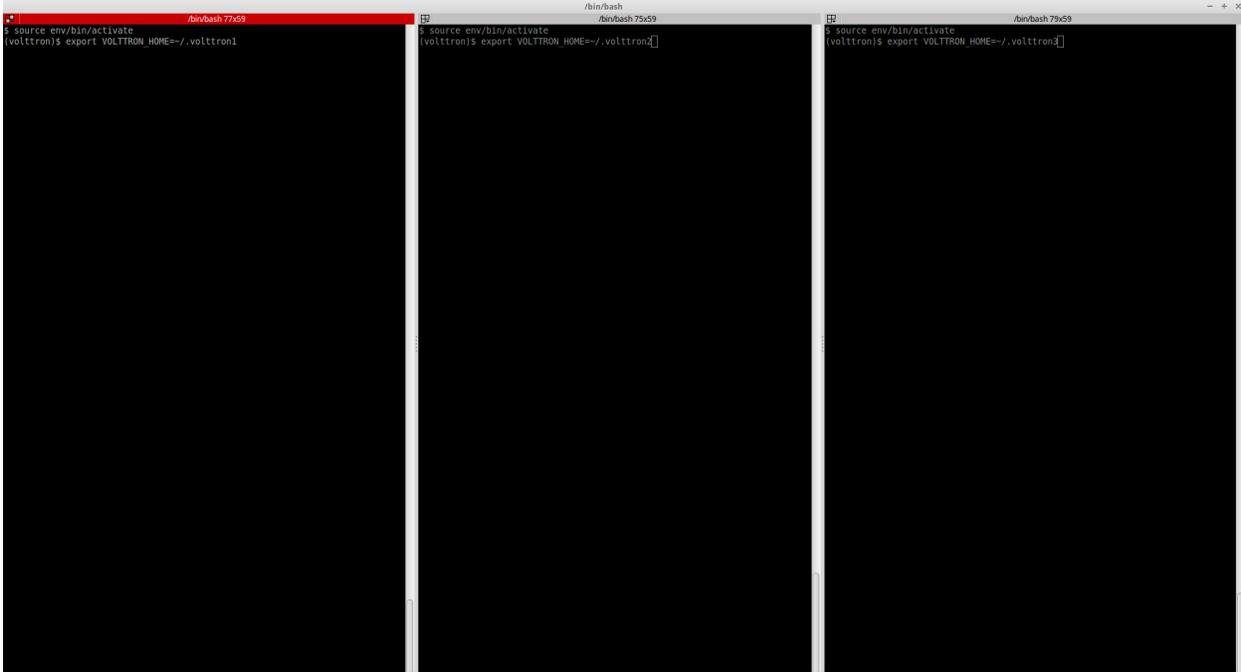
- *Getting Started*
- *Remote Platform Configuration*
- *Starting the Demo*
- *Stopping the Demo*
- *Log In*
- *Log Out*
- *Platforms Tree*
- *Loading the Tree*
- *Health Status*
- *Filter the Tree*
- *Platforms Screen*
- *Register New Platform*
- *Deregister Platform*
- *Platform View*
- *Add Charts*
- *Dashboard Charts*
- *Remove Charts*

Getting Started

After *building VOLTTRON*, open three shells with the current directory the root of the VOLTTRON repository. Then activate the VOLTTRON environment and export the VOLTTRON_HOME variable. The home variable needs to be different for each instance.

If you are using Terminator you can right click and select “Split Vertically”. This helps us keep from losing terminal windows or duplicating work.

```
$ source env/bin/activate
$ export VOLTTRON_HOME=~/.volttron1
```



One of our instances will have a VOLTRON Central agent. We will install a platform agent and a historian on all three platforms.

Run `vcfg` in the first shell. This command will ask how the instance should be set up. Many of the options have defaults that will be sufficient. When asked if this instance is a VOLTRON Central enter `y`. Read through the options and use the enter key to accept default options. There are no default credentials for VOLTRON Central. You can have it install the agents at this time. Below is an example configuration. In this case, username is `user` and localhost is `volttron-pc`.

```
(volttron)user@volttron-pc:~/voltron$ vcfg

Your VOLTRON_HOME currently set to: /home/user/.voltron1

Is this the volttron you are attempting to setup? [Y]:
What type of message bus (rmq/zmq)? [zmq]:
What is the vip address? [tcp://127.0.0.1]:
What is the port for the vip address? [22916]:
Is this instance web enabled? [N]: y
What is the protocol for this instance? [https]:
Web address set to: https://volttron-pc
What is the port for this instance? [8443]:
Would you like to generate a new web certificate? [Y]:
WARNING! CA certificate does not exist.
Create new root CA? [Y]:

Please enter the following details for web server certificate:
  Country: [US]:
  State: WA
  Location: Richland
  Organization: PNNL
  Organization Unit: VOLTRON
Created CA cert
Creating new web server certificate.
Is this an instance of volttron central? [N]: y
```

(continues on next page)

(continued from previous page)

```

Configuring /home/user/volttron/services/core/VolttronCentral.
Enter volttron central admin user name: <your volttron central admin_
↪username here>
Enter volttron central admin password: <your volttron central admin password_
↪here>
Retype password: <retype your volttron central admin password here>
Installing volttron central.
Should the agent autostart? [N]: y
Will this instance be controlled by volttron central? [Y]: y
Configuring /home/user/volttron/services/core/VolttronCentralPlatform.
What is the name of this instance? [volttron1]:
Volttron central address set to https://volttron-pc:8443
Should the agent autostart? [N]: y
Would you like to install a platform historian? [N]: y
Configuring /home/user/volttron/services/core/SQLHistorian.
Should the agent autostart? [N]: y
Would you like to install a master driver? [N]: y
Configuring /home/user/volttron/services/core/MasterDriverAgent.
Would you like to install a fake device on the master driver? [N]: y
Should the agent autostart? [N]: y
Would you like to install a listener agent? [N]: y
Configuring examples/ListenerAgent.
Should the agent autostart? [N]: y
Finished configuration!

You can now start the volttron instance.

If you need to change the instance configuration you can edit
the config file is at /home/user/.volttron1/config

(volttron)user@volttron-pc:~/volttron$

```

VOLTRON Central needs to accept the connecting instances' public keys. For this example we'll allow any CURVE credentials to be accepted. After *starting*, the command `vctl auth add` will prompt the user for information about how the credentials should be used. We can simply hit Enter to select defaults on all fields except **credentials**, where we will type `./.*`

```

$ vctl auth add --credentials "./.*/"
added entry domain=None, address=None, mechanism='CURVE', credentials=u'./.*/', user_
↪id='63b126a7-2941-4ebe-8588-711d1e6c70d1'

```

For more information on authorization see *authentication*.

Remote Platform Configuration

The next step is to configure the instances that will connect to VOLTRON Central. In the second and third terminal windows run `vcfg`. Like the `VOLTRON_HOME` variable, these instances need to have unique addresses.

Install a platform agent and a historian as before. Since we used the default options when configuring VOLTRON Central, we can use the default options when configuring these platform agents as well. The configuration will be a little different.

```

(volttron)user@volttron-pc:~/volttron$ vcfg

Your VOLTRON_HOME currently set to: /home/user/.volttron2

```

(continues on next page)

(continued from previous page)

```

Is this the volttron you are attempting to setup? [Y]:
What type of message bus (rmq/zmq)? [zmq]:
What is the vip address? [tcp://127.0.0.1]: tcp://127.0.0.2
What is the port for the vip address? [22916]:
Is this instance web enabled? [N]:
Is this an instance of volttron central? [N]:
Will this instance be controlled by volttron central? [Y]: y
Configuring /home/user/volttron/services/core/VolttronCentralPlatform.
What is the name of this instance? [volttron1]:
What is the hostname for volttron central? [https://volttron-pc]:
What is the port for volttron central? [8443]:
Should the agent autostart? [N]: y
Would you like to install a platform historian? [N]: y
Configuring /home/user/volttron/services/core/SQLHistorian.
Should the agent autostart? [N]: y
Would you like to install a master driver? [N]:
Would you like to install a listener agent? [N]:
Finished configuration!

You can now start the volttron instance.

If you need to change the instance configuration you can edit
the config file is at /home/user/.volttron2/config

(volttron)user@volttron-pc:~/volttron$

```

Starting the Demo

Start each Volttron instance after configuration. The “-l” option in the following command tells volttron to log to a file. The file name should be different for each instance.

```
$ volttron -l log1&
```

Note: If you choose to not start your agents with their platforms they will need to be started by hand.

List the installed agents with

```
$ vctl status
```

A portion of each agent’s uuid makes up the leftmost column of the status output. This is all that is needed to start or stop the agent. If any installed agents share a common prefix then more of the uuid will be needed to identify it.

```
$ vctl start uuid
```

or

```
$ vctl start --tag tag
```

Note: In each of the above examples one could use * suffix to match more than one agent.

Open your browser to `localhost:8443/vc/index.html` and log in with the credentials you provided. The platform agents should be automatically register with VOLTTRON central.

Note: localhost is the local host of your machine. In the above examples, this was `voltron-pc`.

Stopping the Demo

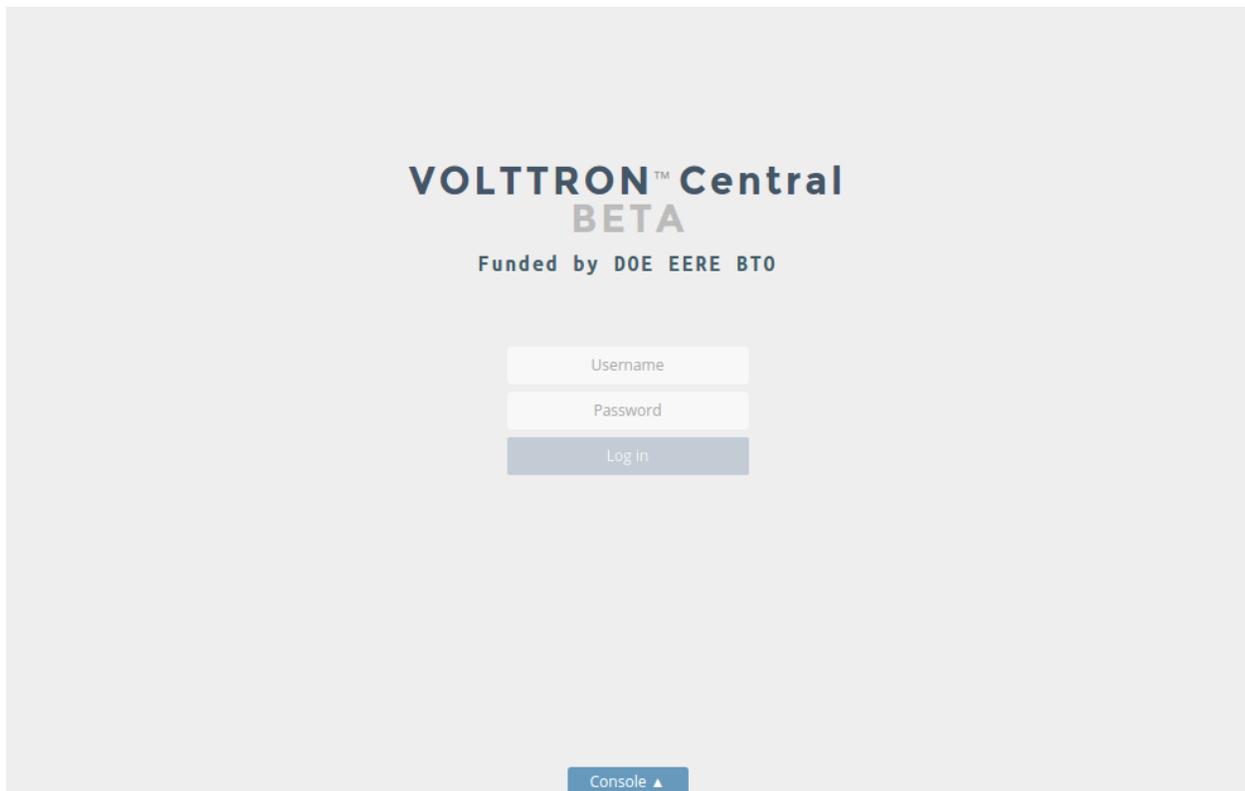
Once you have completed your walk through of the different elements of the VOLTTRON Central demo you can stop the demos by executing the following command in each terminal window.

```
$ vctl shutdown --platform
```

Once the demo is complete you may wish to see the *VOLTTRON Central Management Agent* page for more details on how to configure the agent for your specific use case.

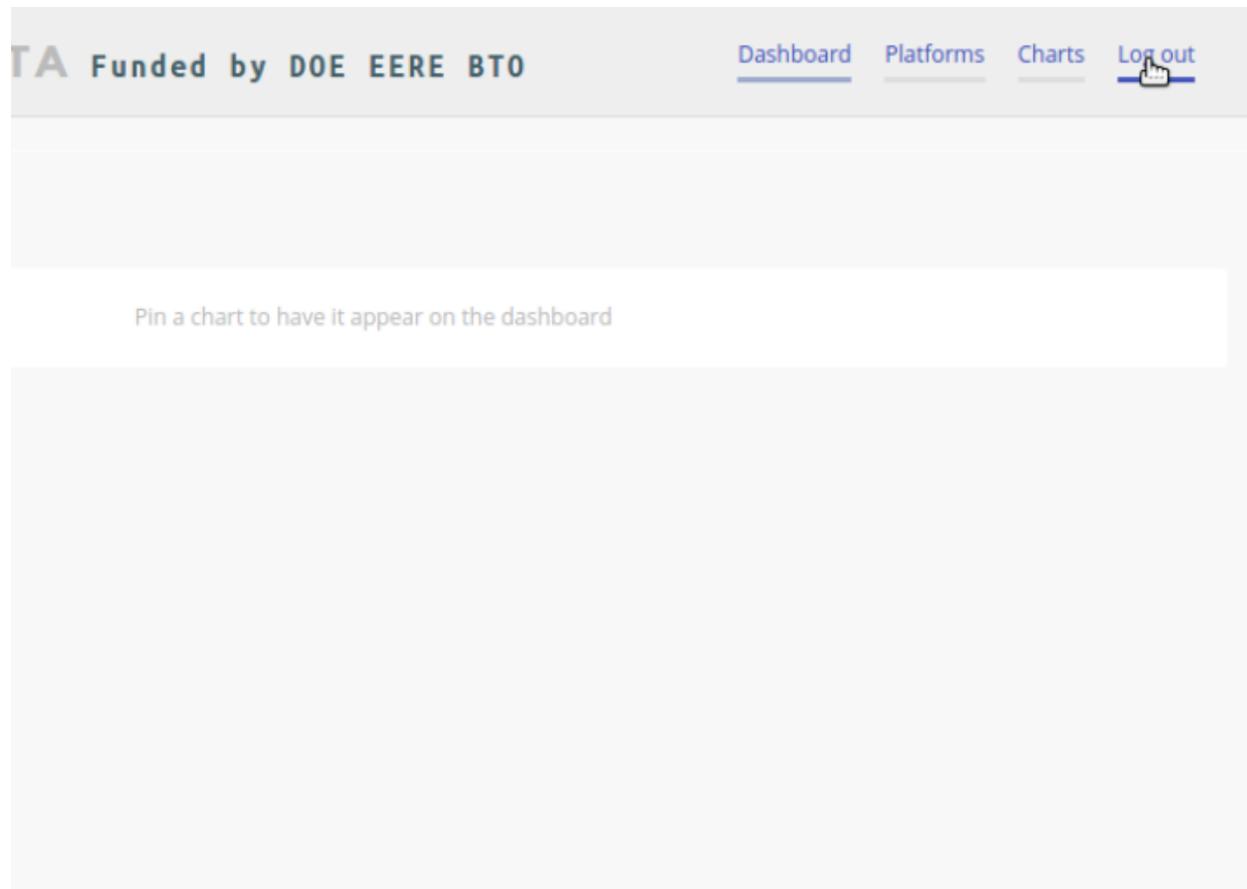
Log In

To log in to VOLTTRON Central, navigate in a browser to `localhost:8443/vc/index.html`, and enter the user name and password on the login screen.



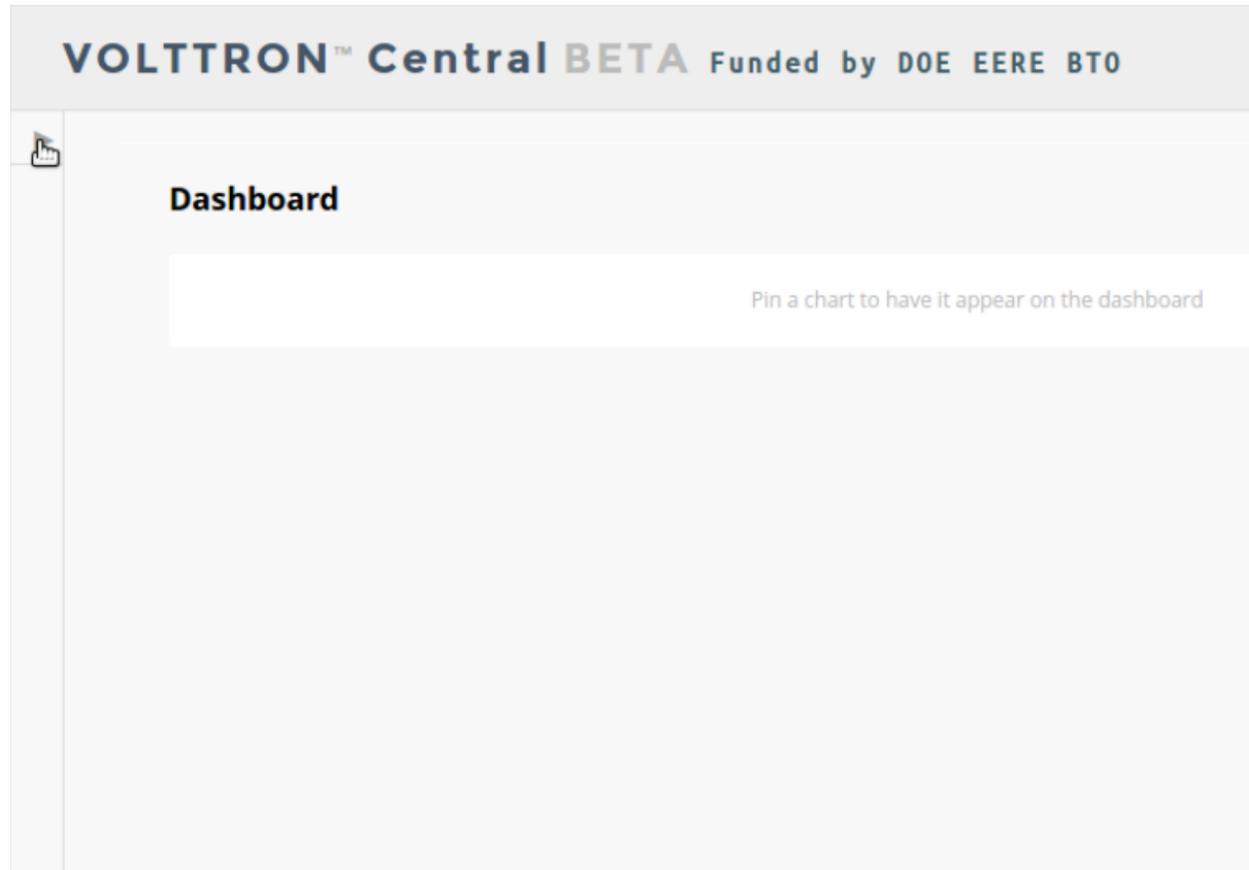
Log Out

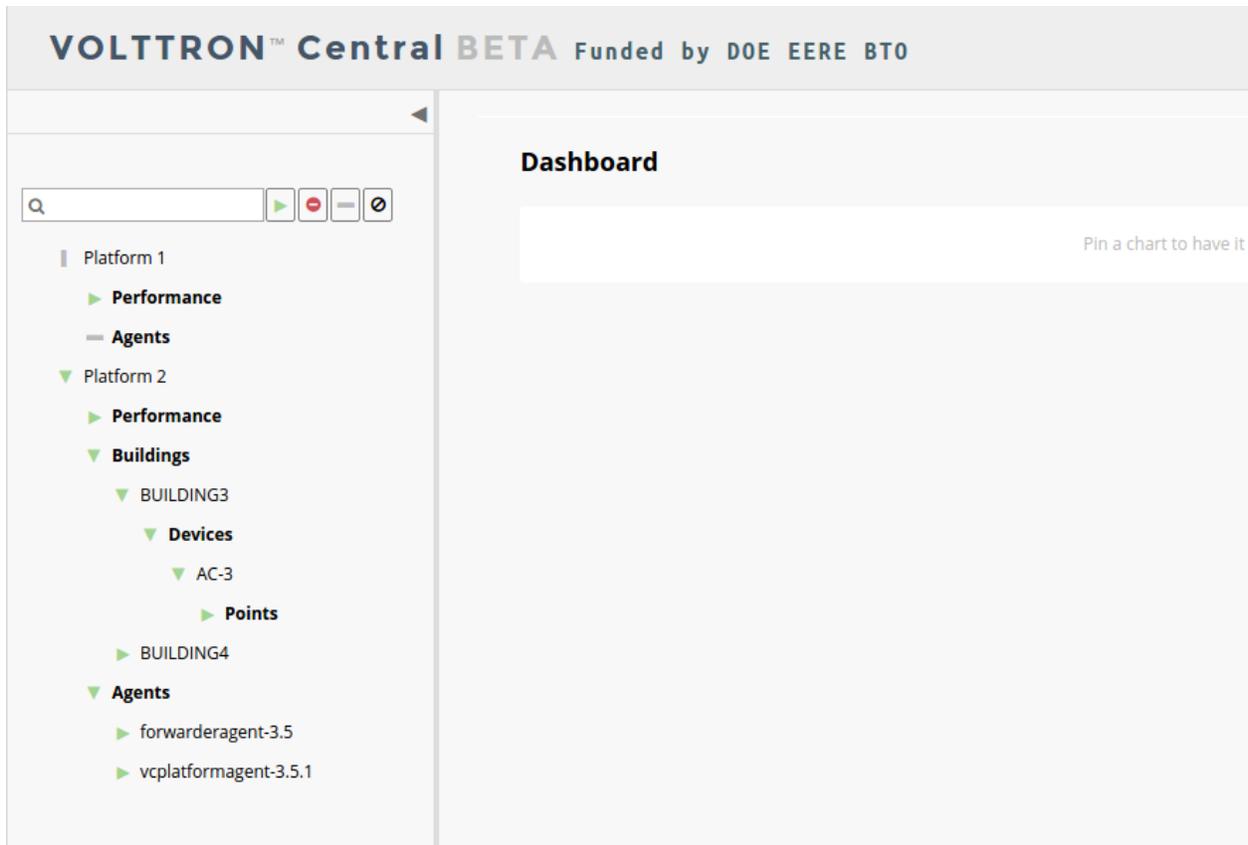
To log out of VOLTTRON Central, click the link at the top right of the screen.



Platforms Tree

The side panel on the left of the screen can be extended to reveal the tree view of registered platforms.

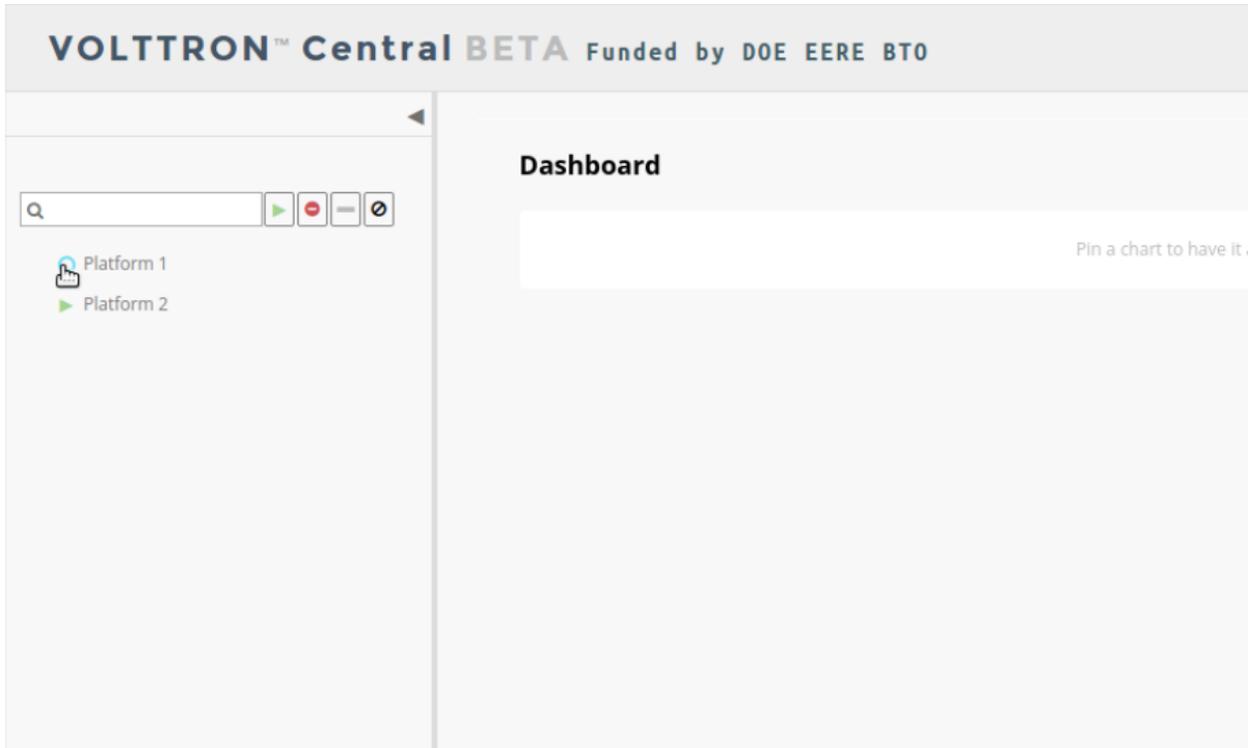




Top-level nodes in the tree are platforms. Platforms can be expanded in the tree to reveal installed agents, devices on buildings, and performance statistics about the platform instances.

Loading the Tree

The initial state of the tree is not loaded. The first time a top-level node is expanded is when the items for that platform are loaded.

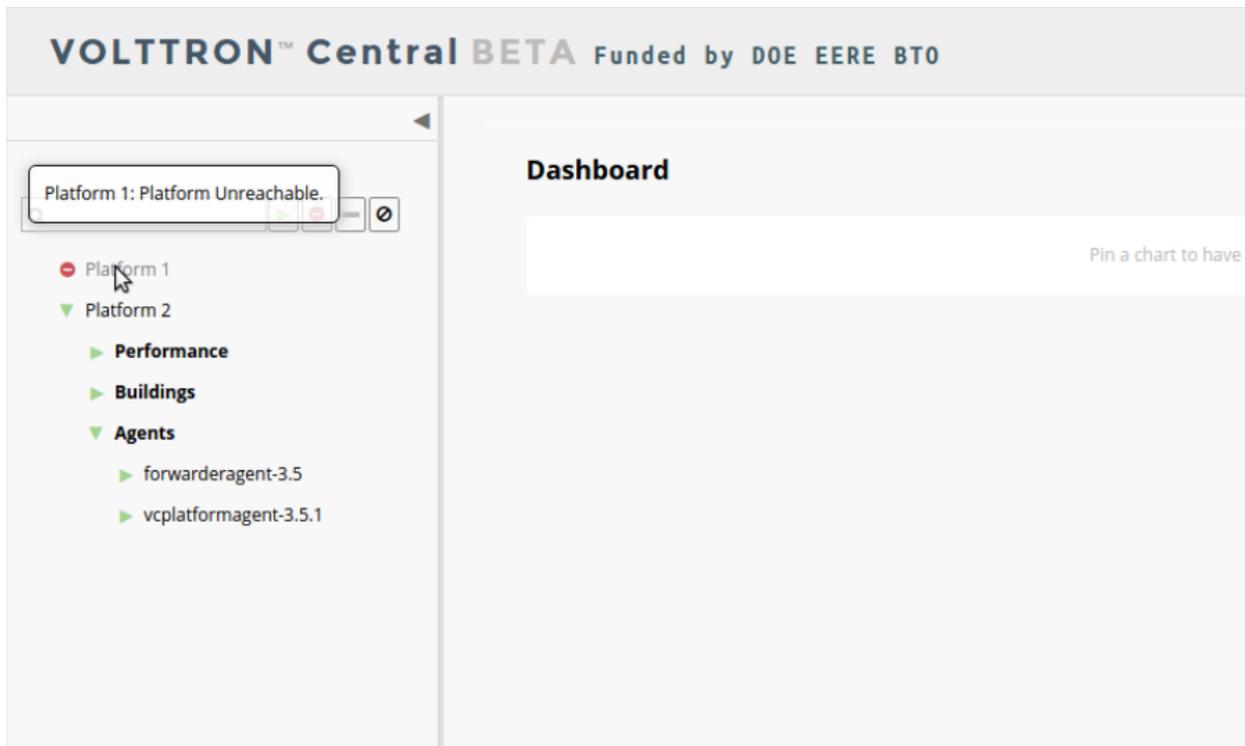


After a platform has been loaded in the tree, all the items under a node can be quickly expanded by double-clicking on the node.

Health Status

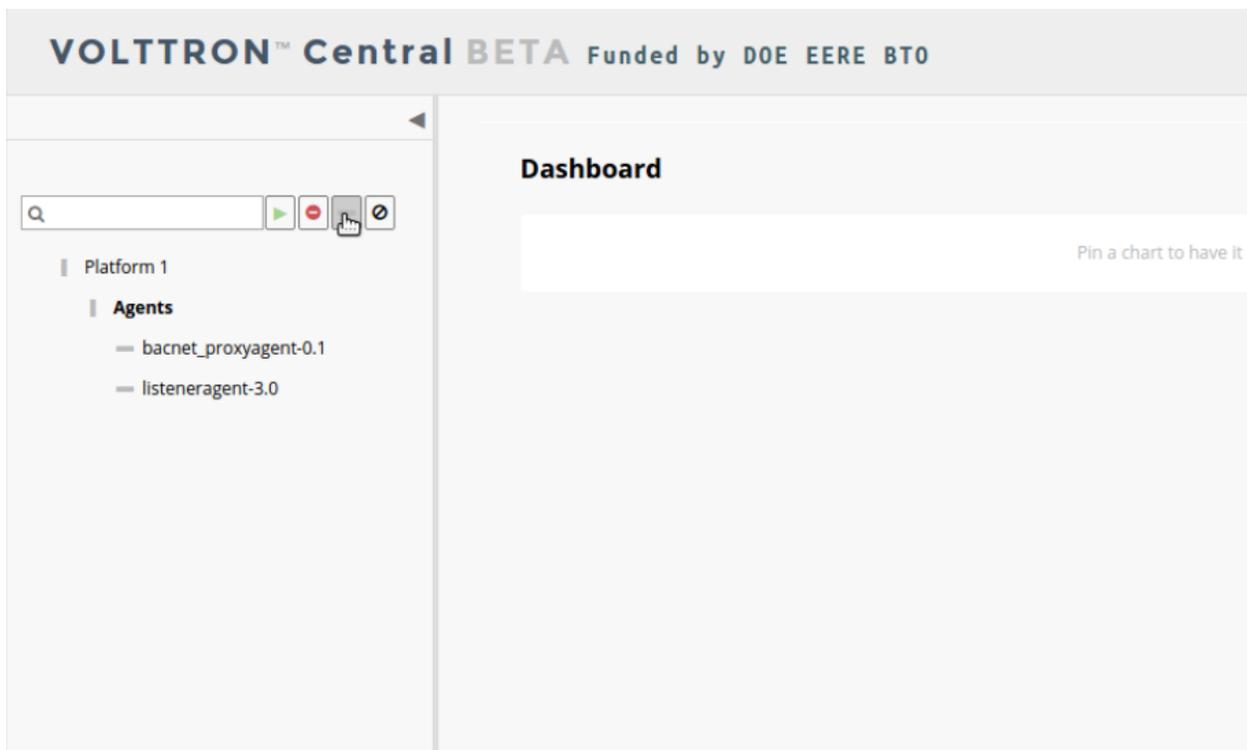
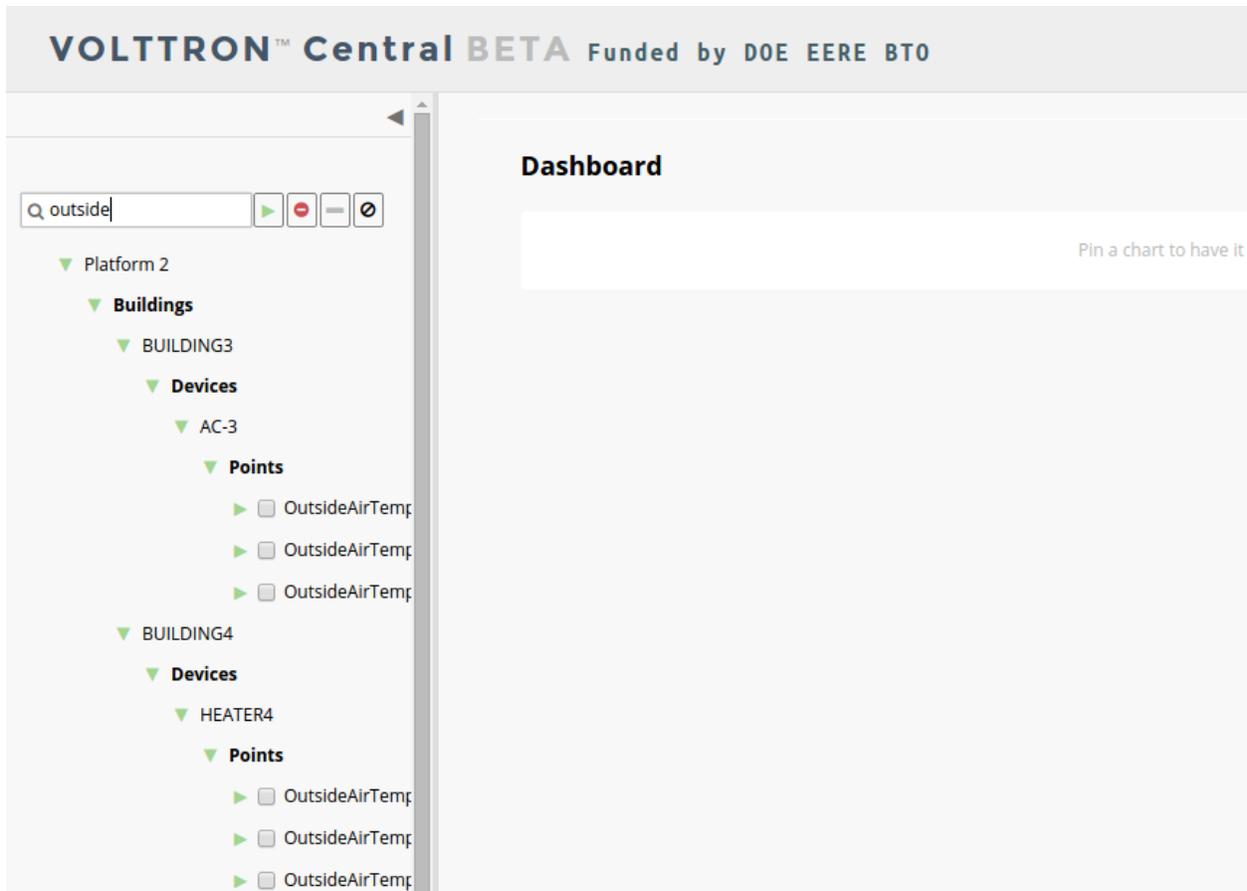
The health status of an item in the tree is indicated by the color and shape next to it. A green triangle means healthy, a red circle means there's a problem, and a gray rectangle means the status can't be determined.

Information about the health status also may be found by hovering the cursor over the item.



Filter the Tree

The tree can be filtered by typing in the search field at the top or clicking on a status button next to the search field.



Meta terms such as “status” can also be used as filter keys. Type the keyword “status” followed by a colon, and then

the word “good,” “bad,” or “unknown.”

The screenshot shows the VOLTRON Central BETA dashboard. The header includes the logo and text 'Funded by DOE EERE BTO'. The main content area is titled 'Dashboard' and contains a search bar with the text 'status: bad'. Below the search bar, there is a section for 'Platform 1' with a sub-section for 'Performance'. This section lists various performance metrics, each with a red minus sign and a checkbox:

- times_percent / irq
- cpu / percent
- times_percent / guest_nice
- times_percent / idle
- times_percent / iowait
- times_percent / guest
- times_percent / nice
- times_percent / softirq
- times_percent / steal
- times_percent / system
- times_percent / user

Platforms Screen

This screen lists the registered VOLTRON platforms and allows new platforms to be registered by clicking the Register Platform button. Each platform is listed with its unique ID and the number and status of its agents. The platform’s name is a link that can be clicked on to go to the platform management view.

The screenshot shows the VOLTRON Central BETA 'Platforms' screen. The header includes the logo and text 'Funded by DOE EERE BTO'. The main content area is titled 'Platforms' and contains a 'Register platform' button. Below the button, there is a list of registered platforms:

Platform Name	Unique ID	Agents Status
Platform 1	aae27d82-d66b-4c72-9899-be6166a44b5f	3 running, 0 stopped, 5 installed
Platform 2	7859e705-d604-4b5b-8f2b-4018780a28b8	2 running, 0 stopped, 2 installed

Platform View

From the platforms screen, click on the name link of a platform to manage it. Managing a platform includes installing, starting, stopping, and removing its agents.

VOLTRON™ Central BETA Funded by DOE EERE BT0 [Dashboard](#) [Platforms](#) [Charts](#) [Log out](#)

Platforms / Platform 1 (aae27d82-d66b-4c72-9899-be6166a44b5f)

Agents			
Name	UUID	Status	Action
bacnet_proxyagent-0.1	6a3be214-27a7-476d-a1d3-f3aa6c66db6e	Never started	Start Remove
listeneragent-3.0	0385b12e-2a5a-4d8a-abbf-6c3b6bab1acb	Never started	Start Remove
sqlhistorianagent-3.5.0	1fcb2c10-cb8e-4715-b248-69b4ad5d62b2	Running (PID 14296)	Stop Remove
vcplatformagent-3.5.1	dc720b2a-403a-4961-b316-74a964e5e038	Running (PID 14263)	Stop Remove
volttroncentralagent-3.5.3	f321f0b1-247a-476a-933b-d03cfdc81500	Running (PID 12933)	Stop Remove

Install agents
 No file chosen

To install a new agent, all you need is the agent's wheel file. Click on the button and choose the file to upload it and install the agent.

To start, stop, or remove an agent, click on the button next to the agent in the list. Buttons may be disabled if the user lacks the correct permission to perform the action or if the action can't be performed on a specific type of agent. For instance, platform agents and VOLTRON Central agents can't be removed or stopped, but they can be restarted if they've been interrupted.

Add Charts

Performance statistics and device points can be added to charts either from the Charts page or from the platforms tree in the side panel.

Click the Charts link at the top-right corner of the screen to go to the Charts page.

Dashboard

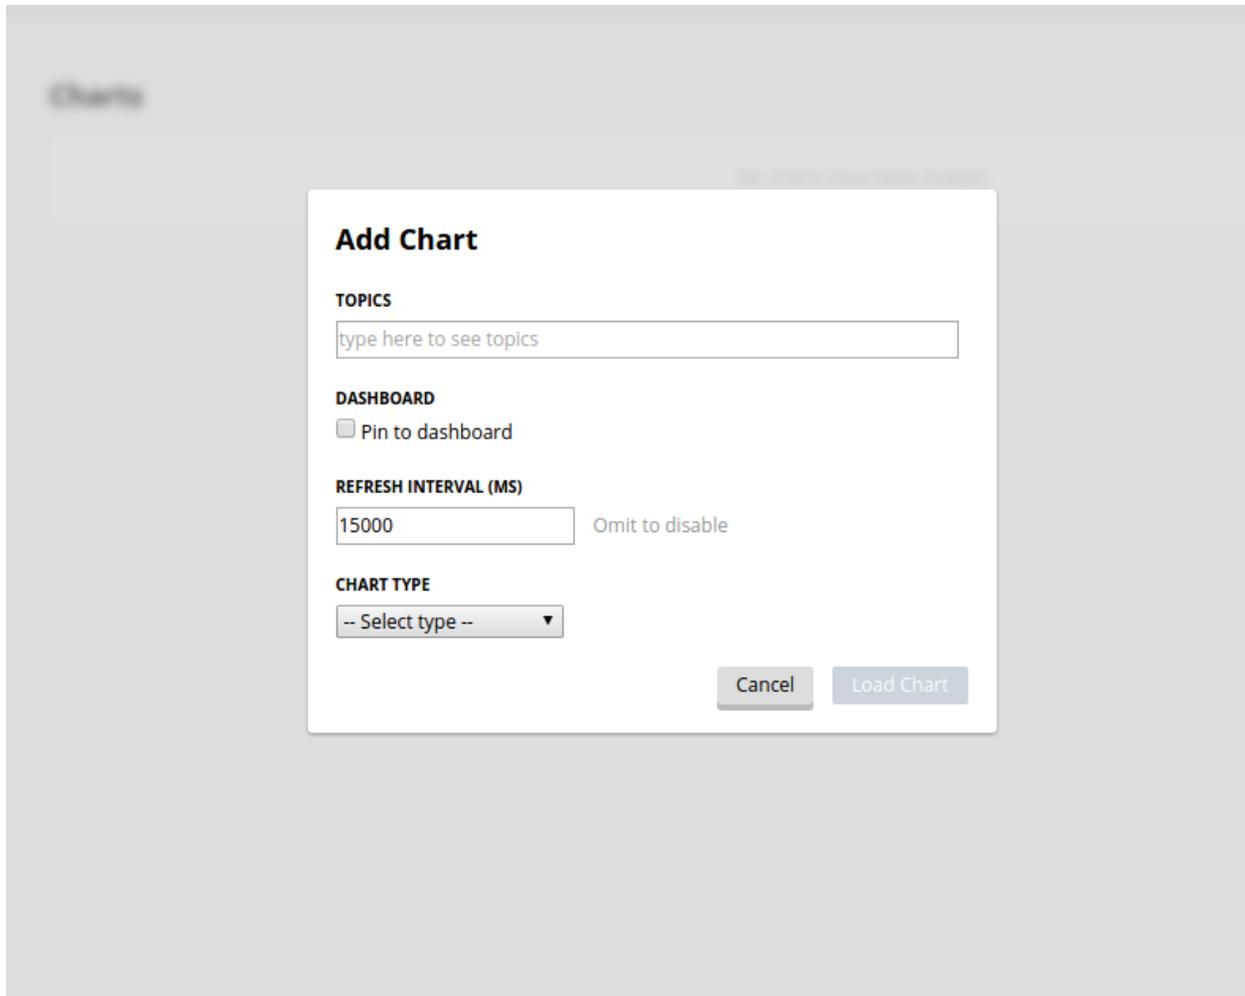
Pin a chart to have it appear on the dashboard

From the Charts page, click the Add Chart button to open the Add Chart window.

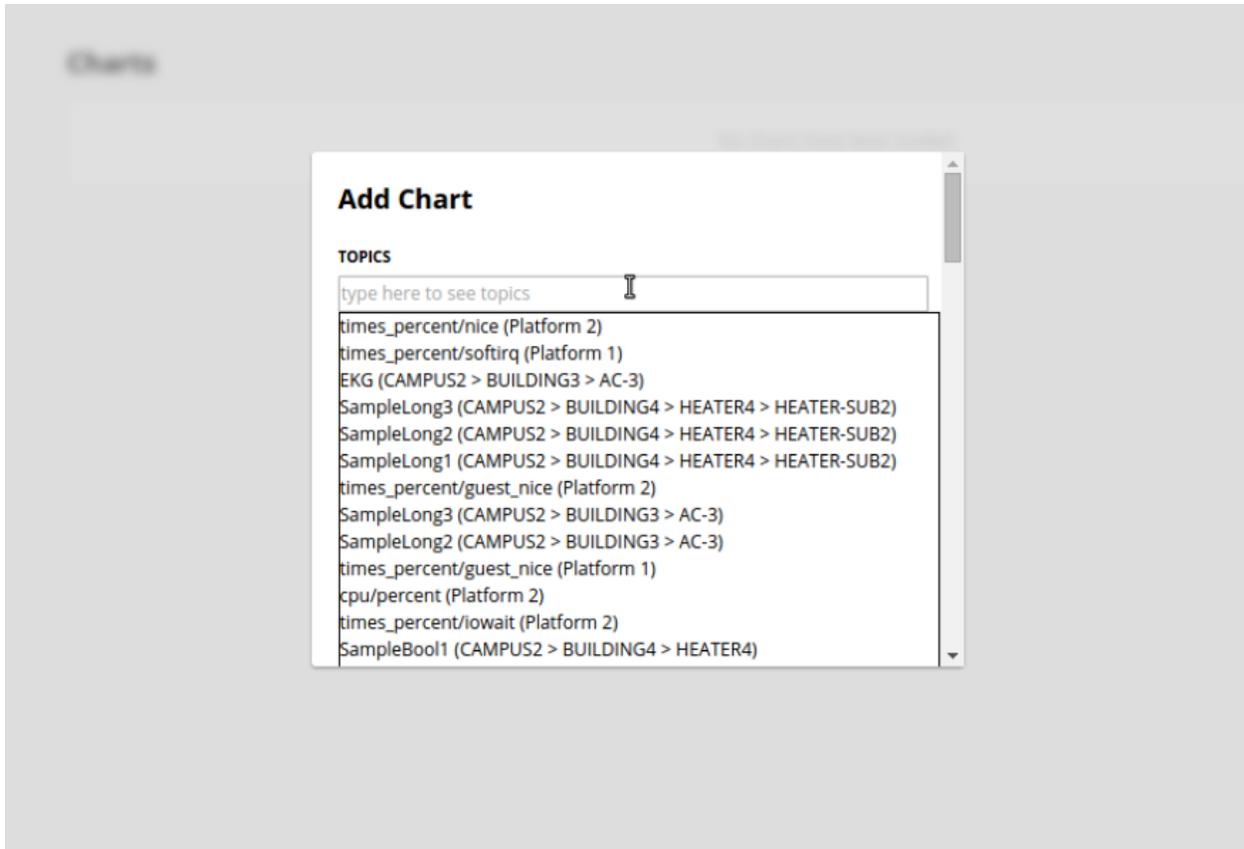
Charts

Add Chart

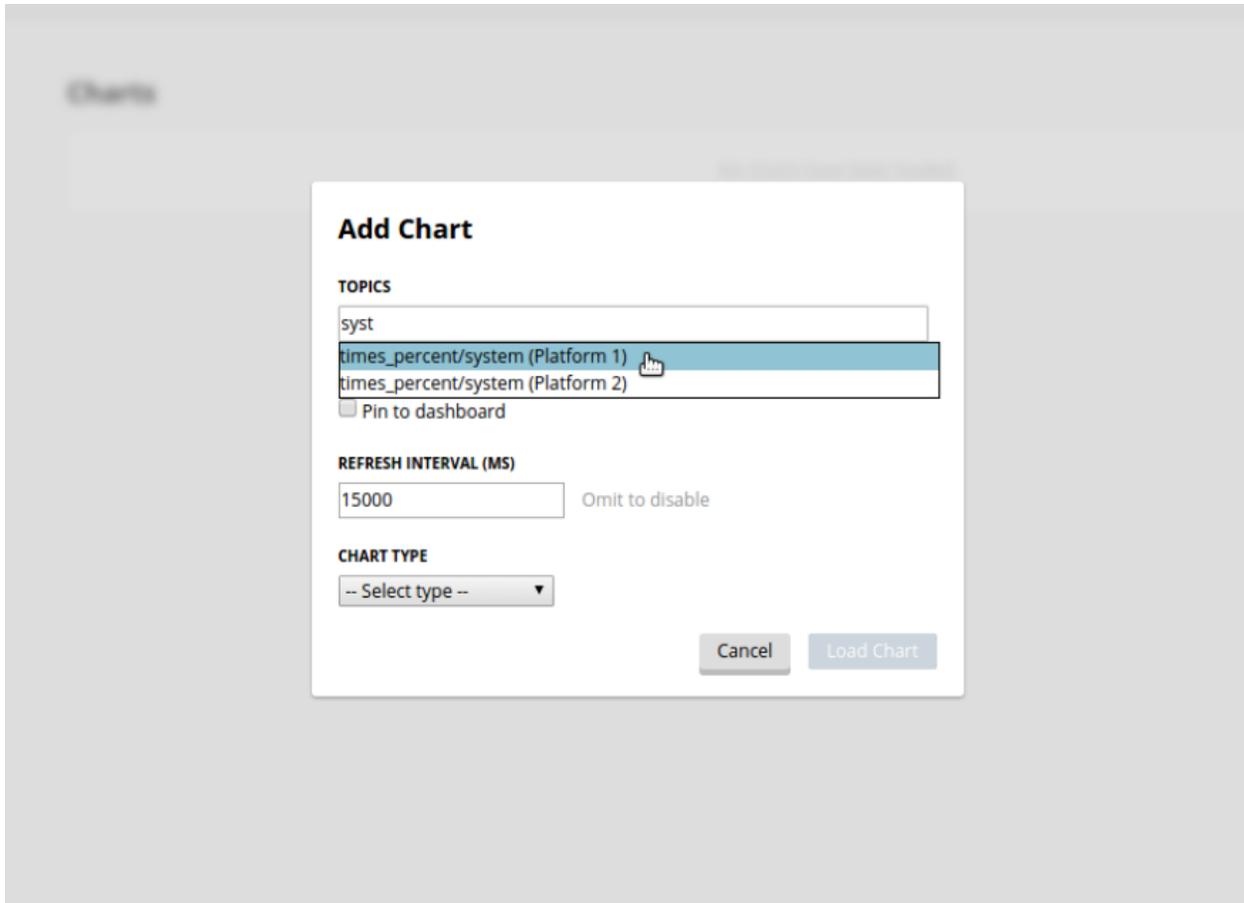
No charts have been loaded.



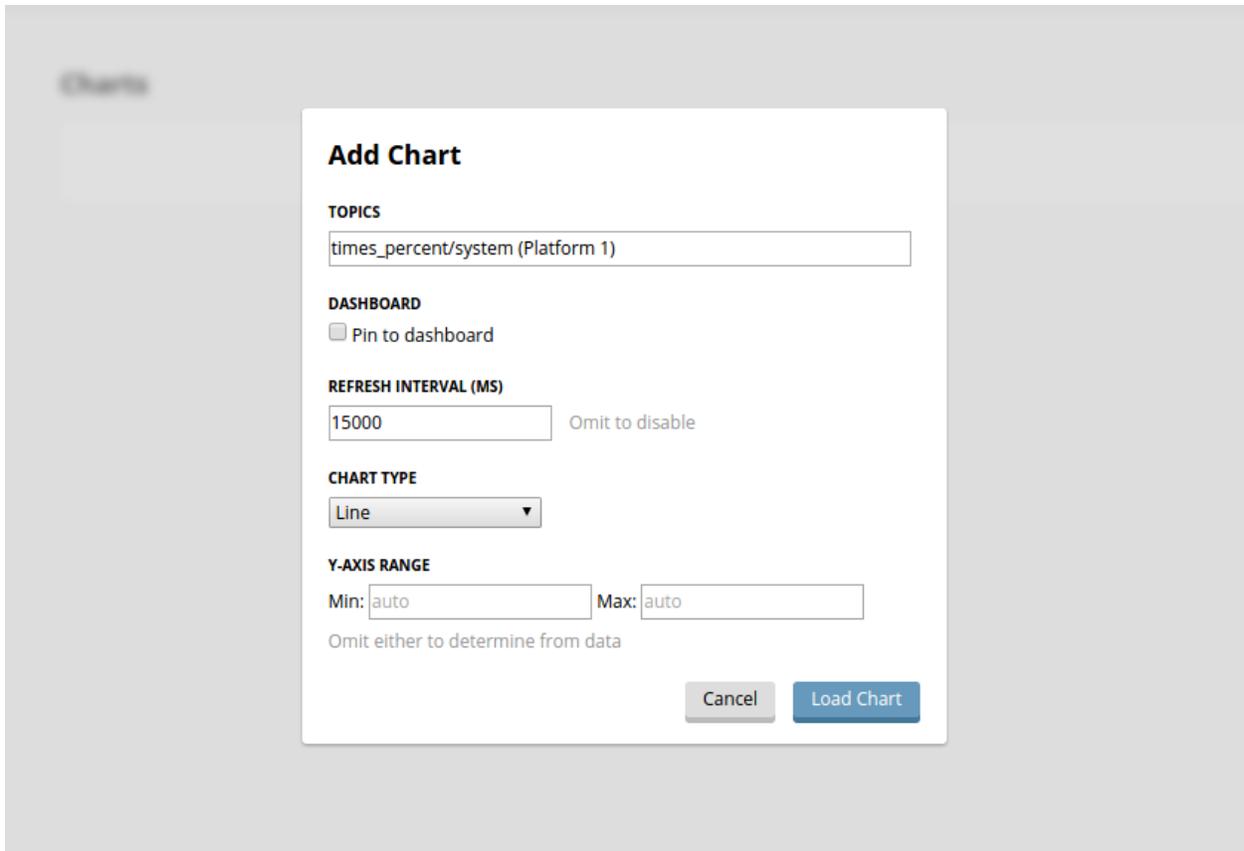
Click in the topics input field to make the list of available chart topics appear.



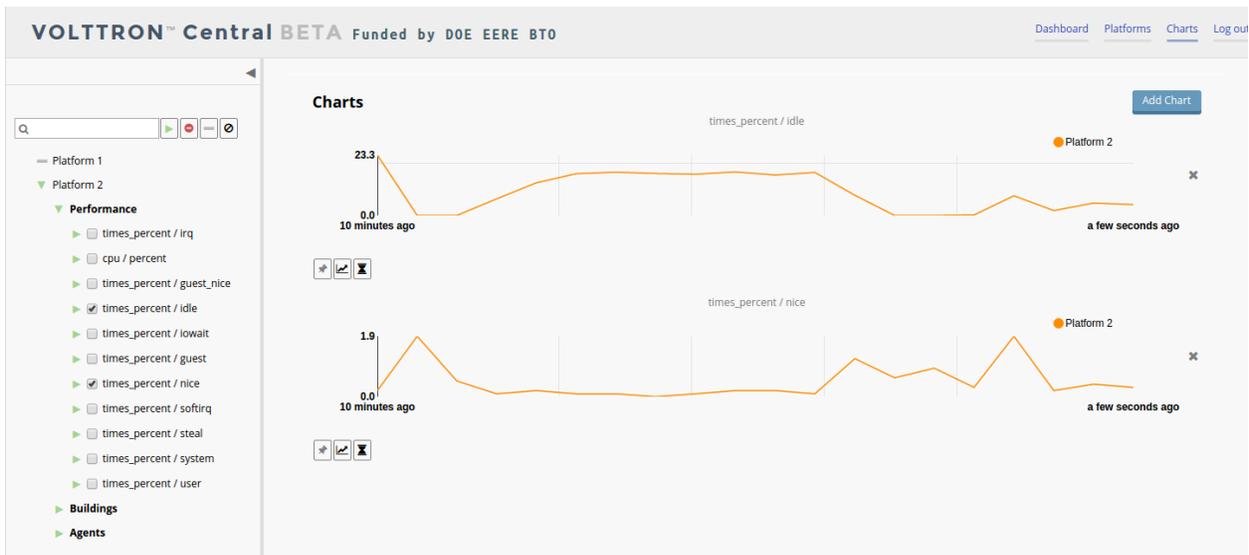
Scroll and select from the list, or type in the field to filter the list, and then select.



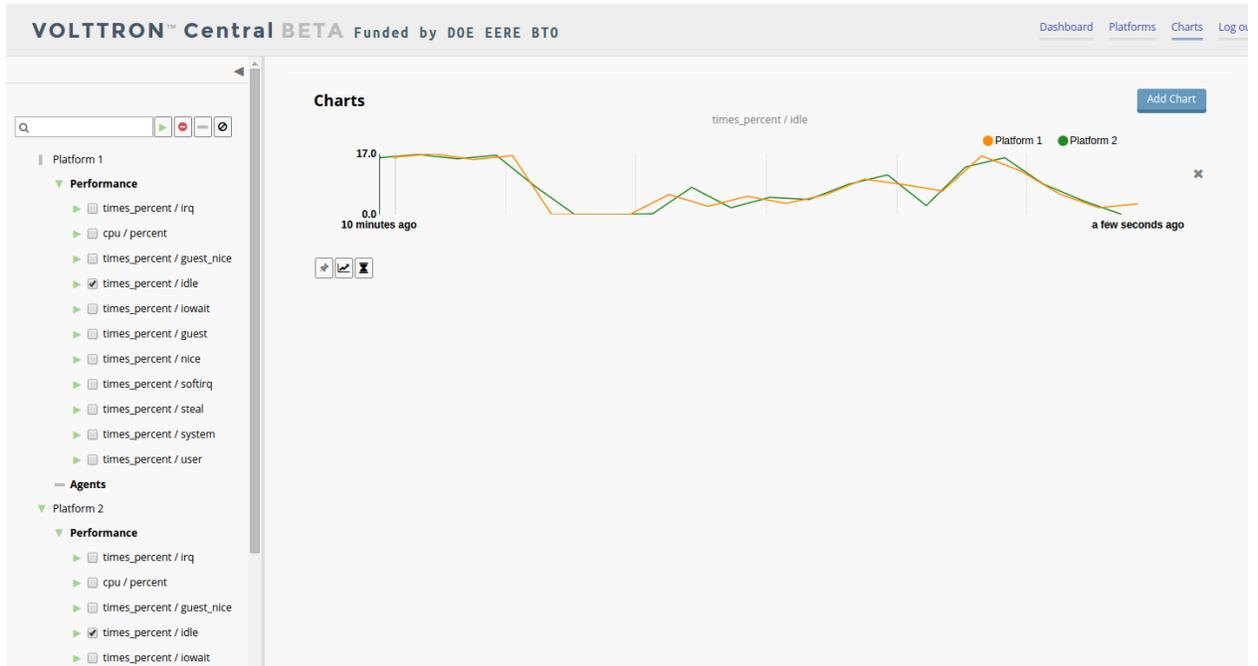
Select a chart type and click the Load Chart button to close the window and load the chart.



To add charts from the side panel, check boxes next to items in the tree.



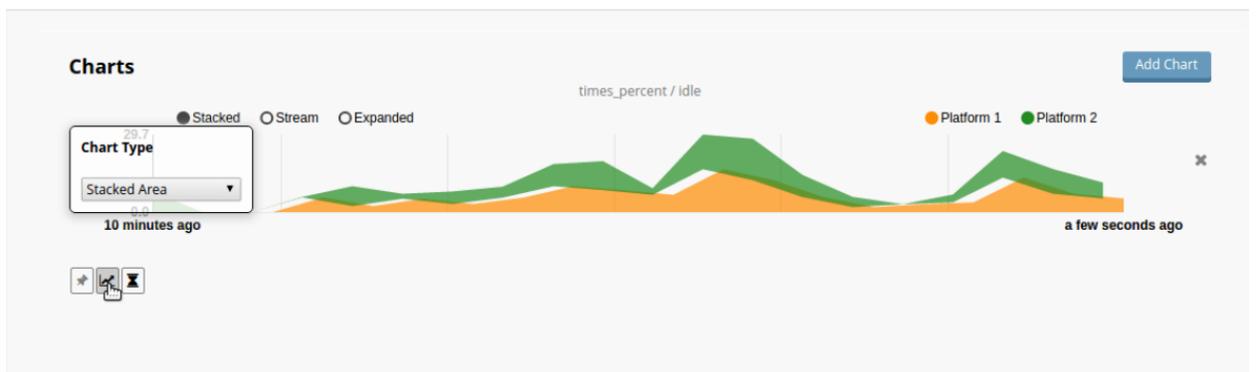
Choose points with the same name from multiple platforms or devices to plot more than one line in a chart.



Move the cursor arrow over the chart to inspect the graphs.

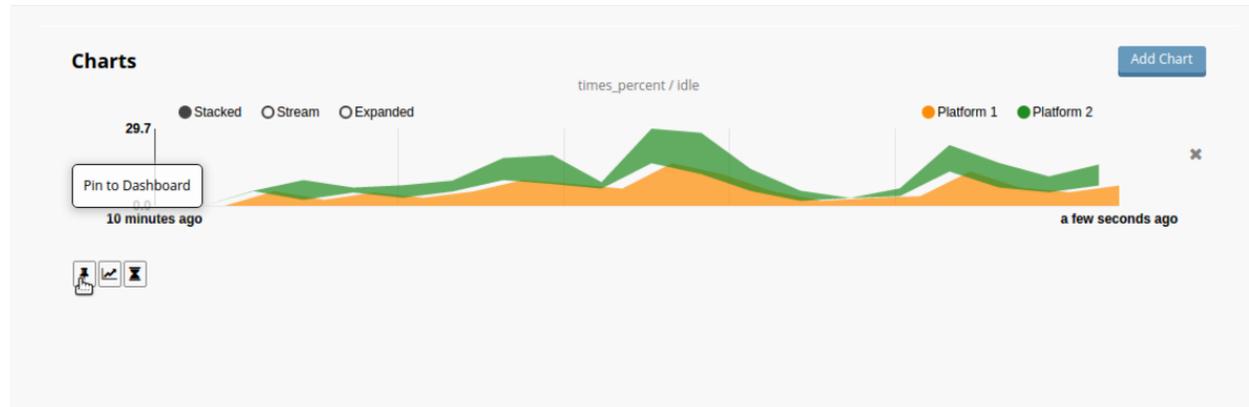


To change the chart's type, click on the Chart Type button and choose a different option.



Dashboard Charts

To pin a chart to the Dashboard, click the Pin Chart button to toggle it. When the pin image is black and upright, the chart is pinned; when the pin image is gray and diagonal, the chart is not pinned and won't appear on the Dashboard.



Charts that have been pinned to the Dashboard are saved to the database and will automatically load when the user logs in to VOLTRON Central. Different users can save their own configurations of dashboard charts.

Remove Charts

To remove a chart, uncheck the box next to the item in the tree or click the X button next to the chart on the Charts page. Removing a chart removes it from the Charts page and the Dashboard.

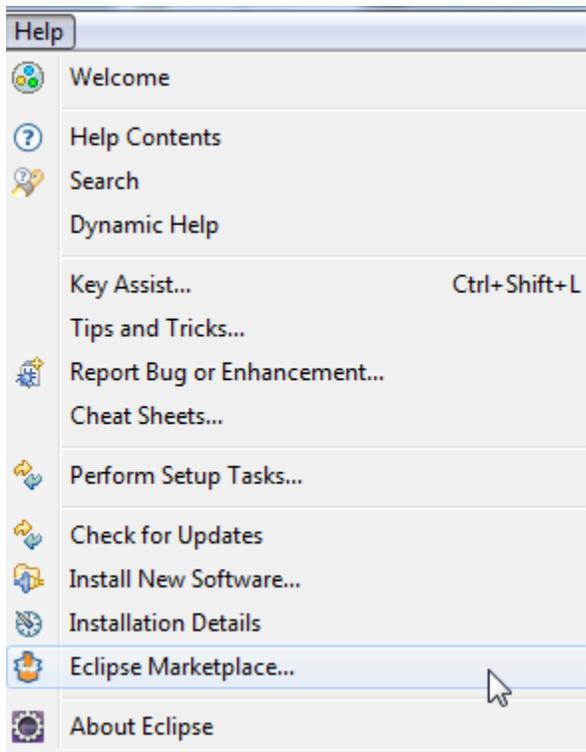
3.4.4 Eclipse IDE Setup

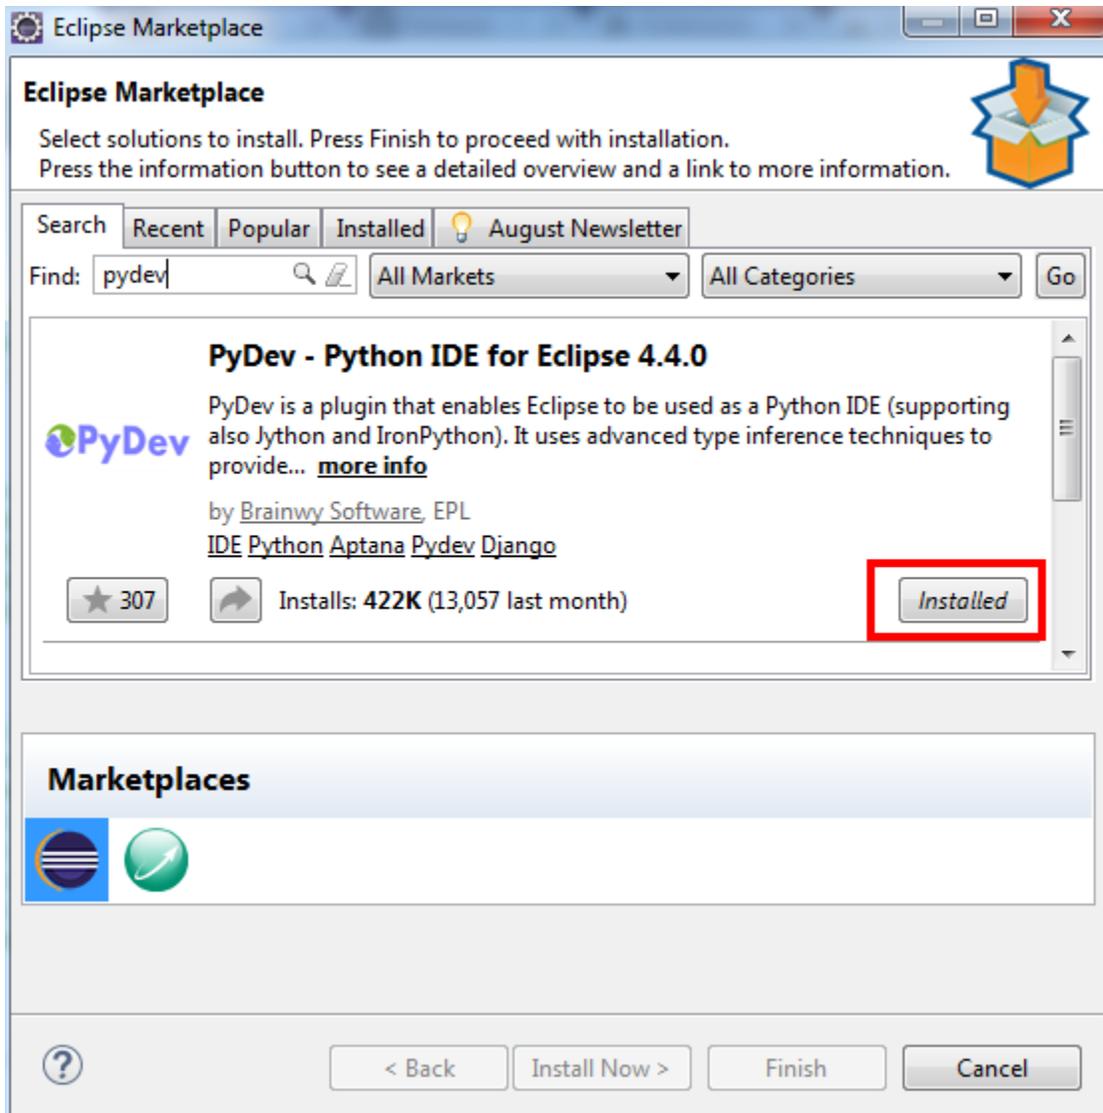
The only thing that is necessary to create a VOLTRON agent is a text editor and the shell. However, we have found the Eclipse Development Environment (IDE) to be a valuable tool for helping to develop VOLTRON agents. You can obtain the latest (MARS as fo 10/7/15) from <http://www.eclipse.org/>. Once downloaded the *PyDev Plugin* is a valuable tool for executing the platform as well as debugging agent code.

- *Install PyDev Plugin*
- *Clone the VOLTRON Source*
- *Build VOLTRON*
- *Link Eclipse to VOLTRON Python Environment*
- *Make Project a PyDev Project*
- *Testing the Installation*
- *Execute VOLTRON Through Shell*
- *Execute VOLTRON Through Eclipse*
- *Start a ListenerAgent*

PyDev Plugin

Installing the PyDev plugin from the Eclipse Market place There is a python plugin for eclipse that makes development much easier. Install it from the eclipse marketplace.





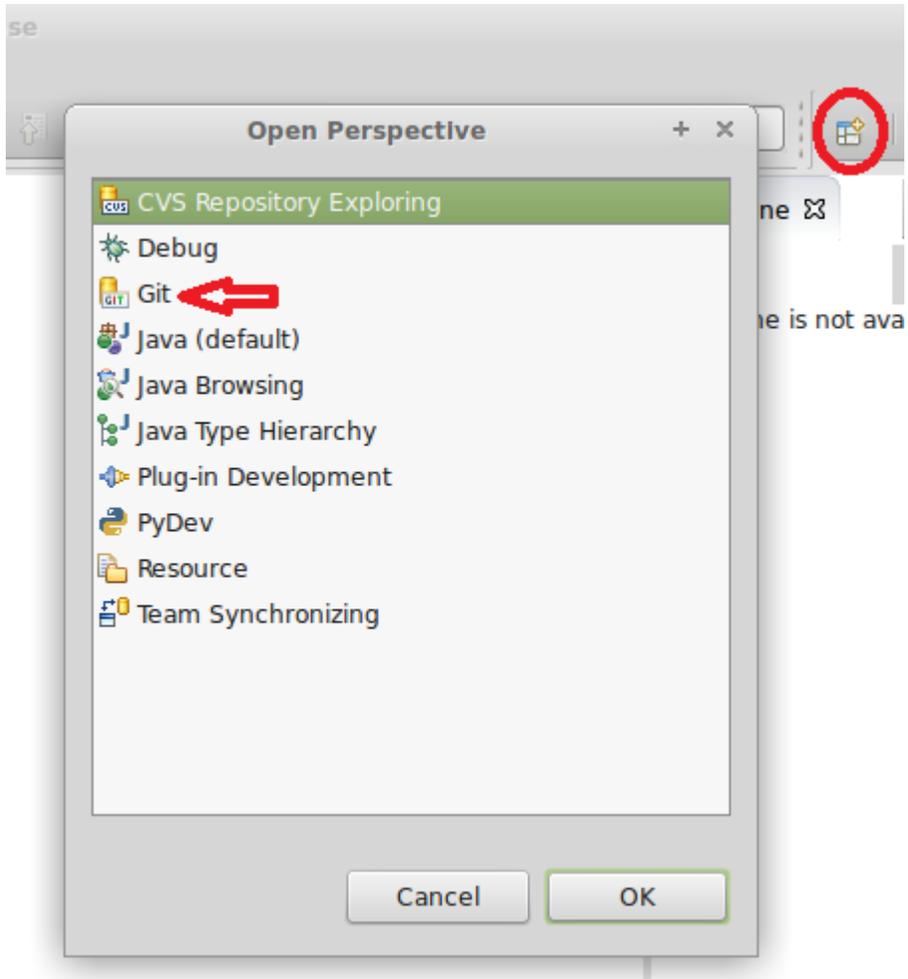
Cloning the Source Code

The VOLTRON code is stored in a git repository. Eclipse (Luna and Mars) come with a git plugin out of the box. For other versions the plugin is available for Eclipse that makes development more convenient (note: you must have Git *already installed* on the system and have *built VOLTRON*):

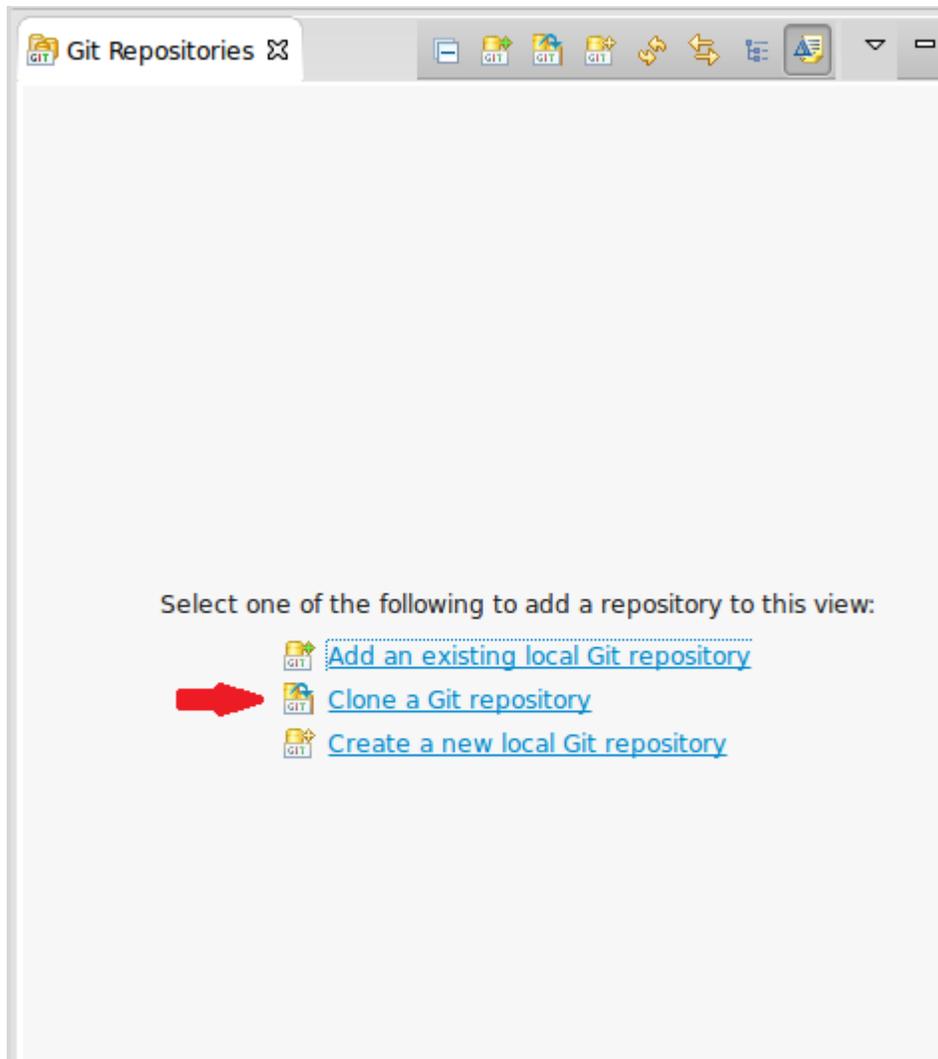
If your version of Eclipse does not have the marketplace follow these instructions.

The project can now be checked out from the repository into Eclipse.

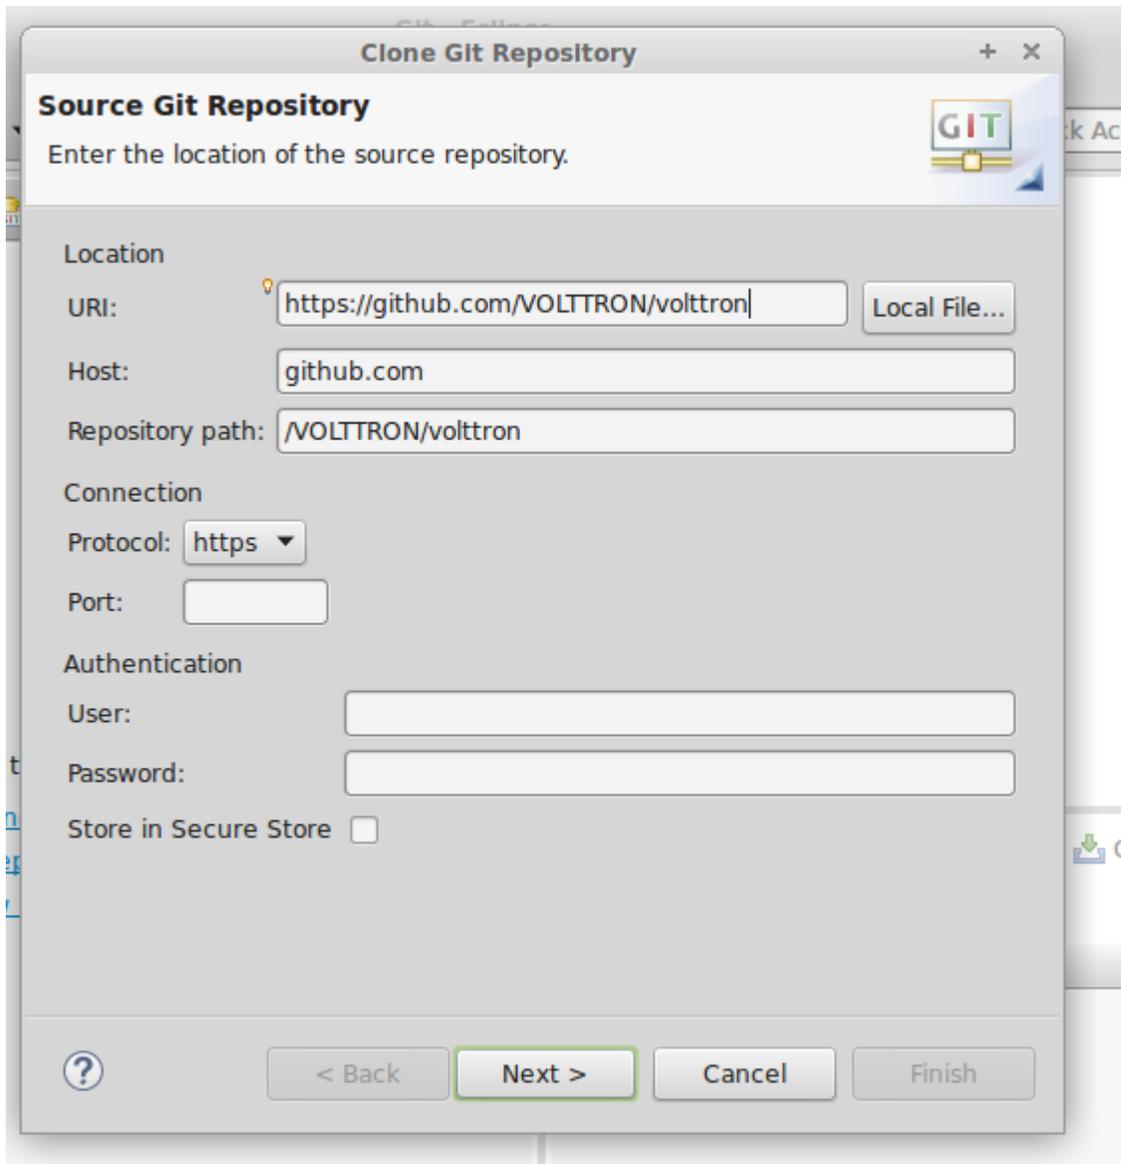
1. Open the Git view



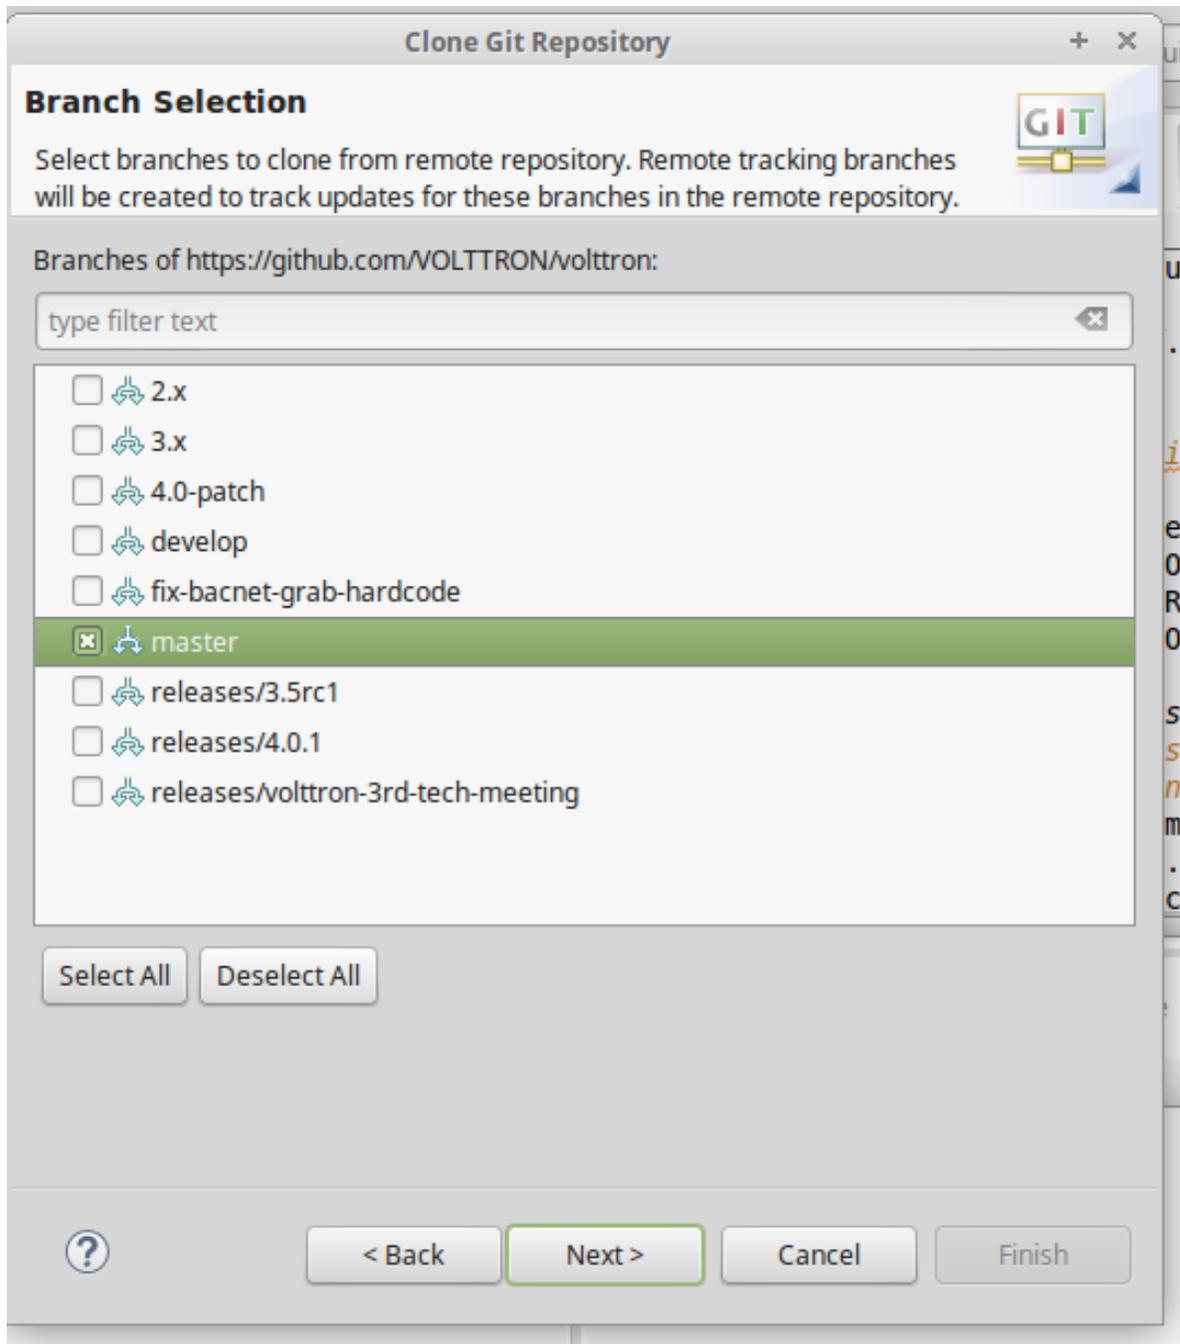
2. Clone a Git Repository



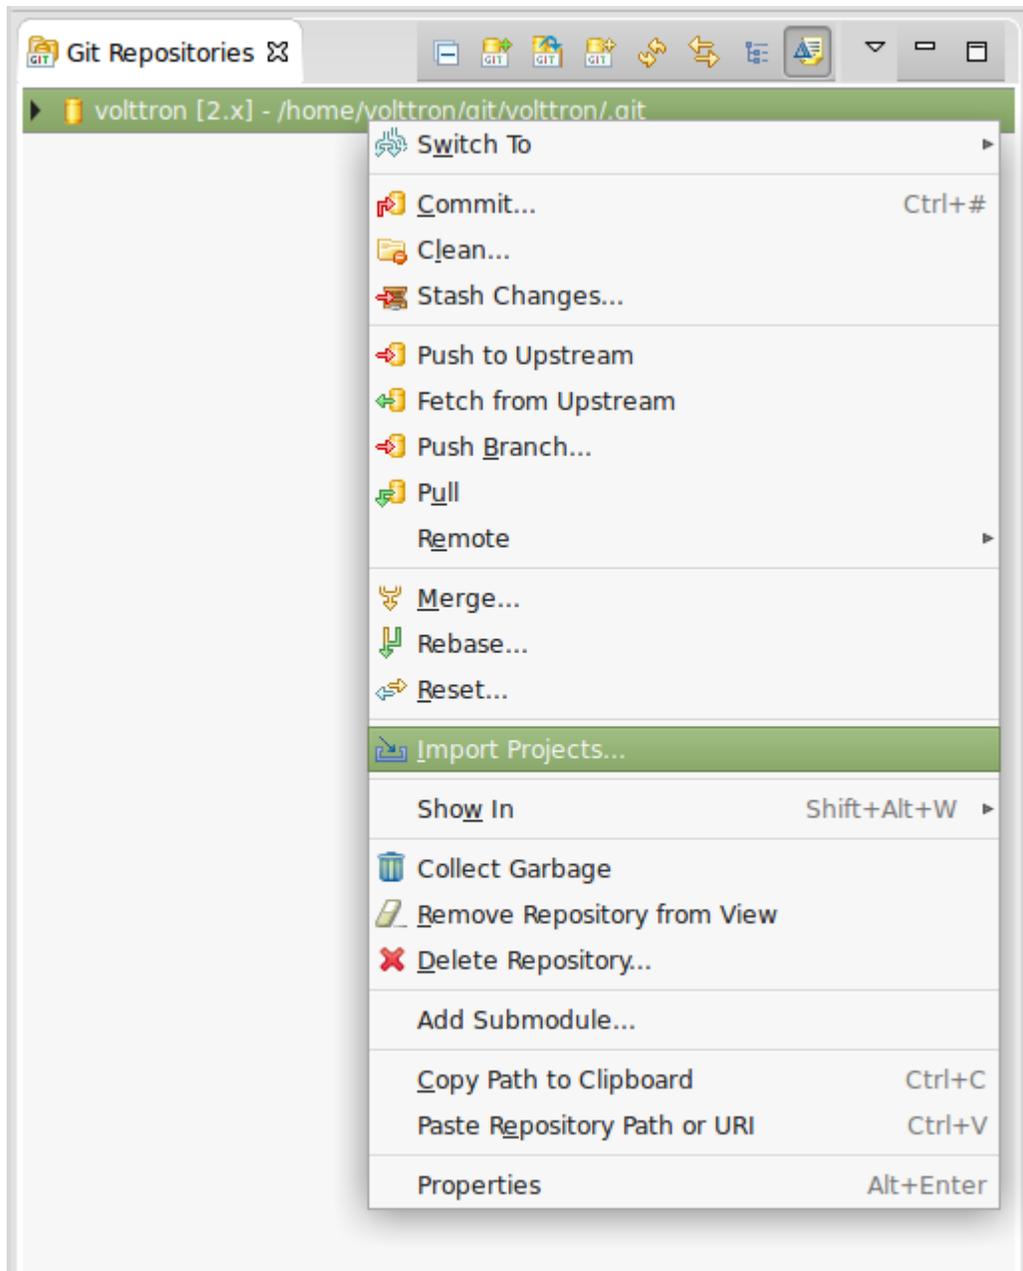
3. Fill out the URI: `https://github.com/VOLTRON/voltron`



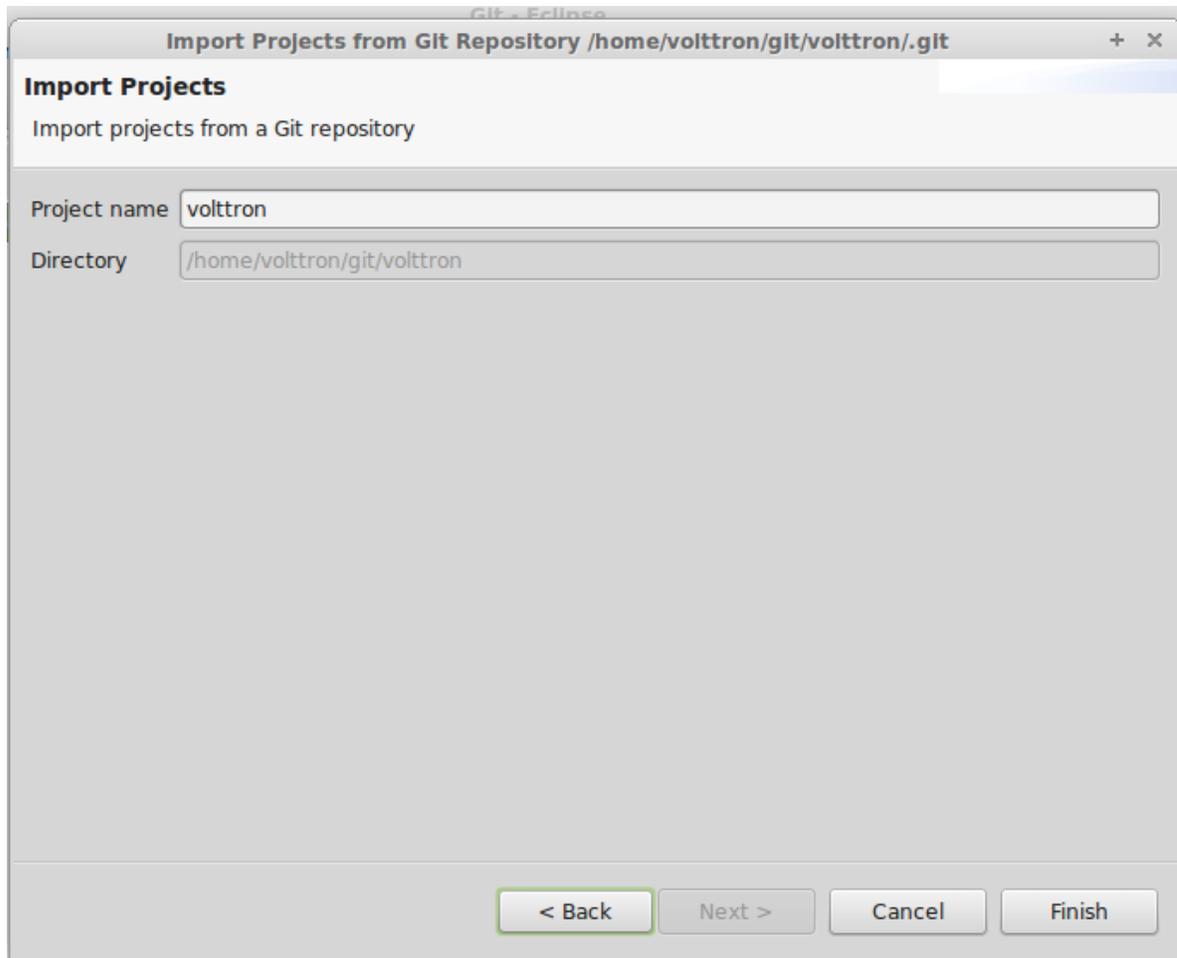
4. Select master for latest stable version



5. Import the cloned repository as a general project



6. Pick a project name (default volttron) and hit Finish



7. Switch to the PyDev perspective

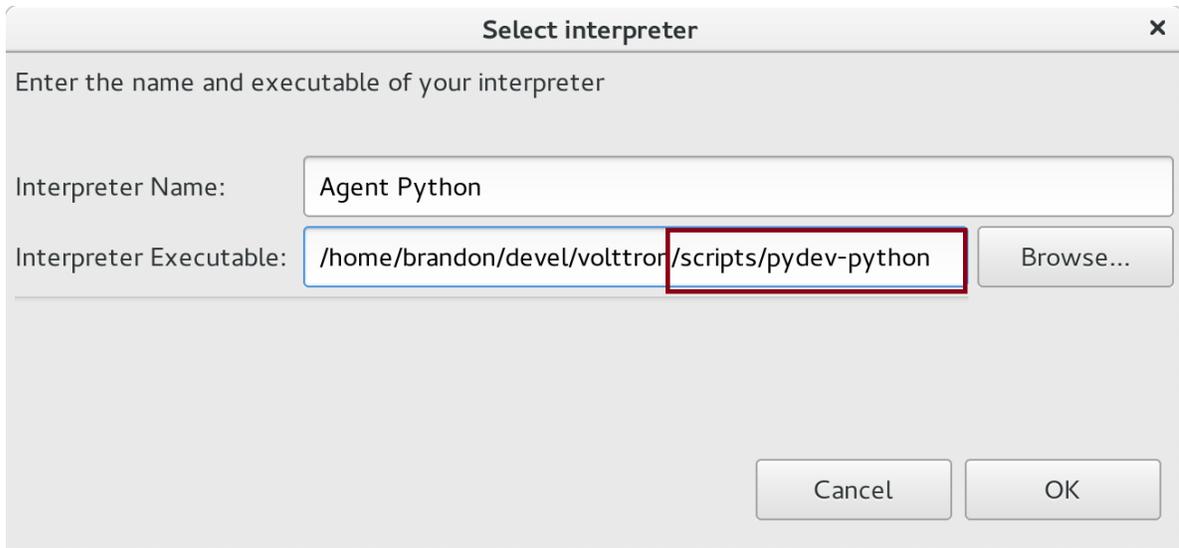
Build VOLTRON

Continue the setup process by opening a command shell. Make the current directory the root of your cloned VOLTRON directory. Follow the instructions in our [Building VOLTRON](#) section of the wiki and then continue below.

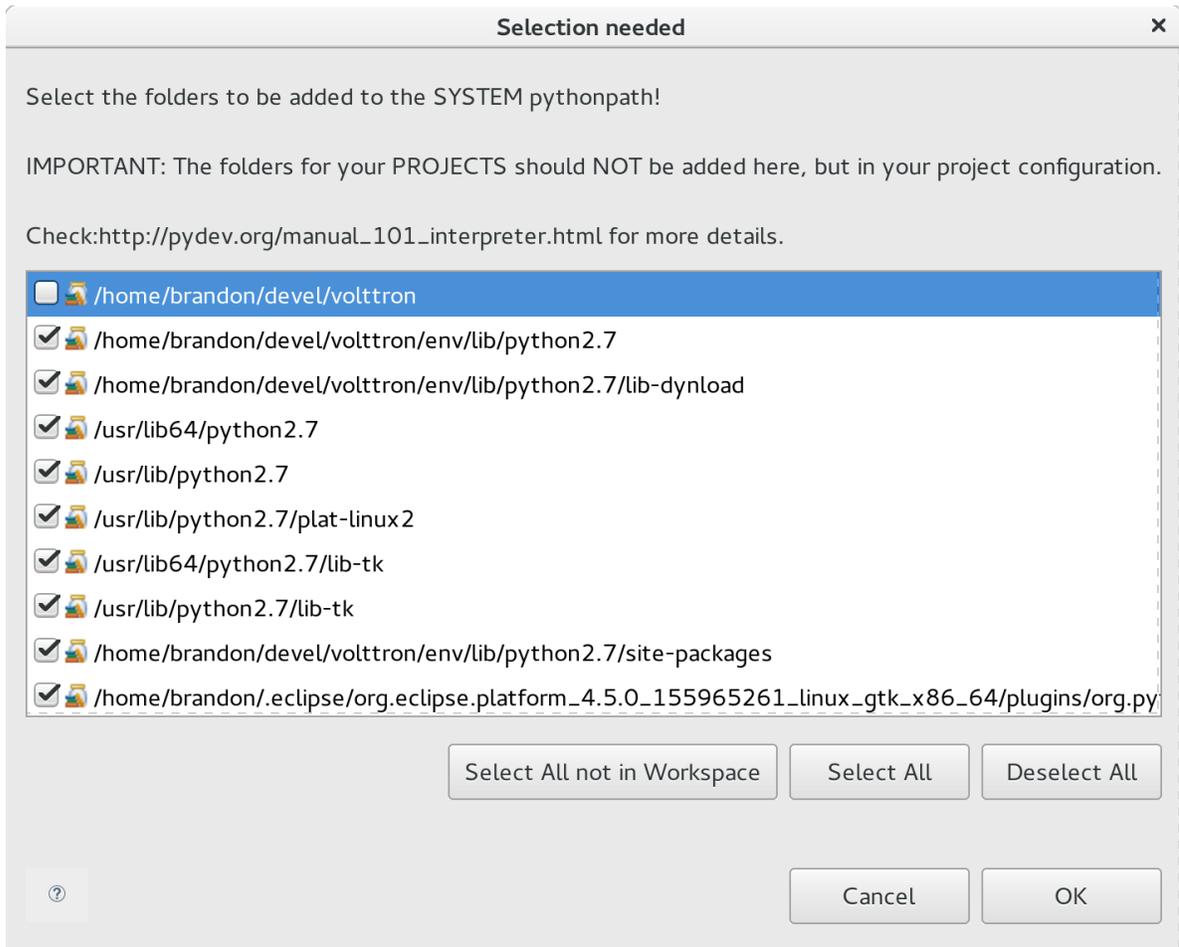
Linking Eclipse and the VOLTRON Python Environment

From the Eclipse IDE right click on the project name and select Refresh so eclipse will be aware of the file system changes. The next step will define the python version that PyDev will use for VOLTRON

1. Choose Window -> Preferences
2. Expand the PyDev tree
3. Select Interpreters -> Python Interpreter
4. Click New
5. Click Browse and browse to the pydev-python file located in scripts directory off of the volttron source
6. Click Ok



7. Select All, then uncheck the VOLTRON root like the picture below

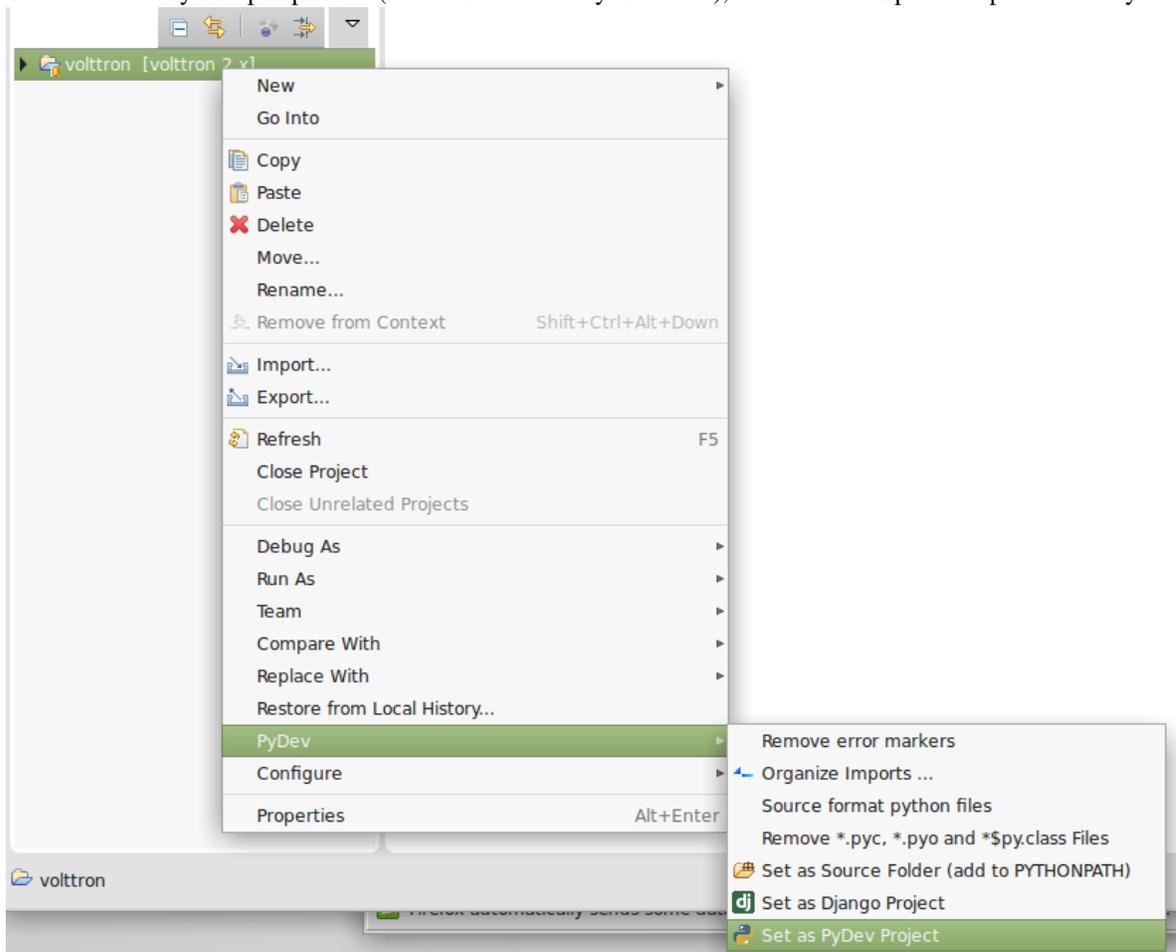


8. Click Ok

Note: You may need redo this stage after platform updates

Make Project a PyDev Project

1. In the Project/PackageExplorer view on the left, right-click on the project, PyDev-> Set as PyDev Project
2. Switch to the PyDev perspective (if it has not already switched), Window -> Open Perspective -> PyDev



Eclipse should now be configured to use the project's environment.

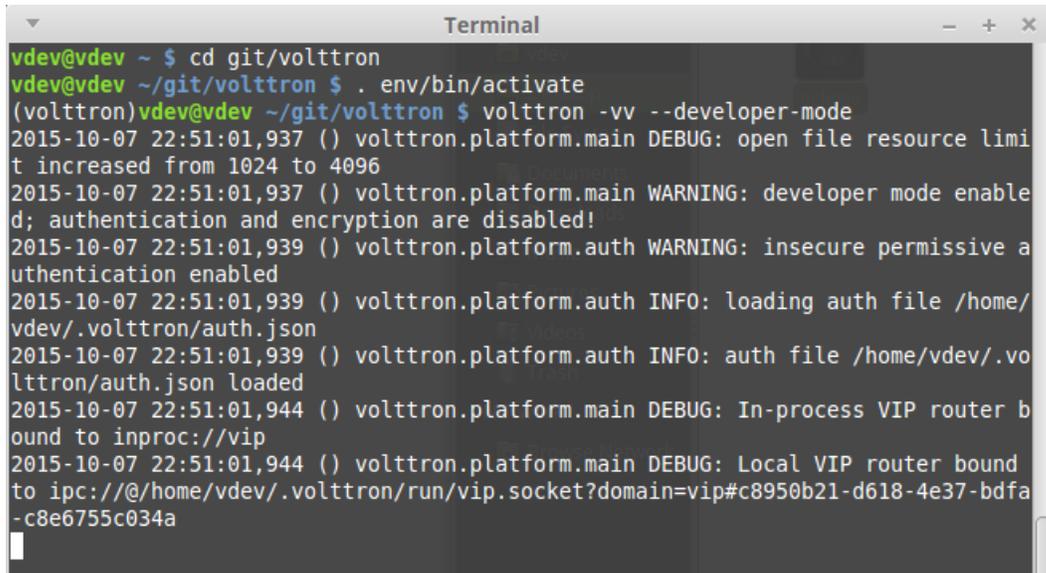
Testing the Installation

In order to test the installation the VOLTRON platform must be running. You can do this either through *the shell* or *through Eclipse*.

Execute VOLTRON Through Shell

1. Open a console and cd into the root of the volttron repository.
2. Execute `source env/bin/activate`

3. Execute `voltron -vv`

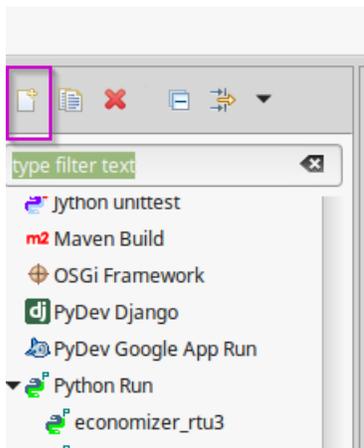


```
Terminal
vdev@vdev ~ $ cd git/voltron
vdev@vdev ~/git/voltron $ . env/bin/activate
(voltron)vdev@vdev ~/git/voltron $ voltron -vv --developer-mode
2015-10-07 22:51:01,937 () voltron.platform.main DEBUG: open file resource limit
increased from 1024 to 4096
2015-10-07 22:51:01,937 () voltron.platform.main WARNING: developer mode enable
d; authentication and encryption are disabled!
2015-10-07 22:51:01,939 () voltron.platform.auth WARNING: insecure permissive a
uthentication enabled
2015-10-07 22:51:01,939 () voltron.platform.auth INFO: loading auth file /home/
vdev/.voltron/auth.json
2015-10-07 22:51:01,939 () voltron.platform.auth INFO: auth file /home/vdev/.vo
ltron/auth.json loaded
2015-10-07 22:51:01,944 () voltron.platform.main DEBUG: In-process VIP router b
ound to inproc://vip
2015-10-07 22:51:01,944 () voltron.platform.main DEBUG: Local VIP router bound
to ipc://@/home/vdev/.voltron/run/vip.socket?domain=vip#c8950b21-d618-4e37-bdfa
-c8e6755c034a
```

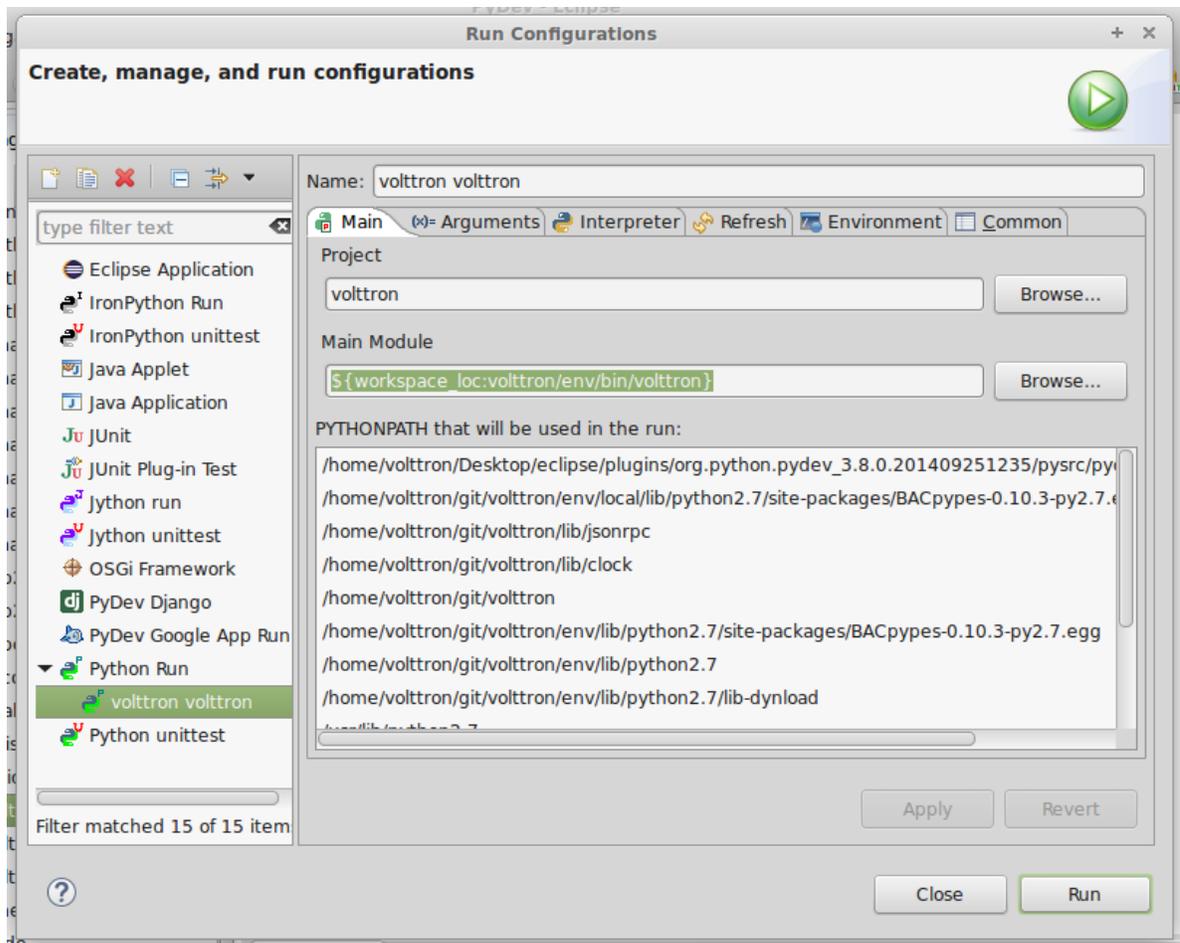
You now have a running VOLTRON logging to standard out. The next step to verifying the installation is to *start a listener agent*.

Execute VOLTRON Through Eclipse

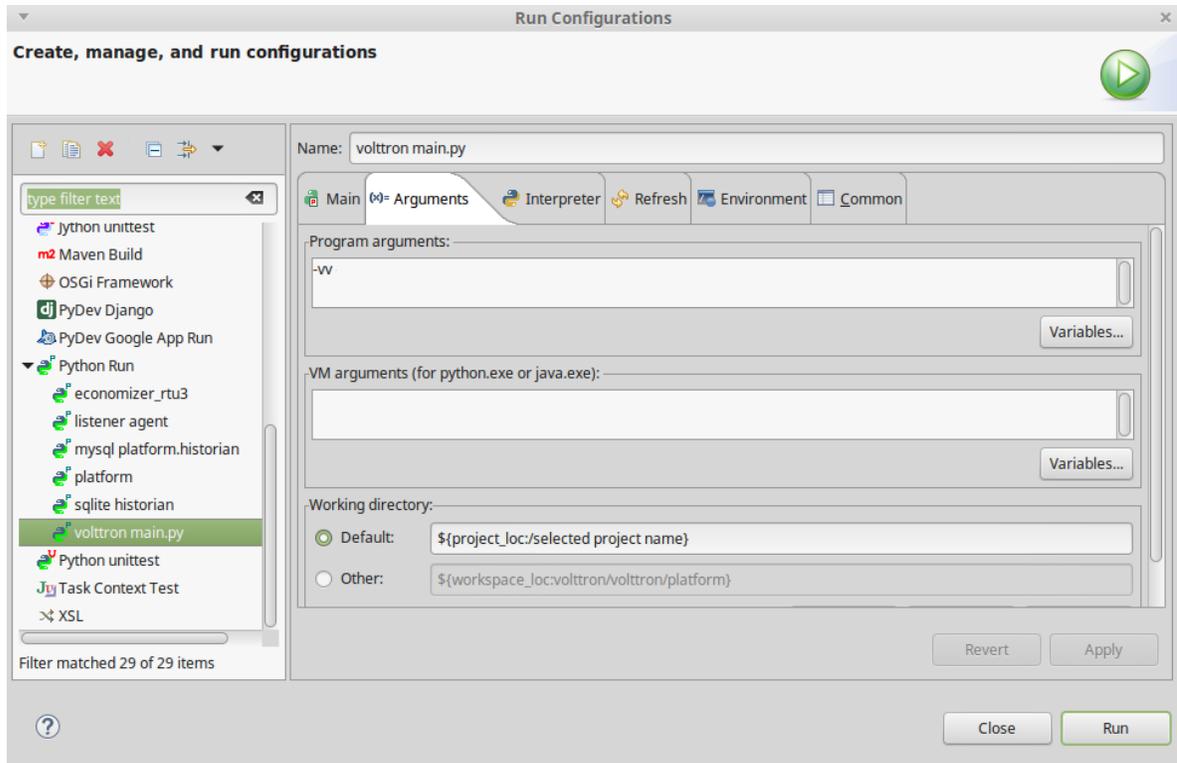
1. Click Run -> Run Configuration from the Eclipse Main Menu
2. Click the New Launch Configuration button



3. Change the name and select the main module `voltron/platform/main.py`



4. Click the Arguments Tab add '-vv' to the arguments and change the working directory to default



5. Click Run. The following image displays the output of a successfully started platform

```

--developer-mode
2015-10-07 22:26:01,812 () main.py WARNING: developer mode enabled; authentication and encryption are disabled!
2015-10-07 22:26:01,817 () voltron.platform.auth WARNING: insecure permissive authentication enabled
2015-10-07 22:26:01,817 () voltron.platform.auth INFO: loading auth file /home/vdev/.voltron/auth.json
2015-10-07 22:26:01,818 () voltron.platform.auth INFO: auth file /home/vdev/.voltron/auth.json loaded
2015-10-07 22:26:01,819 () main.py DEBUG: In-process VIP router bound to inproc://vip
2015-10-07 22:26:01,823 () main.py DEBUG: Local VIP router bound to ipc://@/home/vdev/.voltron/run/vip.socket?dom

```

ref `_Start-Listener-Eclipse`:

Start a ListenerAgent

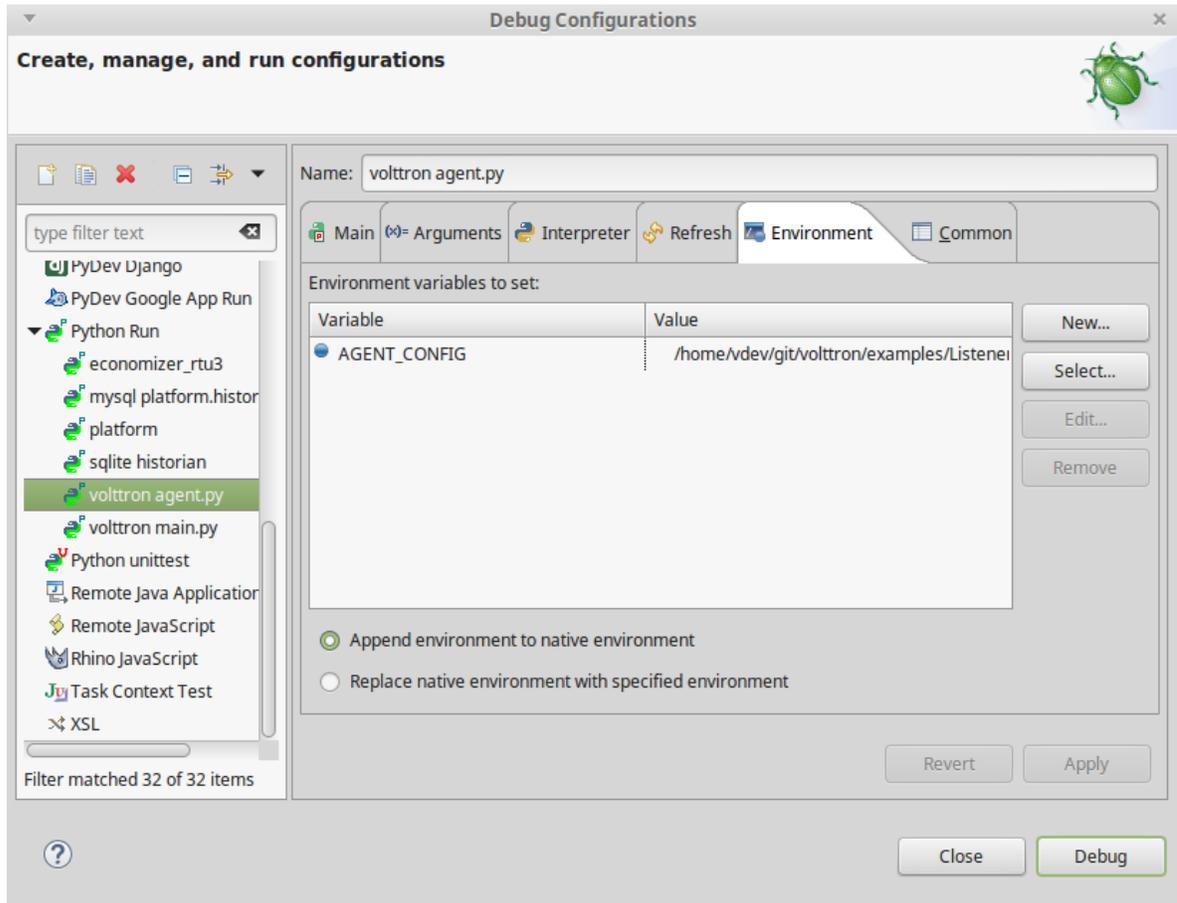
Warning: Before attempting to run an agent in Eclipse, please see the note in: [AgentDevelopment](#)

The listener agent will listen to the message bus for any published messages. It will also publish a heartbeat message ever 10 seconds (by default).

Create a new run configuration entry for the listener agent.

1. In the Package Explorer view, open examples -> ListenerAgent -> listener
2. Right-click on agent.py and select Run As -> Python Run (this will create a run configuration but fail)
3. On the menu bar, pick Run -> Run Configurations...
4. Under Python Run pick "voltron agent.py"
5. Click on the Arguments tab and Change Working Directory to Default

- In the Environment tab, click new set the variable to AGENT_CONFIG with the value of /home/git/voltron/examples/ListenerAgent/config



- Click Run, this launches the agent

You should see the agent start to publish and receive its own heartbeat message in the console.

3.4.5 Pycharm Development Environment

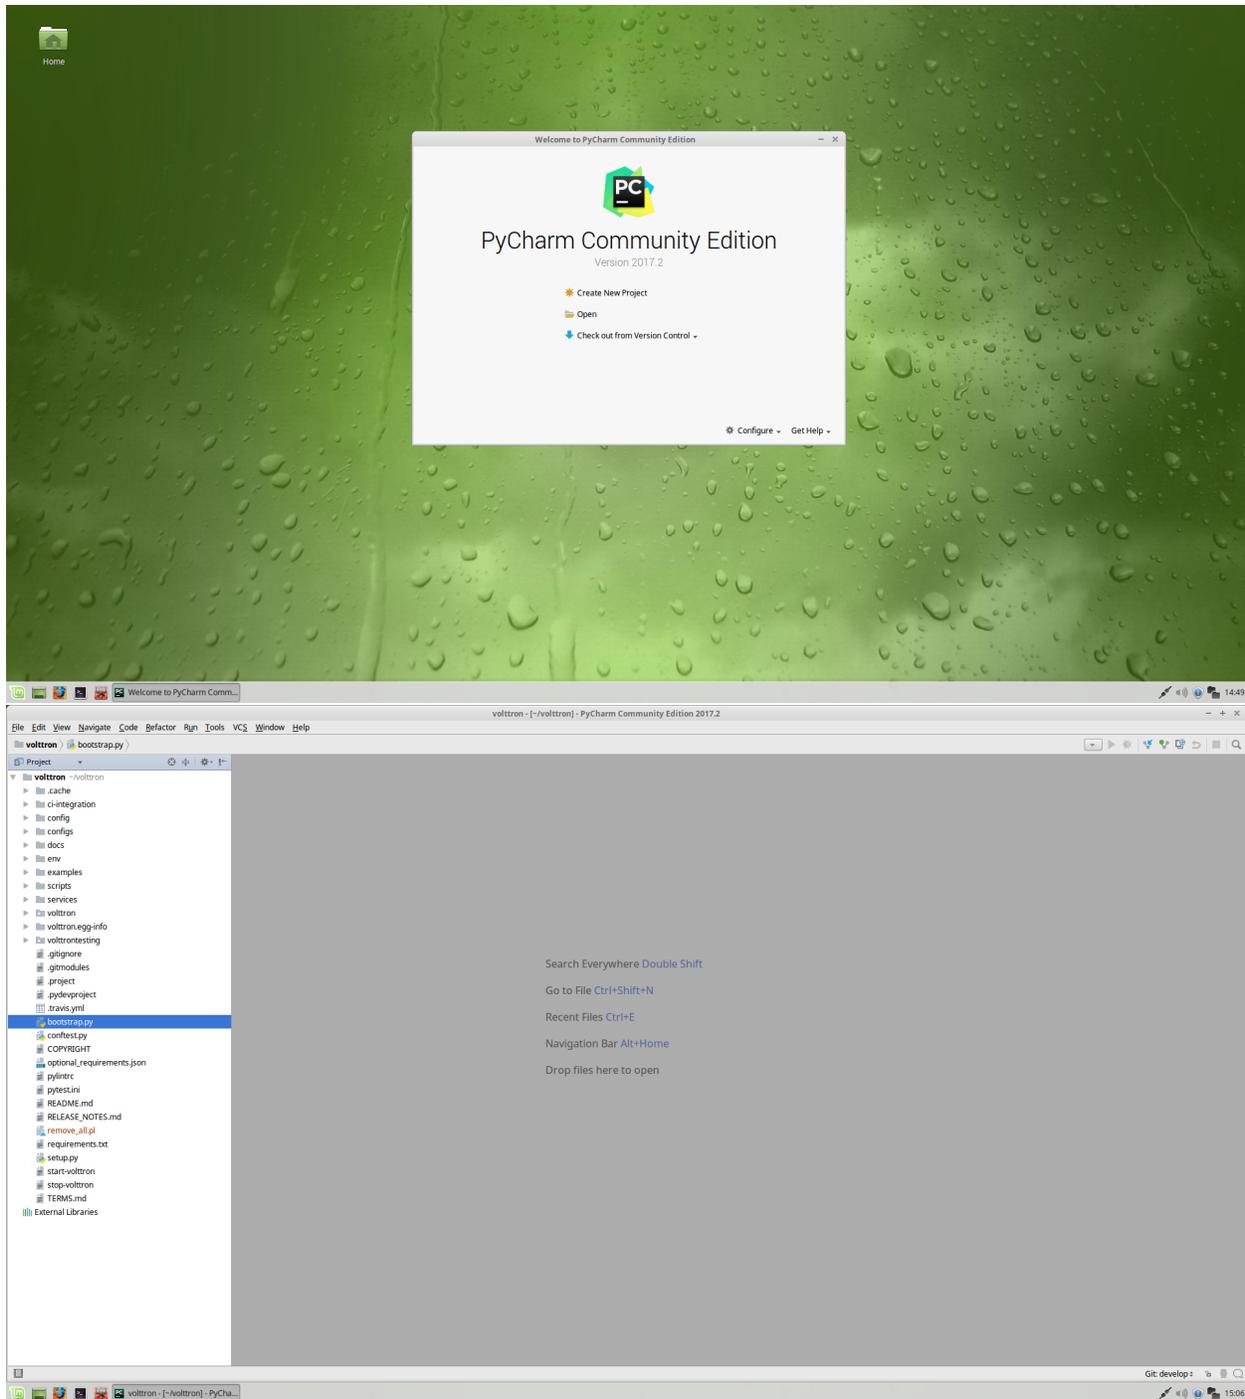
Pycharm is an IDE dedicated to developing python projects. It provides coding assistance and easy access to debugging tools as well as integration with py.test. It is a popular tool for working with VOLTRON. JetBrains provides a free community version that can be downloaded from <https://www.jetbrains.com/pycharm/>

Open Pycharm and Load VOLTRON

When launching Pycharm for the first time we have to tell it where to find the VOLTRON source code. If you have already cloned the repo then point Pycharm to the cloned project. Pycharm also has options to access remote repositories.

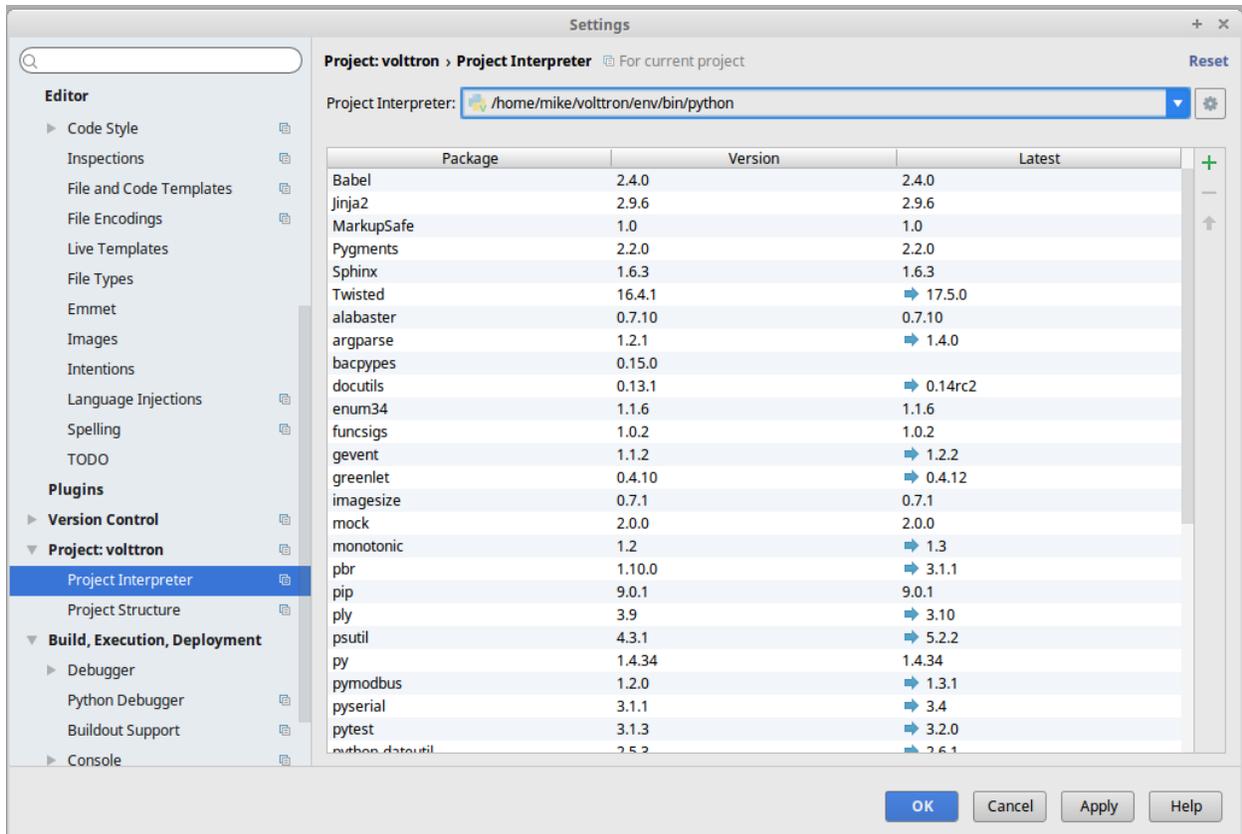
Subsequent instances of Pycharm will automatically load the VOLTRON project.

Note: When getting started make sure to search for *gevent* in the settings and ensure that support for it is enabled.



Set the Project Interpreter

This step should be completed after running the bootstrap script in the VOLTRON source directory. Pycharm needs to know which python environment it should use when running and debugging code. This also tells Pycharm where to find python dependencies. Settings menu can be found under the File option in Pycharm.



Running the VOLTRON Process

If you are not interested in running the VOLTRON process itself in Pycharm then this step can be skipped.

In **Run > Edit Configurations** create a configuration that has `<your source dir>/env/bin/voltron` in the script field, `-vv` in the script parameters field (to turn on verbose logging), and set the working directory to the top level source directory.

VOLTRON can then be run from the Run menu.

Name: Share Single instance only

Configuration | Logs

Script: ...

Script parameters: ...

Environment

Environment variables: ...

Python interpreter: ▾

Interpreter options: ...

Working directory: ...

Add content roots to PYTHONPATH

Add source roots to PYTHONPATH

Emulate terminal in output console

Show command line afterwards

▾ Before launch: Activate tool window

+ - ↗ ↖

Running an Agent

Running an agent is configured similarly to running VOLTTRON proper. In **Run > Edit Configurations** add a configuration and give it the same name as your agent. The script should be the path to *scripts/pycharm-launch.py* and the script parameter must be the path to your agent's *agent.py* file.

In the Environment Variables field add the variable *AGENT_CONFIG* that has the path to the agent's configuration file as its value, as well as *AGENT_VIP_IDENTITY*, which must be unique on the platform.

A good place to keep configuration files is in a directory called *config* in top level source directory; git will ignore changes to these files.

Note: There is an issue with imports in Pycharm when there is a secondary file (i.e. not *agent.py* but another module within the same package). When that happens right click on the directory in the file tree and select **Mark Directory As -> Source Root**

Name: Share Single instance only

Configuration Logs

Script: ...

Script parameters: ...

Environment

Environment variables: ...

Python interpreter: ...

Interpreter options: ...

Working directory: ...

Add content roots to PYTHONPATH

Add source roots to PYTHONPATH

Emulate terminal in output console

Show command line afterwards

Before launch: Activate tool window

Environment Variables

Name	Value
AGENT_VIP_IDENTITY	listener
PYTHONUNBUFFERED	1
AGENT_CONFIG	examples/ListenerAgent/config

Include parent environment variables [Show](#)

There are no tasks to run before launch

File Edit View Navigate Code Refactor Run Tools VCS Window Help

Project: examples ListenerAgent listener agent.py

listener

```

58 from __future__ import absolute_import
59
60 from datetime import datetime
61 import logging
62 import sys
63
64 from pprint import pformat
65
66 from voltron.platform.messaging.health import STATUS_GOOD
67 from voltron.platform.vip.agent import Agent, Core, PubSub, compat
68 from voltron.platform.agent import utils
69 from voltron.platform.messaging import headers as headers_mod
70
71
72 utils.setup_logging()
73 log = logging.getLogger(__name__)
74 version = 3.2
75 DEFAULT_MESSAGE = 'Listener Message'
76 DEFAULT_AGENTID = 'listener'
77 DEFAULT_HEARTBEAT_PERIOD = 5
78
79
80 class ListenerAgent(Agent):
81     """Listens to everything and publishes a heartbeat according to the
82     heartbeat period specified in the settings module.
83     """
84
85     def __init__(self, config_path, **kwargs):
86         super(ListenerAgent, self).__init__(**kwargs)
87         self.config = utils.load_config(config_path)
88         self.agent_id = self.config.get('agentid', DEFAULT_AGENTID)
89         self.message = self.config.get('message', DEFAULT_MESSAGE)
90         self.heartbeat_period = self.config.get('heartbeat_period',
91                                               DEFAULT_HEARTBEAT_PERIOD)
92
93         try:
94             self.heartbeat_period = int(self.heartbeat_period)
95         except:
96             log.warn('Invalid heartbeat period specified setting to default')
97             self.heartbeat_period = DEFAULT_HEARTBEAT_PERIOD
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000

```

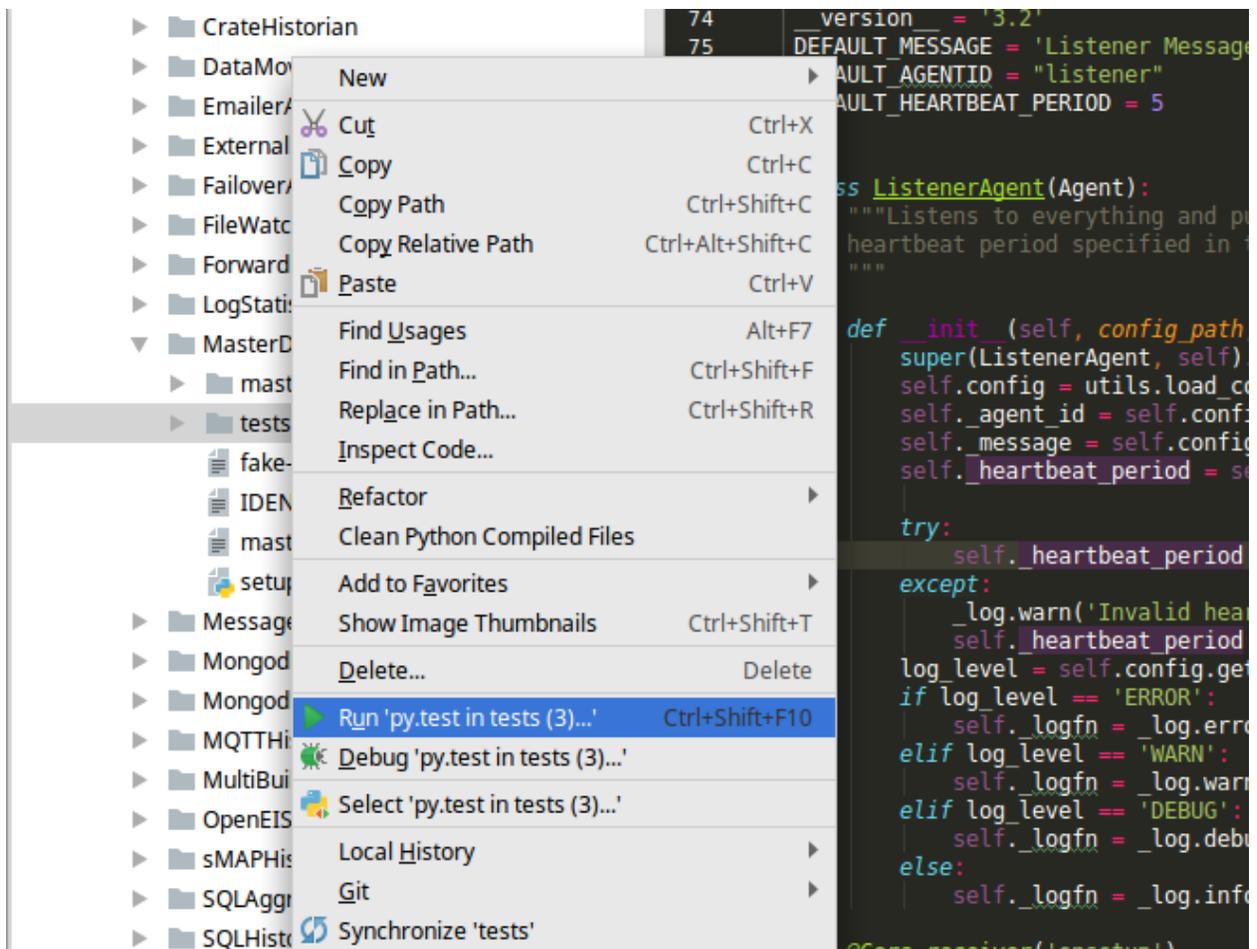
```

2017-08-01 15:59:09,286 main INFO: hello
2017-08-01 15:59:09,301 voltron.platform.vip.agent.core INFO: Connected to platform: router: db37c971-9f6e-43f4-b526-c395b9a0552 version: 1.0 identity: f8c77edf-c34b-46e7-8311-b7358891d300
2017-08-01 15:59:09,301 voltron.platform.vip.agent.core DEBUG: Running onstart methods.
2017-08-01 15:59:09,301 main DEBUG: VERSION: 3.2
2017-08-01 15:59:09,305 main INFO: Peer: 'pubsub', Sender: 'f8c77edf-c34b-46e7-8311-b7358891d300', Bus: 'u', Topic: 'heartbeat/f8c77edf-c34b-46e7-8311-b7358891d300', Headers: {'Date': '2017-08-01T22:59:09.302179+00:00', 'max_compatible_version': 'u', 'min_compatible_version': 'u'}
(context: 'hello',
last updated: '2017-08-01T22:59:09.301925+00:00',
status: 'GOOD')
2017-08-01 15:59:09,306 voltron.platform.vip.agent.subsystems.configstore DEBUG: Processing callbacks for affected files: {}
2017-08-01 15:59:14,307 main INFO: Peer: 'pubsub', Sender: 'f8c77edf-c34b-46e7-8311-b7358891d300', Bus: 'u', Topic: 'heartbeat/f8c77edf-c34b-46e7-8311-b7358891d300', Headers: {'Date': '2017-08-01T22:59:14.305576+00:00', 'max_compatible_version': 'u', 'min_compatible_version': 'u'}
(context: 'hello',
last updated: '2017-08-01T22:59:09.301925+00:00',
status: 'GOOD')
2017-08-01 15:59:19,319 main INFO: Peer: 'pubsub', Sender: 'f8c77edf-c34b-46e7-8311-b7358891d300', Bus: 'u', Topic: 'heartbeat/f8c77edf-c34b-46e7-8311-b7358891d300', Headers: {'Date': '2017-08-01T22:59:19.305561+00:00', 'max_compatible_version': 'u', 'min_compatible_version': 'u'}
(context: 'hello',
last updated: '2017-08-01T22:59:09.301925+00:00',
status: 'GOOD')

```

Testing an Agent

Agent tests written in `pytest` can be run simply by right-clicking the `tests` directory and selecting `Run 'pytest in tests'`, so long as the root directory is set as the VOLTRON source root.



3.4.6 Scalability Experiments

Scalability Setup

Core Platform

- VIP router - how many messages per second can the router pass
- A single agent can connect and send messages to itself as quickly as possible
- Repeat but with multiple agents
- Maybe just increase the number of connected but inactive agents to test lookup times
- Inject faults to test impact of error handling
- Agents
- How many can be started on a single platform?
- How does it affect memory?
- How is CPU affected?

Socket types

- inproc - lockless, copy-free, fast
- ipc - local, reliable, fast
- tcp - remote, less reliable, possibly much slower
- test with different
 - latency
 - throughput
 - jitter (packet delay variation)
 - error rate

Subsystems

- ping - simple protocol which can provide baseline for other subsystems
- RPC - requests per second
- Pub/Sub - messages per second
- How does it scale as subscribers are added

Core Services

- historian
- How many records can be processed per second?
- drivers
 - BACnet drivers use a virtual BACnet device as a proxy to do device communication. Currently there is no known upper limit to the number of devices that can be handled at once. The BACnet proxy opens a single UDP port to do all communication. In theory the upper limit is the point when UDP packets begin to be lost due to network congestion. In practice we have communicated with ~190 devices at once without issue.
 - ModBUS opens up a TCP connection for each communication with a device and then closes it when finished. This has the potential to hit the limit for open file descriptors available to the master driver process. (Before, each driver would run in a separate process, but that quickly uses up sockets available to the platform.) To protect from this the master driver process raises the total allowed open sockets to the hard limit. The number of concurrently open sockets is throttled at 80% of the max sockets. On most Linux systems this is about 3200. Once that limit is hit additional device communications will have to wait in line for a socket to become available.

Twaking tests

- Configure message size
- Perform with/without encryption
- Perform with/without authentication

Hardware profiling

- Perform tests on hardware of varying resources: Raspberry Pi, NUC, Desktop, etc.

Scenarios

- One platform controlling large numbers of devices
- One platform managing large numbers of platforms
- Peer communication (Hardware demo type setup)

Impact on Platform

What is the impact of a large number of devices being scraped on a platform (and how does it scale with the hardware)?

- Historians
 - At what point are historians unable to keep up with the traffic being generated?
 - Is the bottleneck the sqlite cache or the specific implementation (SQLite, MySQL, sMAP)
 - Do historian queues grow so large we have a memory problem?
 - Large number of devices with small number of points vs small number of devices with large number of points
 - How does a large message flow affect the router?
 - Examine effects of the watermark (does increasing help)
 - Response time for volttron-ctl commands (for instance: status)
 - Affect on round trip times (Agent A sends message, Agent B replies, Agent A receives reply)
 - Do messages get lost at some point (EAGAIN error)?
 - What impact does security have? Are things significantly faster in developer-mode? (Option to turn off encryption, no longer available)
- Regulation Agent
 - Every 10 minutes there is an action the master node determines.
 - Duty cycle cannot be faster than that but is set to 2 seconds for simulation. | Some clients miss duty cycle signal | Mathematically each node solves ODE. | Model notes accept switch on/off from master. | Bad to lose connection to clients in the field

Chaos router to introduce delays and dropped packets.

MasterNode needs to have vip address of clients.

Experiment capture historian - not listening to devices, just capturing results

- Go straight to db to see how far behind other historians

Improvements Based on Results

Here is the list of scalability improvements so far:

Reduced the overhead of the base historian by removing unneeded writes to the backup db. Significantly improved performance on low end devices.

Added options to reduce the number of publishes per device per scrape. Common cases where per point publishes and breadth first topics are not needed the driver can be configured only publish the depth first “all” or any combination per device the operator needs. This dramatically decreases the platform hardware requirements while increasing the number of devices that can be scraped.

Added support for staggering device scrapes to reduce CPU load during a scrape.

Further ideas:

Determine if how we use ZMQ is reducing its efficiency.

Remove an unneeded index in historian backup db.

Increase backup db page count.

Scalability Planning

Goals

- Determine the limits of the number of devices that can be interacted with via a single Volttron platform.
- Determine how scaling out affects the rate at which devices are scraped. i.e. How long from the first device scrape to the last?
- Determine the effects of socket throttling in the master driver on the performance of Modbus device scraping.
- Measure total memory consumption of the Master Driver Agent at scale.
- Measure how well the base history agent and one or more of the concrete agents handle a large amount of data.
- Determine the volume of messages that can be achieved on the pubsub before the platform starts rejecting them.

Test framework

Test Devices

Simple, command line configured virtual devices to test against in both Modbus and BACnet flavors. Devices should create 10 points to read that generate either random or easily predictable (but not necessarily constant) data. Process should be completely self contained.

Test devices will be run on remote hosts from the Volttron test deployment.

Launcher Script

- The script will be configurable as to the number and type of devices to launch.
- The script will be configurable as to the hosts to launch virtual devices on.
- The script (probably a fabric script) will push out code for and launch one or more test devices on one or more machines for the platform to scrape.
- The script will generate all of the master driver configuration files to launch the master driver.
- The script may launch the master driver.
- The script may launch any other agents used to measure performance.

Shutdown Script

- The script (probably the same fabric script run with different options) will shutdown all virtual drivers on the network.
- The script may shutdown the master driver.
- The script may shutdown any related agents.

Performance Metrics Agent

This agent will track the publishes by the different drivers and generate data in some form to indicate:

- Total time for all devices to be scraped
- Any devices that were not successfully scraped.
- Performance of the message bus.

Additional Benefits

Most parts of a test bed run should be configurable. If a user wanted to verify that the Master Driver worked, for instance, they could run the test bed with only a few virtual device to confirm that the platform is working correctly.

Running a simple test

You will need 2 open terminals to run this test. (3 if you want to run

the platform in it's own terminal) | Checkout the feature/scalability branch.

Start the platform.

Go to the `voltron/scripts/scalability-testing` directory in two different terminals. (Both with the environment activated)

In one terminal run:

```
python config_builder.py --count=1500 --scalability-test --scalability-test-  
↪iterations=6 fake fake18.csv localhost
```

Change the path to `fake.csv` as needed.

(Optional) After it finishes run:

```
./launch_fake_historian.sh
```

to start the null historian.

In a separate terminal run:

```
./launch_scalability_drivers.sh
```

to start the scalability test.

This will emulate the scraping of 1500 devices with 18 points each 6 times, log the timing, and quit.

Redirecting the driver log output to a file can help improve performance. Testing should be done with and without the null historian.

Currently only the depth first all is published by drivers in this branch. Uncomment the other publishes in driver.py to test out full publishing. fake.csv has 18 points.

Optionally you can run two listener agents from the volttron/scripts directory in two more terminals with the command:

```
./launch_listener.sh
```

and rerun the test to see the how it changes the performance.

Real Driver Benchmarking

Scalability testing using actual MODBUS or BACnet drivers can be done using the virtual device applications in the scripts/scalability-testing/virtual-drivers/ directory. The configuration of the master driver and launching of these virtual devices on a target machine can be done automatically with fabric.

Setup

This requires two computers to run: One for the VOLTRON platform to run the tests on (“the platform”) and a target machine to host the virtual devices (“the target”).

Target setup

The target machine must have the VOLTRON source with the feature/scalability branch checked out and bootstrapped. Make a note of the directory of the VOLTRON code.

Platform setup

With the VOLTRON environment activated install fabric.

```
pip install fabric
```

Edit the file scripts/scalability-testing/test_settings.py as needed.

- virtual_device_host (string) - Login name and IP address of the target machine. This is used to remotely start and stop virtual devices via ssh. “volttron@10.0.0.1”
- device_types - map of driver types to tuple of the device count and registry config to use for the virtual devices. Valid device types are “bacnet” and “modbus”.
- volttron_install - location of volttron code on the target.

To configure the driver on the platform and launch the virtual devices on the target run

```
fab deploy_virtual_devices
```

When prompted enter the password for the target machine. Upon completion virtual devices will be running on the target and configuration files written for the master driver.

Launch Test

If your test includes virtual BACnet devices be sure to configure and launch the BACnet Proxy before launching the scalability driver test.

(Optional)

```
./launch_fake_historian.sh
```

to start the null historian.

In a separate terminal run:

```
./launch_scalability_drivers.sh
```

to start the scalability test.

To stop the virtual devices run

```
fab stop_virtual_devices
```

and enter the user password when prompted.

3.4.7 Examples/Samples

Agents

CAgent

The C Agent uses the ctypes module to load a shared object into memory so its functions can be called from python.

There are two versions of the C Agent. The first is a standard agent that can be installed with the make agent script. The other is a driver interface for the master driver.

Building the Shared Object

The shared object library must be built before installing C Agent examples. Running *make* in the C Agent source directory will compile the provided C code using the position independent flag; a requirement for creating shared objects.

Files created by make can be removed by running *make clean*.

Agent Installation

After building the shared object library the standard agent can be installed with the make-agent script.

The driver interface example must be copied or moved to the master driver's interface directory. The C Driver configuration tells the interface where to find the shared object. An example is available in the C Agent's *driver* directory.

CSVHistorian

The CSV Historian Agent is an example historian agent that writes device data to the CSV file specified in the configuration file.

This is the code created during Kyle Monson's presentation on VOLTTRON Historians at the 2017 VOLTTRON Technical Meeting.

Explanation of CSVHistorian

Setup logging for later.

```
utils.setup_logging()
_log = logging.getLogger(__name__)
```

The *historian* method is called by *utils.vip_main* when the agents is started (see below). *utils.vip_main* expects a callable object that returns an instance of an Agent. This method of dealing with a configuration file and instantiating an Agent is common practice.

```
def historian(config_path, **kwargs):
    if isinstance(config_path, dict):
        config_dict = config_path
    else:
        config_dict = utils.load_config(config_path)

    output_path = config_dict.get("output", "~/historian_output.csv")

    return CSVHistorian(output_path = output_path, **kwargs)
```

All historians must inherit from *BaseHistorian*. The *BaseHistorian* class handles the capturing and caching of all device, logging, analysis, and record data published to the message bus.

```
class CSVHistorian(BaseHistorian):
```

The Base Historian creates a separate thread to handle publishing data to the data store. In this thread the Base Historian calls two methods on the created historian, *historian_setup* and *publish_to_historian*.

The Base Historian created the new thread in it's *__init__* method. This means that any instance variables must assigned in *__init__* before calling the Base Historian's *__init__* method.

```
def __init__(self, output_path="", **kwargs):
    self.output_path = output_path
    self.csv_dict = None
    super(CSVHistorian, self).__init__(**kwargs)
```

Historian setup is called shortly after the new thread starts. This is where a Historian sets up a connect the first time. In our example we create the Dictwriter object that we will use to create and add lines to the CSV file.

We keep a reference to the file object so that we may flush its contents to disk after writing the header and after we have written new data to the file.

The CSV file we create will have 4 columns: *timestamp*, *source*, *topic*, and *value*.

```
def historian_setup(self):
    self.f = open(self.output_path, "wb")
    self.csv_dict = csv.DictWriter(self.f, ["timestamp", "source", "topic", "value"])
    self.csv_dict.writeheader()
    self.f.flush()
```

publish_to_historian is called when data is ready to be published. It is passed a list of dictionaries. Each dictionary contains a record of a single value that was published to the message bus.

The dictionary takes the form:

```
{
    '_id': 1,
```

(continues on next page)

(continued from previous page)

```

'timestamp': timestamp1.replace(tzinfo=pytz.UTC), #Timestamp in UTC
'source': 'scrape', #Source of the data point.
'topic': "pnnl/isbl/hvac1/thermostat", #Topic that published to without prefix.
'value': 73.0, #Value that was published
'meta': {"units": "F", "tz": "UTC", "type": "float"} #Meta data published with_
↪the topic
}

```

Once the data is written to the historian we call `self.report_all_handled()` to inform the *BaseHistorian* that all data we received was successfully published and can be removed from the cache. Then we can flush the file to ensure that the data is written to disk.

```

def publish_to_historian(self, to_publish_list):
    for record in to_publish_list:
        row = {}
        row["timestamp"] = record["timestamp"]

        row["source"] = record["source"]
        row["topic"] = record["topic"]
        row["value"] = record["value"]

        self.csv_dict.writerow(row)

    self.report_all_handled()
    self.f.flush()

```

This agent does not support the Historian Query interface.

Agent Testing

The CSV Historian can be tested by running the included `launch_my_historian.sh` script.

Agent Installation

This Agent may be installed on the platform using the standard method.

Example Agents Overview

Some example agents are included with the platform to help explore its features.

- *DataPublisher*
- *ListenerAgent*
- *Process Agent*
- *SchedulerExampleAgent*
- *CAgent*
- *DDSAgent*
- *CSVHistorian*

More complex agents contributed by other researchers can also be found in the examples directory. It is recommended that developers new to VOLTRON understand the example agents first before diving into the other agents.

Example Agent Conventions

Some of the example agent classes are defined inside a method, for instance:

```
def ScheduleExampleAgent (config_path, **kwargs):
    config = utils.load_config(config_path)
    campus= config['campus']
```

This allows configuration information to be extracted from an agent config file for use in topics.

```
@Pubsub.subscribe('pubsub', DEVICES_VALUE (campus=campus))
def actuate(self, peer, sender, bus, topic, headers, message):
```

Fake Driver

The FakeDriver is included as a way to quickly see data published to the message bus in a format that mimics what a true Driver would produce. This is an extremely simple implementation of the *VOLTRON driver framework*

Make a script to build and deploy the fake driver.

- Create a config directory (if one doesn't already exist). All local config files will be worked on here.
- `cp examples/configurations/drivers/fake.config config/`
- Edit `registry_config` for the paths on your system

fake.config:

```
{
  "driver_config": {},
  "registry_config": "config://fake.csv",
  "interval": 5,
  "timezone": "US/Pacific",
  "heart_beat_point": "Heartbeat",
  "driver_type": "fakedriver",
  "publish_breadth_first_all": false,
  "publish_depth_first": false,
  "publish_breadth_first": false
}
```

- `cp examples/configurations/drivers/master-driver.agent config/fake-master-driver.config`
- Add `fake.csv` and `fake.config` to the *configuration store*.
- Edit `fake-master-driver.config` to reflect paths on your system

fake-master-driver.config:

```
{
  "driver_scrape_interval": 0.05
}
```

- Create a script to simplify installation. The following will stop and remove any existing instances of agents create with the script, then package, install, and start the new instance. You will need to make the file executable: `chmod +x make-fakedriver`

make-fakedriver:

```
export SOURCE=services/core/MasterDriverAgent
export CONFIG=config/fake-master-driver.config
export TAG=fake-driver
./scripts/core/make-agent.sh
```

- If you have a Listener Agent already installed, you should start seeing data being published to the bus.

ListenerAgent

The ListenerAgent subscribes to all topics and is useful for testing that agents being developed are publishing correctly. It also provides a template for building other agents as it expresses the requirements of a platform agent.

Explanation of ListenerAgent

Use utils to setup logging which we'll use later.

```
utils.setup_logging()
_log = logging.getLogger(__name__)
```

The Listener agent extends (inherits from) the Agent class for its default functionality such as responding to platform commands:

```
class ListenerAgent(Agent):
    '''Listens to everything and publishes a heartbeat according to the
    heartbeat period specified in the settings module.'''
    '''
```

After the class definition, the Listener agent reads the configuration file, extracts the configuration parameters, and initializes any Listener agent instance variable. This is done the agents `init` method:

```
def __init__(self, config_path, **kwargs):
    super(ListenerAgent, self).__init__(**kwargs)
    self.config = utils.load_config(config_path)
    self._agent_id = self.config.get('agentid', DEFAULT_AGENTID)
    log_level = self.config.get('log-level', 'INFO')
    if log_level == 'ERROR':
        self._logfn = _log.error
    elif log_level == 'WARN':
        self._logfn = _log.warn
    elif log_level == 'DEBUG':
        self._logfn = _log.debug
    else:
        self._logfn = _log.info
```

Next, the Listener agent will run its setup method. This method is tagged to run after the agent is initialized by the decorator `@Core.receiver('onsetup')`. This method accesses the configuration parameters, logs a message to the platform log, and sets the agent ID.

```
@Core.receiver('onsetup')
def onsetup(self, sender, **kwargs):
    # Demonstrate accessing a value from the config file
    _log.info(self.config.get('message', DEFAULT_MESSAGE))
    self._agent_id = self.config.get('agentid')
```

The Listener agent subscribes to all topics published on the message bus. Subscribe/publish interactions with the message bus are handled by the PubSub module located at:

```
~/voltrtron/voltrtron/platform/vip/agent/subsystems/pubsub.py
```

The Listener agent uses an empty string to subscribe to all messages published. This is done in a [decorator](#) for simplifying subscriptions.

It also checks for the sender being `pubsub.compat` in case there are any VOLTRON 2.0 agents running on the platform.

```
@PubSub.subscribe('pubsub', '')
def on_match(self, peer, sender, bus, topic, headers, message):
    '''Use match_all to receive all messages and print them out.'''
    if sender == 'pubsub.compat':
        message = compat.unpack_legacy_message(headers, message)
    self._logfn(
        "Peer: %r, Sender: %r, Bus: %r, Topic: %r, Headers: %r, "
        "Message: %r", peer, sender, bus, topic, headers, message)
```

Node Red Example

Node Red is a visual programming wherein users connect small units of functionality “nodes” to create “flows”.

There are two example nodes that allow communication between Node Red and VOLTRON. One node reads subscribes to messages on the VOLTRON message bus and the other publishes to it.

Dependencies

The example nodes depend on *python-shell* to be installed and available to the Node Red environment.

Installation

Copy all files from *voltrtron/examples/NodeRed* to your *~/node-red/nodes* directory. *~/node-red* is the default directory for Node Red files. If you have set a different directory use that instead.

Set the variables at the beginning of the *voltrtron.js* file to be a valid VOLTRON environment, VOLTRON home, and python path.

Valid CURVE keys need to be added to the *settings.py* file. If they are generated with the *voltrtron-ctl auth keypair* command then the public key should be added to VOLTRON’s authorization file with the following:

```
$ voltrtron-ctl auth add
```

The serverkey can be found with

```
$ voltrtron-ctl auth serverkey
```

Usage

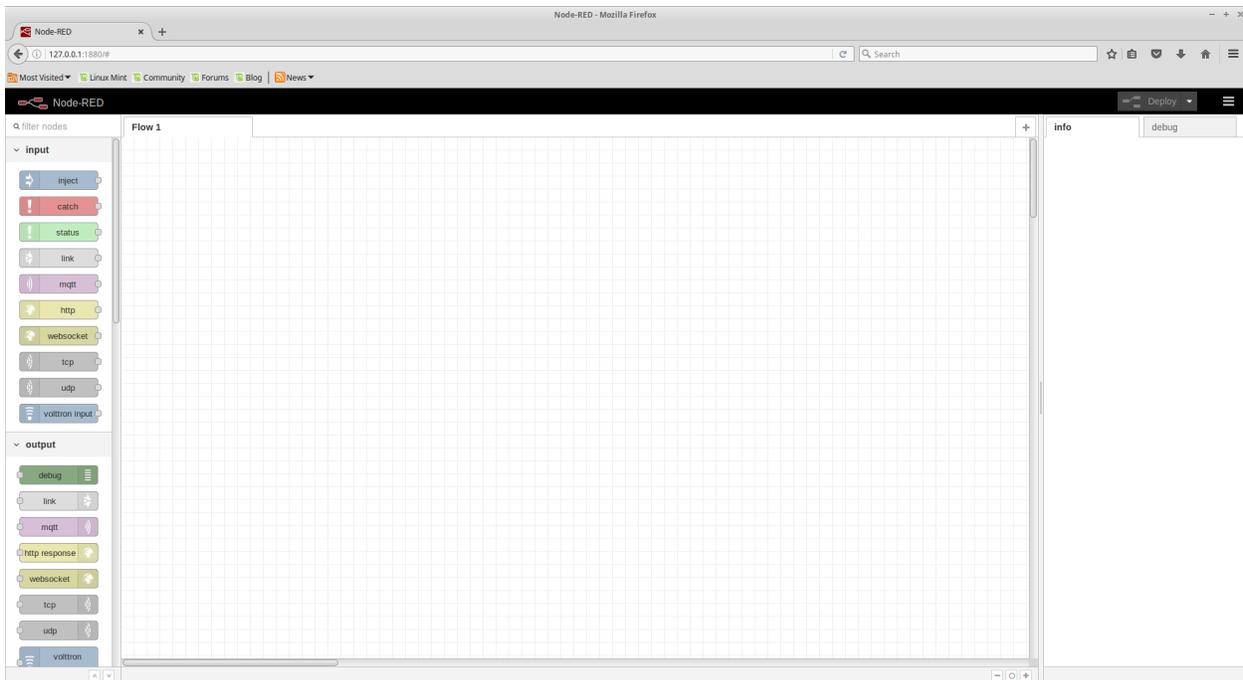
Start VOLTRON and Node Red.

```
$ node-red

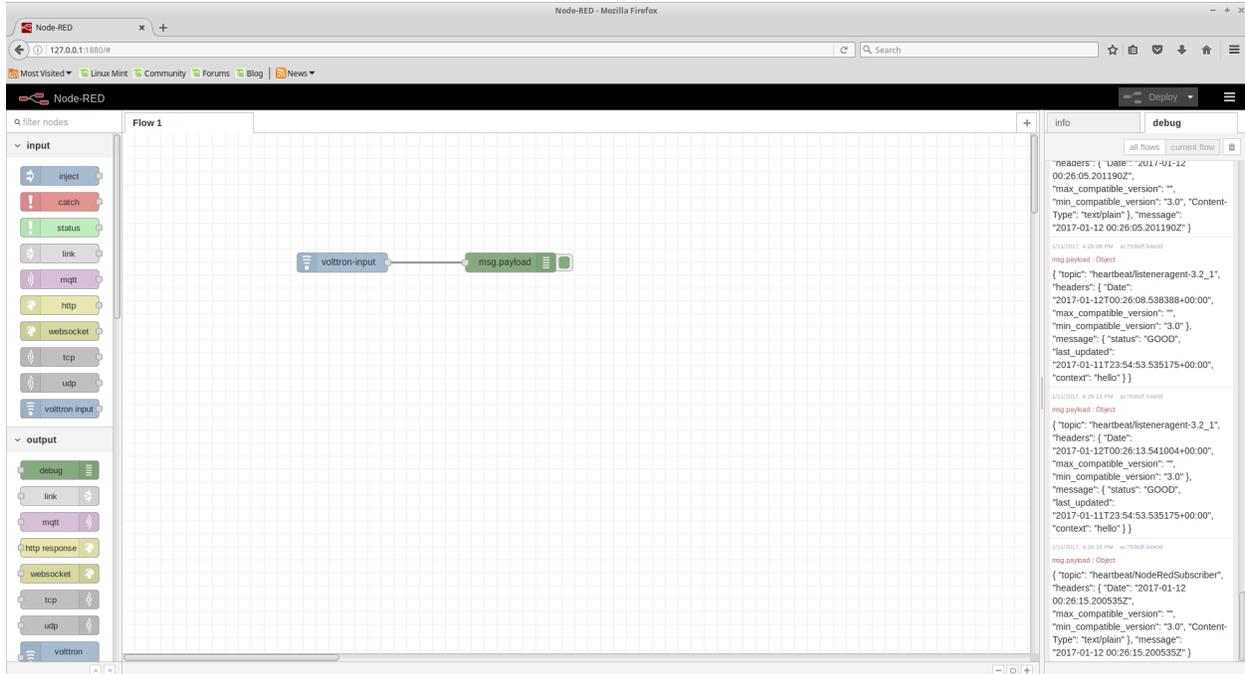
Welcome to Node-RED
=====

11 Jan 15:26:49 - [info] Node-RED version: v0.14.4
11 Jan 15:26:49 - [info] Node.js version: v0.10.25
11 Jan 15:26:49 - [info] Linux 3.16.0-38-generic x64 LE
11 Jan 15:26:49 - [info] Loading palette nodes
11 Jan 15:26:49 - [warn] -----
11 Jan 15:26:49 - [warn] [rpi-gpio] Info : Ignoring Raspberry Pi specific node
11 Jan 15:26:49 - [warn] -----
11 Jan 15:26:49 - [info] Settings file   : /home/volttron/.node-red/settings.js
11 Jan 15:26:49 - [info] User directory  : /home/volttron/.node-red
11 Jan 15:26:49 - [info] Flows file     : /home/volttron/.node-red/flows_volttron.json
11 Jan 15:26:49 - [info] Server now running at http://127.0.0.1:1880/
11 Jan 15:26:49 - [info] Starting flows
11 Jan 15:26:49 - [info] Started flows
```

The output from the Node Red command indicates the address of its web interface. Nodes available for use are in the left sidebar.



We can now use the VOLTTRON nodes to read from and write to VOLTTRON.



SchedulerExampleAgent

The SchedulerExampleAgent demonstrates how to use the scheduling feature of the [[ActuatorAgent]] as well as how to send a command. This agent publishes a request for a reservation on a (fake) device then takes an action when it's scheduled time appears. The ActuatorAgent must be running to exercise this example.

Note: Since there is no actual device, an error is produced when the agent attempts to take its action.

```
def publish_schedule(self):
    '''Periodically publish a schedule request'''
    headers = {
        'AgentID': agent_id,
        'type': 'NEW_SCHEDULE',
        'requesterID': agent_id, #The name of the requesting agent.
        'taskID': agent_id + "-ExampleTask", #The desired task ID for this task. It
        ↪ must be unique among all other scheduled tasks.
        'priority': 'LOW', #The desired task priority, must be 'HIGH', 'LOW', or 'LOW_
        ↪ PREEMPT'
    }

    start = str(datetime.datetime.now())
    end = str(datetime.datetime.now() + datetime.timedelta(minutes=1))

    msg = [
        ['campus/building/unit', start, end]
    ]
    self.vip.pubsub.publish(
        'pubsub', topics.ACTUATOR_SCHEDULE_REQUEST, headers, msg)
```

The agent listens to schedule announcements from the actuator and then issues a command

```
@PubSub.subscribe('pubsub', topics.ACTUATOR_SCHEDULE_ANNOUNCE (campus='campus',
                                                                building='building',unit='unit'))
def actuate(self, peer, sender, bus, topic, headers, message):
    print ("response:",topic,headers,message)
    if headers[headers_mod.REQUESTER_ID] != agent_id:
        return
    '''Match the announce for our fake device with our ID
    Then take an action. Note, this command will fail since there is no
    actual device'''
    headers = {
        'requesterID': agent_id,
    }
    self.vip.pubsub.publish(
        'pubsub', topics.ACTUATOR_SET (campus='campus',
                                       building='building',unit='unit',
                                       point='point'),
        headers, 0.0)
```

Utilities

BaseAgent

The BaseAgent class subscribes to required topics, responds to platform messages, and provides hooks for application-specific logic. While it is not required, it is recommended that agents extend this class to ease development.

DataPublisher

This is a simple agent that plays back data either from the config store or a CSV to the configured topic. It can also provide basic emulation of the actuator agent for testing agents that expect to be able to set points on a device in response to device publishes.

Installation notes

In order to simulate the actuator you must install the agent with the VIP identity of *platform.actuator*.

Configuration

```
{
    # basetopic can be devices, analysis, or custom base topic
    "basepath": "devices/PNNL/ISB1",

    # use_timestamp uses the included in the input_data if present.
    # Currently the column must be named `Timestamp`.
    "use_timestamp": true,

    # Only publish data at most once every max_data_frequency seconds.
    # Extra data is skipped.
    # The time windows are normalized from midnight.
    # ie 900 will publish one value for every 15 minute window starting from
    # midnight of when the agent was started.
```

(continues on next page)

(continued from previous page)

```

# Only used if timestamp in input file is used.
"max_data_frequency": 900,

# The meta data published with the device data is generated
# by matching point names to the unittype_map.
"unittype_map": {
  ".*Temperature": "Farenheit",
  ".*SetPoint": "Farenheit",
  "OutdoorDamperSignal": "On/Off",
  "SupplyFanStatus": "On/Off",
  "CoolingCall": "On/Off",
  "SupplyFanSpeed": "RPM",
  "Damper*": "On/Off",
  "Heating*": "On/Off",
  "DuctStatic*": "On/Off"
},
# Path to input CSV file.
# May also be a list of records or reference to a CSV file in the config store.
# Large CSV files should be referenced by file name and not
# stored in the config store.
"input_data": "econ_test2.csv",
# Publish interval in seconds
"publish_interval": 1,

# Tell the playback to maintain the location a the file in the config store.
# Playback will be resumed from this point
# at agent startup even if this setting is changed to false before restarting.
# Saves the current line in line_marker in the DataPublishers's config store
# as plain text.
# default false
"remember_playback": true,

# Start playback from 0 even if the line_marker configuration is set a non 0_
↪value.
# default false
"reset_playback": false,

# Repeat data from the start if this flag is true.
# Useful for data that does not include a timestamp and is played back in real_
↪time.
"replay_data": false
}

```

CSV File Format

The CSV file must have a single header line. The column names are appended to the *basepath* setting in the configuration file and the resulting topic is normalized to remove extra */'s*. The values are all treated as floating point values and converted accordingly.

The corresponding device for each point is determined and the values are combined together to create an *all* topic publish for each device.

If a *Timestamp* column is in the input it may be used to set the timestamp in the header of the published data.

Timestamp	centrifugal_chiller/OutsideAirTemperature	centrifugal_chiller/DischargeAirTemperatureSetPoint
2012/05/19 05:07:00	0	56
2012/05/19 05:08:00	0	56
2012/05/19 05:09:00	0	56
2012/05/19 05:10:00	0	56
2012/05/19 05:11:00	0	56
2012/05/19 05:12:00	0	56
2012/05/19 05:13:00	0	56
2012/05/19 05:14:00	0	56
2012/05/19 05:15:00	0	56
2012/05/19 05:16:00	0	56
2012/05/19 05:17:00	0	56
2012/05/19 05:18:00	0	56
2012/05/19 05:19:00	0	56
2012/05/19 05:20:00	0	56
2012/05/19 05:21:00	0	56
2012/05/19 05:22:00	0	56
2012/05/19 05:23:00	0	56
2012/05/19 05:24:00	0	56
2012/05/19 05:25:00	48.78	56
2012/05/19 05:26:00	48.88	56
2012/05/19 05:27:00	48.93	56
2012/05/19 05:28:00	48.95	56
2012/05/19 05:29:00	48.92	56
2012/05/19 05:30:00	48.88	56
2012/05/19 05:31:00	48.88	56
2012/05/19 05:32:00	48.99	56
2012/05/19 05:33:00	49.09	56
2012/05/19 05:34:00	49.11	56
2012/05/19 05:35:00	49.07	56
2012/05/19 05:36:00	49.05	56
2012/05/19 05:37:00	49.09	56
2012/05/19 05:38:00	49.13	56
2012/05/19 05:39:00	49.09	56
2012/05/19 05:40:00	49.01	56
2012/05/19 05:41:00	48.92	56
2012/05/19 05:42:00	48.86	56
2012/05/19 05:43:00	48.92	56
2012/05/19 05:44:00	48.95	56
2012/05/19 05:45:00	48.92	56
2012/05/19 05:46:00	48.86	56
2012/05/19 05:47:00	48.78	56
2012/05/19 05:48:00	48.69	56
2012/05/19 05:49:00	48.65	56
2012/05/19 05:50:00	48.65	56
2012/05/19 05:51:00	48.65	56
2012/05/19 05:52:00	48.61	56
2012/05/19 05:53:00	48.59	56
2012/05/19 05:54:00	48.55	56
2012/05/19 05:55:00	48.63	56

Timestamp	centrifugal_chiller/OutsideAirTemperature	centrifugal_chiller/DischargeAirTemperatureSetPoint
2012/05/19 05:56:00	48.76	56
2012/05/19 05:57:00	48.95	56
2012/05/19 05:58:00	49.24	56
2012/05/19 05:59:00	49.54	56
2012/05/19 06:00:00	49.71	56
2012/05/19 06:01:00	49.79	56
2012/05/19 06:02:00	49.94	56
2012/05/19 06:03:00	50.13	56
2012/05/19 06:04:00	50.18	56
2012/05/19 06:05:00	50.15	56

Driven Agents

Configuration for running openeis applications within voltron.

The configuration of an agent within voltron requires a small modification to the imports of the openeis application and a couple of configuration parameters.

Import and Extend

```
from voltron.platform.agent import (AbstractDrivenAgent, Results)
...
class OpeneisApp(AbstractDrivenAgent):
```

Configuration

The two parameters that are necessary in the json configuration file are “application” and “device”. An optional but recommended argument should also be added “agentid”.

```
{
  "agentid": "drivenlogger1",
  "application": "drivenlogger.logdevice.LogDevice",
  "device": "pnnl/isbl/oat",
  ...
}
```

Any other keys will be passed on to the openeis application when it is run.

Process Agent

This agent can be used to launch non-Python agents in the VOLTRON platform. The agent handles keeping track of the process so that it can be started and stopped with platform commands. Edit the configuration file to specify how to launch your process.

This agent was originally created for launching sMAP along with the platform, but can be used for any process.

Note: Currently this agent does not respond to a blanket “shutdown” request and must be stopped with the “stop” command.

Scripts

In order to make repetitive tasks less repetitive the VOLTTRON team has create several scripts in order to help. These tasks are available in the scripts directory. Before using these scripts you should become familiar with the *Agent Development* process.

In addition to the scripts directory, the VOLTTRON team has added the config directory to the .gitignore file. By convention this is where we store customized scripts and configuration that will not be made public. Please feel free to use this convention in your own processes.

The scripts/core directory is laid out in such a way that we can build scripts on top of a base core. For example the scripts in sub-folders such as the historian-scripts and demo-comms use the scripts that are present in the core directory.

The most widely used script is scripts/core/pack_install.sh. The pack_install.sh script will remove an agent if the tag is already present, create a new agent package, and install the agent to VOLTTRON_HOME. This script has three required arguments and has the following signature

```
# Agent to Package must have a setup.py in the root of the directory.
scripts/core/pack_install.sh <Agent to Package> <Config file> <Tag>
```

The pack_install.sh script will respect the VOLTTRON_HOME specified on the command line or set in the global environment. An example of setting VOLTTRON_HOME is as follows.

```
# Sets VOLTTRON_HOME to /tmp/vlhome
VOLTTRON_HOME=/tmp/vlhome scripts/core/pack_install.sh <Agent to Package> <Config_
↪file> <Tag>
```

Use the following scripts as examples that can be modified for your own agents.

- scripts/core/make-listener can be modified for any agent and make it one command to stop, remove, build, install, configure, tag, start, and (optionally) enable an agent for autostart. Fill out the script with the location of the agent source, config file, and tag name. The optional parameter enable can be passed to the make-agent script to set the agent to autostart with the platform.
- make-listener-enc-auth is similar to make-listener but uses *encryption and authentication*.

Applications

These resources summarize the use of the sample applications that are pre-packaged with VOLTTRON. For detailed information on these applications, refer to the report Transactional Network Platform: Applications available at http://www.pnl.gov/main/publications/external/technical_reports/PNNL-22941.pdf.

Note, as of VOLTTRON 4.0, applications are now in their own repository at: <https://github.com/VOLTTRON/volttron-applications>

Acquiring Third Party Agent Code

Add the volttron-applications repository under the volttron/applications directory by using following command:

```
git subtree add --prefix applications https://github.com/VOLTTRON/volttron-applications.git develop
--squash
```

Passive Automated Fault Detection and Diagnostic Agent

The Passive Automated Fault Detection and Diagnostic (Passive AFDD) agent is used to identify problems in the operation and performance of air-handling units (AHUs) or packaged rooftop units (RTUs). Air-side economizers modulate controllable dampers to use outside air to cool instead of (or to supplement) mechanical cooling, when outdoor-air conditions are more favorable than the return-air conditions. Unfortunately, economizers often do not work properly, leading to increased energy use rather than saving energy. Common problems include incorrect control strategies, diverse types of damper linkage and actuator failures, and out-of-calibration sensors. These problems can be detected using sensor data that is normally used to control the system.

The Passive AFDD requires the following data fields to perform the fault detection and diagnostics:

- Outside-air temperature
- Return-air temperature
- Mixed-air temperature
- Outside-air damper position/signal
- Supply fan status
- Mechanical cooling status
- Heating status.

The AFDD supports both real-time data via a Modbus or BACnet device, or input of data from a csv style text document.

The following section describes how to configure the Passive AFDD agent, methods for data input (real-time data from a device or historical data in a comma separated value formatted text file), and launching the Passive AFDD agent.

Note: A proactive version of the Passive AFDD exists as a PNNL application (AFDDAgent). This application requires active control of the RTU for fault detection and diagnostics to occur. The Passive AFDD was created to allow more users a chance to run diagnostics on their HVAC equipment without the need to actively modify the controls of the system.

Configuring the Passive AFDD Agent

Before launching the Passive AFDD agent, several parameters require configuration. The AFDD utilizes the same JSON style configuration file used by the Actuator, Listener, and Weather agents. The threshold parameters used for the fault detection algorithms are pre-configured and will work well for most RTUs or AHUs. Figure 1 shows an example configuration file for the AFDD agent.

The parameters boxed in black (in Figure 1) are the pre-configured fault detection thresholds; these do not require any modification to run the Passive AFDD agent. The parameters in the example configuration that are boxed in red will require user input. The following list describes each user configurable parameter and their possible values:

- **agentid** – This is the ID used when making schedule, set, or get requests to the Actuator agent; usually a string data type.
- **campus** – Campus name as configured in the sMAP driver. This parameter builds the device path that allows the Actuator agent to set and get values on the device; usually a string data type.
- **building** – Building name as configured in the sMAP driver. This parameter builds the device path that allows the Actuator agent to set and get values on the device; usually a string data type.
- **unit** – Device name as configured in the sMAP driver. This parameter builds the device path that allows the Actuator agent to set and get values on the device; usually a string data type. Note: The campus, building, and unit parameters are used to build the device path (campus/building/unit). The device path is used for communication on the message bus.

- **controller point names** – When using real-time communication, the Actuator agent identifies what registers or values to set or get by the point name you specify. This name must match the “Point Name” given in the Modbus registry file, as specified in VOLTTRON Core Services.
- **aggregate_data** – When using real-time data sampled at an interval of less than 1 hour or when inputting data via a csv file sampled at less than 1 hour intervals, set this flag to “1.” Value should be an integer or floating-point number (i.e., 1 or 1.0)
- **csv_input** – Flag to indicate if inputting data from a csv text file. Set to “0” for use with real-time data from a device or “1” if data is input from a csv text file. It should be an integer or floating point number (i.e., 1 or 1.0)
- **EER** – Energy efficiency ratio for the AHU or RTU. It should be an integer or floating-point number (i.e., 10 or 10.0)
- **tonnage** – Cooling capacity of the AHU or RTU in tons of cooling. It should be an integer or floating-point number (i.e., 10 or 10.0)
- **economizer_type** – This field indicates what type of economizer control is used. Set to “0” for differential dry-bulb control or to “1” for high limit dry-bulb control. It should be an integer or floating-point number.
- **high_limit** – If the economizer is using high-limit dry-bulb control, this value indicates what the outside-air temperature high limit should be. The input should be floating-point number (i.e., 60.0)
- **matemp_missing** – Flag used to indicate if the mixed-air temperature is missing for this system. If utilizing csv data input, simply set this flag to “1” and replace the mixed-air temperature column with discharge-air temperature data. If using real-time data input, change the field “mat_point_name” under Point Names section to the point name indicating the discharge-air temperature. It should be an integer or floating-point number (i.e., 1 or 1.0)
- **OAE6** – This section contains the schedule information for the AHU or RTU. The default is to indicate a 24-hour schedule for each day of the week. To modify this, change the numbers in the bracketed list next to the corresponding day with which you are making operation schedule modifications. For example: “Saturday”: [0,0] (This indicates the system is off on Saturdays).

```

{
  "agent": {
    "exec": "passiveafdd-0.1-py2.7.egg --config \"%c\" --sub \"%s\" --pub \"%p\""}
  },
  "agentid": "afdd1",
  "campus": "campus1",
  "building": "building1",
  "unit": "device1",
  "smap_path": "datalogger/log/afdd1/campus1/building1/device1", #/datalogger/log/your sMAP path here

  #[Controller point names]
  "oat_point_name": "OutsideAirTemp",
  "mat_point_name": "MixedAirTemp", #"DischargeAirTemp"
  "dat_point_name": "DischargeAirTemperature",
  "rat_point_name": "ReturnAirTemp",
  "damper_point_name": "Damper",
  "cool_call1_point_name": "CoolCall",
  "cool_cmd1_point_name": "CompressorStatus",
  "fan_status_point_name": "FanStatus",
  "heat_command1_point_name": "Heating",

  #[Input Variables]
  "aggregate_data": 1,
  "csv_input": 1,
  "EER": 10,
  "tonnage": 10
  "high_limit": 70,
  "economizer_type": 0,
  "matemp_missing": 0,

  #[oaf]
  "oaf_temp_threshold": 4.0,

  #[OAE1]
  "mat_low": 50,
  "mat_high": 90,
  "rat_low": 50,
  "rat_high": 90,
  "oat_low": 30,
  "oat_high": 120,

  #[OAE2]
  "oae2_damper_threshold": 30.0,
  "oae2_oaf_threshold": 0.25,

  #[OAE3]
  "damper_minimum": 20,

  #[OAE4]
  "minimum_oa": 0.1,
  "oae4_oaf_threshold": 0.25,

  #[OAE5]
  "oae5_oaf_threshold": 0.0,

  #[OAE6]
  "Sunday": [0,23], #this schedule is 24 hours
  "Monday": [0,23],
  "Tuesday": [0,23],
  "Wednesday": [0,23],
  "Thursday": [0,23],
  "Friday": [0,23],
  "Saturday": [0,23],
}

```

Figure 1. Example Passive AFDD Agent Configuration File

Launching the Passive AFDD Agent

The Passive AFDD agent performs passive diagnostics on AHUs or RTUs, monitors and utilizes sensor data, but does not actively control the devices. Therefore, the agent does not require interaction with the Actuator agent. Steps for launching the agent are provided below.

In a terminal window, enter the following commands:

1. Run `pack_install` script on Passive AFDD agent:

```
$ . scripts/core/pack_install.sh applications/PassiveAFDD applications/PassiveAFDD/
↳passiveafdd.launch.json passive-afdd
```

Upon successful completion of this command, the terminal output will show the install directory, the agent UUID (unique identifier for an agent; the UUID shown in red is only an example and each instance of an agent will have a different UUID), and the agent name (blue text):

```
Installed /home/volttron-user/.volttron/packaged/passiveafdd-0.1-py2-none-any.whl as_
↳5df00517-6a4e-4283-8c70-5f0759713c64 passiveafdd-0.1
```

2. Start the agent:

```
$ volttron-ctl start --tag passive-afdd
```

3. Verify that the agent is running:

```
$ volttron-ctl status
$ tail -f volttron.log
```

If changes are made to the Passive AFDD agent’s configuration file after the agent is launched, it is necessary to stop and reload the agent. In a terminal, enter the following commands:

```
$ volttron-ctl stop --tag passive-afdd
$ volttron-ctl remove --tag passive-afdd
```

Then re-build and start the updated agent.

When the AFDD agent is monitoring a device via the message bus, the agent relies on the periodic data published from the sMAP driver. The AFDD agent then aggregates this data each hour and performs the diagnostics on the average hourly data. The result is written to a csv text file, which is appended if the file already exists. This file is in a folder titled “Results” under the (`<project directory>/applications/pnml/PassiveAFDD/passiveafdd`) directory. Below is a key that describes how to interpret the diagnostic results:

Diagnostic Code	Code Message
AFDD-1 (Temperature Sensor Fault)	
20	No faults detected
21	Temperature sensor fault
22	Conditions not favorable for diagnostic
23	Mixed-air temperature outside of expected range
24	Return-air temperature outside of expected range
25	Outside-air temperature outside of expected range
27	Missing data necessary for fault detection
29	Unit is off (No Fault)
AFDD-2 (RTU Economizing When it Should)	
30	No faults detected

Continue

Table 3 – continued from previous page

Diagnostic Code	Code Message
31	Unit is not currently cooling or conditions are not favorable for economizing
32	Insufficient outdoor air when economizing (Fault)
33	Outdoor-air damper is not fully open when the unit should be economizing
36	OAD is open but conditions were not favorable for OAF calculation (No Fault)
37	Missing data necessary for fault detection (No Fault)
38	OAD is open when economizing but OAF calculation led to an unexpected
39	Unit is off (No Fault)
AFDD-3 (Unit Economizing When it Should)	
40	No faults detected
41	Damper should be at minimum position but is not (Fault)
42	Damper is at minimum for ventilation (No Fault)
43	Conditions favorable for economizing (No Fault)
47	Missing data necessary for fault detection (No Fault)
49	Unit is off (No Fault)
AFDD-4 (Excess Outdoor-air Intake)	
50	No faults detected
51	Excessive outdoor-air intake
52	Damper is at minimum but conditions are not favorable for OAF calculation
53	Damper is not at minimum (Fault)
56	Unit should be economizing (No Fault)
57	Missing data necessary for fault detection (No Fault)
58	Damper is at minimum but OAF calculation led to an unexpected value (No
59	Unit is off (No Fault)
AFDD-5 (Insufficient Outdoor-air Ventilation)	
60	No faults detected
61	Insufficient outdoor-air intake (Fault)
62	Damper is at minimum but conditions are not favorable for OAF calculation
63	Damper is not at minimum when is should not be (Fault)
66	Unit should be economizing (No Fault)
67	Missing data necessary for fault detection (No Fault)
68	Damper is at minimum but conditions are not favorable for OAF calculation
69	Unit is off (No Fault)
AFDD-6 (Schedule)	
70	Unit is operating correctly based on input on/off time (No Fault)
71	Unit is operating at a time designated in schedule as “off” time
77	Missing data

Launching the AFDD for CSV Data Input

When utilizing the AFDD agent and inputting data via a csv text file, set the **csv_input** parameter, contained in the AFDD configuration file, to “1.”

- Launch the agent normally.
- A small file input box will appear. Navigate to the csv data file and select the csv file to input for the diagnostic.
- The result will be created for this RTU or AHU in the results folder described.

Figure 2 shows the dialog box that is used to input the csv data file.

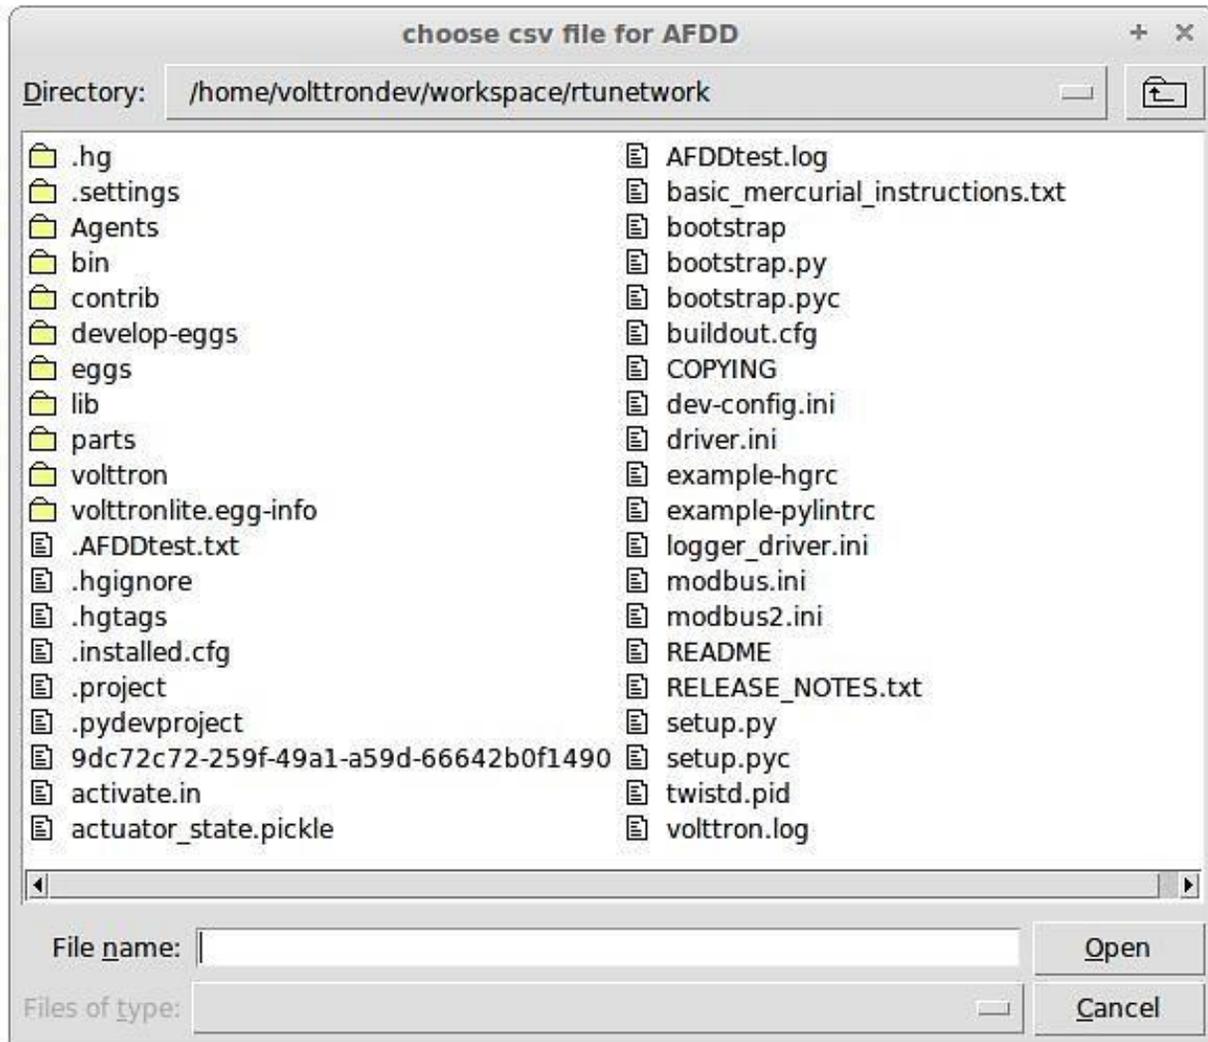


Figure 2 File Selection Dialog Box when Inputting Data in a csv File

If “Cancel” is pushed on the file input dialog box, the AFDD will acknowledge that no file was selected. The Passive AFDD must be restarted to run the diagnostics. If a non-csv file is selected, the AFDD will acknowledge the file selected was not a csv file. The AFDD must be restarted to run the diagnostics.

Figure 3 shows a sample input data in a csv format. The header, or name for each column from the data input csv file used for analysis, should match the name given in the configuration file, as shown in Figure 1, boxed in red.

```
Timestamp,OutsideAirTemp,ReturnAirTemp,MixedAirTemp,CompressorStatus,HeatingStatus,FanStatus,Damper
5/19/2012 6:00,48.902,56.43727273,58.68472222,0,0,0,0
5/19/2012 7:00,51.12316667,59.47933333,59.58916667,0,0,0,0
5/19/2012 8:00,54.70866667,61.1625,64.34266667,0,0,0,0
```

Figure 3 Sample of CSV Data for Passive AFDD Agent

The Demand Response (DR) Agent

Many utilities around the country have or are considering implementing dynamic electrical pricing programs that use time-of-use (TOU) electrical rates. TOU electrical rates vary based on the demand for electricity. Critical peak pricing (CPP), also referred to as critical peak days or event days, is an electrical rate where utilities charge an increased price

above normal pricing for peak hours on the CPP day. CPP times coincide with peak demand on the utility; these CPP events are generally called between 5 to 15 times per year and occur when the electrical demand is high and the supply is low. Customers on a flat standard rate who enroll in a peak time rebate program receive rebates for using less electricity when a utility calls for a peak time event. Most CPP events occur during the summer season on very hot days. The initial implementation of the DR agent addresses CPP events where the RTU would normally be cooling. This implementation can be extended to handle CPP events for heating during the winter season as well. This implementation of the DR agent is specific to the CPP, but it can easily be modified to work with other incentive signals (real-time pricing, day ahead, etc.).

The main goal of the building owner/operator is to minimize the electricity consumption during peak summer periods on a CPP day. To accomplish that goal, the DR agent performs three distinct functions:

- **Step 1 – Pre-Cooling:** Prior to the CPP event period, the cooling and heating (to ensure the RTU is not driven into a heating mode) set points are reset lower to allow for pre-cooling. This step allows the RTU to cool the building below its normal cooling set point while the electrical rates are still low (compared to CPP events). The cooling set point is typically lowered between 3 and 5°F below the normal. Rather than change the set point to a value that is 3 to 5°F below the normal all at once, the set point is gradually lowered over a period of time.
- **Step 2 – Event:** During the CPP event, the cooling set point is raised to a value that is 4 to 5°F above the normal, the damper is commanded to a position that is slightly below the normal minimum (half of the normal minimum), the fan speed is slightly reduced (by 10% to 20% of the normal speed, if the unit has a variable-frequency drive (VFD)), and the second stage cooling differential (time delay between stage one and stage two cooling) is increased (by few degrees, if the unit has multiple stages). The modifications to the normal set points during the CPP event for the fan speed, minimum damper position, cooling set point, and second stage cooling differential are user adjustable. These steps will reduce the electrical consumption during the CPP event. The pre-cooling actions taken in step 1 will allow the temperature to slowly float up to the CPP cooling temperature set point and reduce occupant discomfort during the attempt to shed load.
- **Step 3 – Post-Event:** The DR agent will begin to return the RTU to normal operations by changing the cooling and heating set points to their normal values. Again, rather than changing the set point in one step, the set point is changed gradually over a period of time to avoid the “rebound” effect (a spike in energy consumption after the CPP event when RTU operations are returning to normal).

The following section will detail how to configure and launch the DR agent.

Configuring DR Agent

Before launching the DR agent, several parameters require configuration. The DR utilizes the same JSON style configuration file that the Actuator, Listener, and Weather agent use. A notable limitation of the DR agent is that the DR agent requires active control of an RTU/AHU. The DR agent modifies set points on the controller or thermostat to reduce electrical consumption during a CPP event. The DR agent must be able to **set** certain values on the RTU/AHU controller or thermostat via the Actuator agent. Figure 4 shows a sample configuration file for the DR agent:

```
{
  "agent": {
    "exec": "DemandResponseagent-0.1-py2.7.egg --config \"%c\" --sub \"%s\" --pub \"%p\""}
  },
}
```

```
#Agent Parameters
"agentid": "DRAGENT1", #Agent ID used by actuator agent for control of RTU

"campus": "campus", #campus name as known by Volttron

"building": "building", #Building name as known by Volttron

"unit": "device", #RTU/Controller name as known by Volttron

"smmap_path": "datalogger/log/testing/campus/device" , #/datalogger/log/your path here
```

```
#Catalyst Controller point names
"cooling_stpt": "CoolingTemperatureStPt", # second value in quotes in name from your controller

"heating_stpt": "HeatingTemperatureStPt",

"min_damper_stpt": "MinimumDamperPositionStPt",

"cooling_stage_diff": "CoolingStageDifferential",

"cooling_fan_sp1": "CoolSupplyFanSpeed1",

"cooling_fan_sp2": "CoolSupplyFanSpeed2",

"override_command": "VoltronPBStatus",

"occupied_status": "Occupied",

"space_temp": "SpaceTemp",

"volttron_flag": "VoltronFlag",
```

```

#DR cooling Set Points
"csp_pre": 65.0,    #Pre-cooling zone temperature set point

"csp_cpp": 80.0,    #CPP event zone temperature set point

#Normal set points
"normal_firststage_fanspeed": 90.0,

"normal_secondstage_fanspeed": 90.0,

"normal_damper_stpt": 5.0,

"normal_coolingstpt": 74.0,

"normal_heatingstpt": 67.0,

#DR Parameters
"fan_reduction": 0.1, #fractional reduction 10% = 0.1

"damper_cpp": 0, #minimum damper command during CPP event

"timestep_length": 900, #number of seconds between CSP modifications in Pre and After event (default 900 sec. = 15 min.)

"max_precool_hours": 5, #maximum pre-cooling window in hours

"building_thermal_constant": 4.0, #Building thermal constant F/hr

"cooling_stage_differential": 1.0,

"Schedule": [1,1,1,1,1,1,1] #[Mon, Tue, Wed, Thu, Fri, Sat, Sun]
}

```

Figure 4 Example Configuration File for the DR Agent

The parameters boxed in black (Figure 4) are the demand response parameters; these may require modification to ensure the DR agent and corresponding CPP event are executed as one desires. The parameters in the example configuration that are boxed in red are the controller or thermostat points, as specified in the Modbus or BACnet (depending on what communication protocol your device uses) registry file, that the DR agent will set via the Actuator agent. These device points must be writeable, and configured as such, in the registry (Modbus or BACnet) file. The following list describes each user configurable parameter:

- **agentid** - This is the ID used when making schedule, set, or get requests to the Actuator agent; usually a string data type.
- **campus** - Campus name as configured in the sMAP driver. This parameter builds the device path that allows the Actuator agent to set and get values on the device; usually a string data type.
- **building** - Building name as configured in the sMAP driver. This parameter builds the device path that allows the Actuator agent to set and get values on the device; usually a string data type.
- **unit** - Device name as configured in the sMAP driver. This parameter builds the device path that allows the Actuator agent to set and get values on the device; usually a string data type. Note: The campus, building, and unit parameters are used to build the device path (campus/building/unit). The device path is used for communication on the message bus.
- **csp_pre** - Pre-cooling space cooling temperature set point.
- **csp_cpp** - CPP event space cooling temperature set point.
- **normal_firststage_fanspeed** - Normal operations, first stage fan speed set point.
- **normal_secondstage_fanspeed** - Normal operations, second stage fan speed set point.
- **normal_damper_stpt** - Normal operations, minimum outdoor-air damper set point.
- **normal_coolingstpt** - Normal operations, space cooling temperature set point.
- **normal_heatingstpt** - Normal operations, space heating temperature set point.

- **fan_reduction** - Fractional reduction in fan speeds during CPP event (default: 0.1-10%).
- **damper_cpp** - CPP event, minimum outdoor-air damper set point.
- **max_precool_hours** - Maximum allotted time for pre-cooling, in hours.
- **cooling_stage_differential** - Difference in actual space temperature and set-point temperature before second stage cooling is activated.
- **schedule** - Day of week occupancy schedule “0” indicate unoccupied day and “1” indicate occupied day (e.g., [1,1,1,1,1,1,1] = [Mon, Tue, Wed, Thu, Fri, Sat, Sun]).

OpenADR (Open Automated Demand Response)

Open Automated Demand Response (OpenADR) is an open and standardized way for electricity providers and system operators to communicate DR signals with each other and with their customers using a common language over any existing IP-based communications network, such as the Internet. Lawrence Berkeley National Laboratory created an agent to receive DR signals from an external source (e.g., OpenADR server) and publish this information on the message bus. The DR agent subscribes to the OpenADR topic and utilizes the contents of this message to coordinate the CPP event.

The OpenADR signal is formatted as follows:

```
'openadr/event', {'Content-Type': ['application/json'], 'requesterID': 'openadragent'},  
→ {'status': 'near',  
'start_at': '2013-6-15 14:00:00', 'end_at': '2013-10-15 18:00:00', 'mod_num': 0, 'id':  
'18455630-a5c4-4e4a-9d53-b3cf989ccf1b', 'signals': 'null'}
```

The red text in the signal is the topic associated with CPP events that are published on the message bus. The text in dark blue is the message; this contains the relevant information on the CPP event for use by the DR agent.

If one desires to test the behavior of a device when responding to a DR event, such an event can be simulated by manually publishing a DR signal on the message bus. From the base VOLTTRON directory, in a terminal window, enter the following commands:

1. Activate project:

```
$ source env/bin/activate
```

2. Start Python interpreter:

```
$ python
```

3. Import VOLTTRON modules:

```
$ from volttron.platform.vip.agent import Core, Agent
```

4. Import needed Python library:

```
$ import gevent
```

5. Instantiate agent (agent will publish OpenADR message):

```
$ agent = Agent(address='ipc://@/home/volttron-user/.volttron/run/vip.socket')
```

6. Ensure the setup portion of the agent run loop is executed:

```
$ gevent.spawn(agent.core.run).join(0)
```

7. Publish simulated OpenADR message:

```
$ agent.vip.pubsub.publish(peer='pubsub', topic='openadr/event',headers={},
message={'id': 'event_id','status': 'active', 'start_at': 10-30-15 15:00', 'end_at':
↪'10-30-15
18:00'})
```

To cancel this event, enter the following command:

```
$ agent.vip.pubsub.publish(peer='pubsub', topic='openadr/event',headers={}, message={
↪'id':
'event_id','status': 'cancelled', 'start_at': 10-30-15 15:00', 'end_at': '10-30-15_
↪18:00'})
```

The DR agent will use the most current signal for a given day. This allows utilities/OpenADR to modify the signal up to the time prescribed for pre-cooling.

DR Agent Output to sMAP

After the DR agent has been configured, the agent can be launched. To launch the DR agent from the base VOLTRON directory, enter the following commands in a terminal window:

1. Run *pack_install* script on DR agent:

```
$ . scripts/core/pack_install.sh applications/DemandResponseAgent
applications/DemandResponseAgent/demandresponse.launch.json dr-agent
```

Upon successful completion of this command, the terminal output will show the install directory, the agent UUID (unique identifier for an agent; the UUID shown in red is only an example and each instance of an agent will have a different UUID) and the agent name (blue text):

```
Installed
/home/volttron-user/.volttron/packaged/DemandResponseagent-0.1-py2-none-
any.whlas 5b1706d6-b71d-4045-86a3-8be5c85ce801
DemandResponseagent-0.1
```

2. Start the agent:

```
$ volttron-ctl start --tag dr-agent
```

3. Verify that agent is running:

```
$ volttron-ctl status
$ tail -f volttron.log
```

If changes are made to the DR agent's configuration file after the agent is launched, it is necessary to stop and reload the agent. In a terminal, enter the following commands:

```
$ volttron-ctl stop --tag dr-agent
$ volttron-ctl remove --tag dr-agent
```

Then re-build and start the updated agent.

Jupyter Notebooks

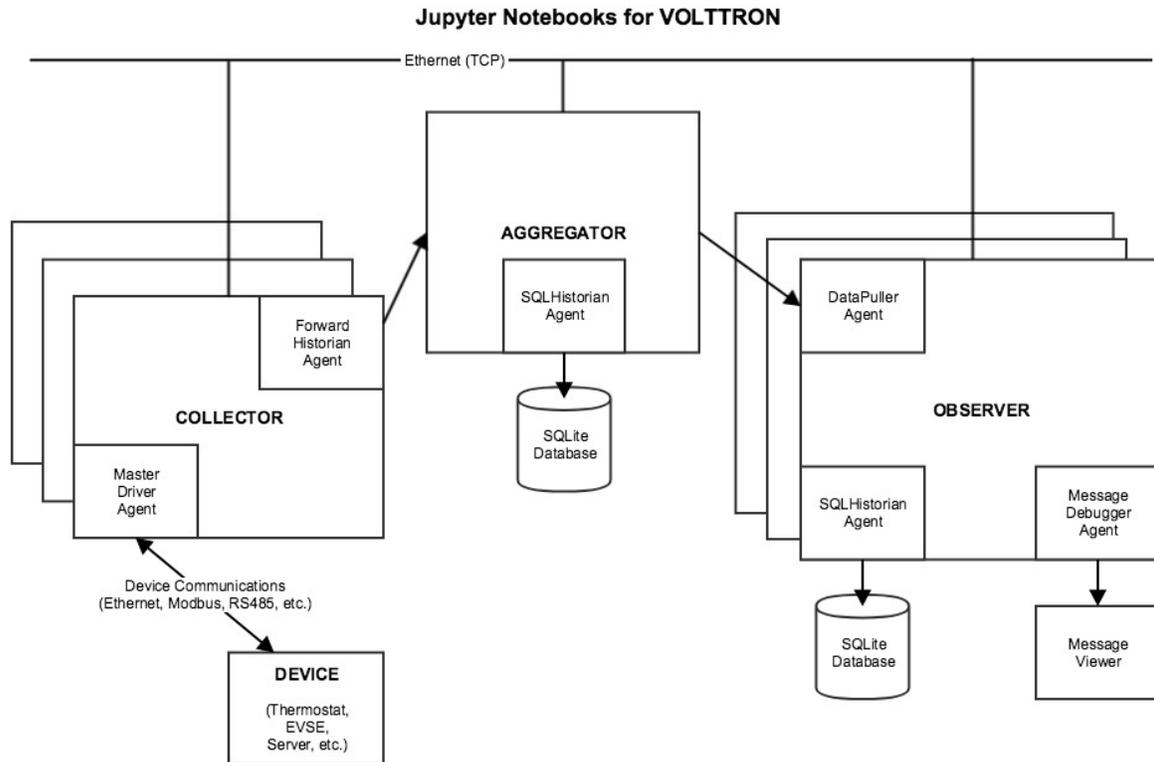
Jupyter is an open-source web application that lets you create and share “notebook” documents. A notebook displays formatted text along with live code that can be executed from the browser, displaying the execution output and preserving it in the document. Notebooks that execute Python code used to be called *iPython Notebooks*. The iPython Notebook project has now merged into Project Jupyter.

Using Jupyter to Manage a Set of VOLTTRON Servers

The following Jupyter notebooks for VOLTTRON have been provided as examples:

- **Collector notebooks.** Each Collector notebook sets up a particular type of device driver and forwards device data to another VOLTTRON instance, the Aggregator.
 - **SimulationCollector notebook.** This notebook sets up a group of Simulation device drivers and forwards device data to another VOLTTRON instance, the Aggregator.
 - **BacnetCollector notebook.** This notebook sets up a Bacnet (or Bacnet gateway) device driver and forwards device data to another VOLTTRON instance, the Aggregator.
 - **ChargePointCollector notebook.** This notebook sets up a ChargePoint device driver and forwards device data to another VOLTTRON instance, the Aggregator.
 - **SEP2Collector notebook.** This notebook sets up a SEP2.0 (IEEE 2030.5) device driver and forwards device data to another VOLTTRON instance, the Aggregator. The Smart Energy Profile 2.0 (“SEP2”) protocol implements IEEE 2030.5, and is capable of connecting a wide array of smart energy devices to the Smart Grid. The standard is designed to run over TCP/IP and is physical layer agnostic.
- **Aggregator notebook.** This notebook sets up and executes aggregation of forwarded data from other VOLTTRON instances, using a historian to record the data.
- **Observer notebook.** This notebook sets up and executes a DataPuller that captures data from another VOLTTRON instance, using a Historian to record the data. It also uses the Message Debugger agent to monitor messages flowing across the VOLTTRON bus.

Each notebook configures and runs a set of VOLTTRON Agents. When used as a set, they implement a multiple-VOLTTRON-instance architecture that captures remote device data, aggregates it, and reports on it, routing the data as follows:



Install VOLTRON and Jupyter on a Server

The remainder of this guide describes how to set up a host for VOLTRON and Jupyter. Use this setup process on a server in order to prepare it to run Jupyter notebook for VOLTRON.

Set Up the Server and Install VOLTRON

The following is a complete, but terse, description of the steps for installing and running VOLTRON on a server. For more detailed, general instructions, see [Installing Volttron](#).

The VOLTRON server should run on the same host as the Jupyter server.

Load third-party software:

```
$ sudo apt-get update
$ sudo apt-get install build-essential python-dev openssl libssl-dev libevent-dev git
$ sudo apt-get install sqlite3
```

Clone the VOLTRON repository from github:

```
$ cd ~
$ mkdir repos
$ cd repos
$ git clone https://github.com/VOLTRON/volttron/
```

Check out the develop (or master) branch and bootstrap the development environment:

```
$ cd volttron
$ git checkout develop
$ python2.7 bootstrap.py
```

Activate and initialize the VOLTTRON virtual environment:

Run the following each time you open a new command-line shell on the server:

```
$ export VOLTTRON_ROOT=~/.repos/volttron
$ export VOLTTRON_HOME=~/.volttron
$ cd $VOLTTRON_ROOT
$ source env/bin/activate
```

Install Extra Libraries

Add Python libraries to the VOLTTRON virtual environment:

These notebooks use third-party software that's not included in VOLTTRON's standard distribution that was loaded by `bootstrap.py`. The following additional packages are required:

- Jupyter
- SQLAlchemy (for the Message Debugger)
- Suds (for the ChargePoint driver)
- Numpy and Matplotlib (for plotted output)

Note: A Jupyter installation also installs and/or upgrades many dependent libraries. Doing so could disrupt other work on the OS, so it's safest to load Jupyter (and any other library code) in a virtual environment. VOLTTRON runs in a virtual environment anyway, so if you're using Jupyter in conjunction with VOLTTRON, it should be installed in your VOLTTRON virtual environment. (In other words, be sure to use `cd $VOLTTRON_ROOT` and `source env/bin/activate` to activate the virtual environment before running `pip install`.)

Install the third-party software:

```
$ pip install SQLAlchemy==1.1.4
$ pip install suds-jurko==0.6
$ pip install numpy
$ pip install matplotlib
$ pip install jupyter
```

Note: If `pip install` fails due to an untrusted cert, try using this command instead:

```
$ pip install --trusted-host pypi.python.org <libraryname>
```

(An `InsecurePlatformWarning` may be displayed, but it typically won't stop the installation from proceeding.)

Configure VOLTTRON

Use the `volttron-cfg` wizard to configure the VOLTTRON instance. By default, the wizard configures a VOLTTRON instance that communicates with agents only on the local host (ip 127.0.0.1). This set of notebooks manages communications among multiple VOLTTRON instances on different hosts. To enable this cross-host communication on VOLTTRON's web server, replace 127.0.0.1 with the host's IP address, as follows:

```
$ volttron-cfg
```

- Accept all defaults, except as follows.
- If a prompt defaults to 127.0.0.1 as an IP address, substitute the host's IP address (this may happen multiple times).

- When asked whether this is a volttron central, answer `Y`.
- When prompted for a username and password, use `admin` and `admin`.

Start VOLTRON

Start the main VOLTRON process, logging to `$VOLTRON_ROOT/volttron.log`:

```
$ volttron -vv -l volttron.log --msgdebug
```

This runs VOLTRON as a foreground process. To run it in the background, use:

```
$ volttron -vv -l volttron.log --msgdebug &
```

This also enables the Message Debugger, a non-production VOLTRON debugging aid that's used by some notebooks. To run with the Message Debugger disabled (VOLTRON's normal state), omit the `--msgdebug` flag.

Now that VOLTRON is running, it's ready for agent configuration and execution. Each Jupyter notebook contains detailed instructions and executable code for doing that.

Configure Jupyter

More detailed information about installing, configuring and using Jupyter Notebooks is available on the Project Jupyter site, <http://jupyter.org/>.

Create a Jupyter configuration file:

```
$ jupyter notebook --generate-config
```

Revise the Jupyter configuration:

Open `~/ .jupyter/jupyter_notebook_config.py` in your favorite text editor. Change the configuration to accept connections from any IP address (not just from localhost) and use a specific, non-default port number:

- Un-comment `c.NotebookApp.ip` and set it to: `'*'` instead of `'localhost'`
- Un-comment `c.NotebookApp.port` and set it to: `'8891'` instead of `'8888'`

Save the config file.

Open ports for TCP connections:

Make sure that your Jupyter server host's security rules allow inbound TCP connections on port 8891.

If the VOLTRON instance needs to receive TCP requests, for example ForwardHistorian or DataPuller messages from other VOLTRON instances, make sure that the host's security rules also allow inbound TCP communications on VOLTRON's port, which is usually 22916.

Launch Jupyter

Start the Jupyter server:

In a separate command-line shell, set up VOLTRON's environment variables and virtual environment, and then launch the Jupyter server:

```
$ export VOLTRON_HOME=(your volttron home directory, e.g. ~/.volttron)
$ export VOLTRON_ROOT=(where volttron was installed; e.g. ~/repos/volttron)
$ cd $VOLTRON_ROOT
$ source env/bin/activate
$ cd examples/JupyterNotebooks
$ jupyter notebook --no-browser
```

Open a Jupyter client in a web browser:

Look up the host's IP address (e.g., using `ifconfig`). Open a web browser and navigate to the URL that was displayed when you started jupyter, replacing `localhost` with that IP address. A Jupyter web page should display, listing your notebooks.

Python for Matlab Users

Matlab is a popular, proprietary programming language and tool suite with built in support for matrix operations and graphically plotting computation results. The purpose of this document is to introduce Python to those already familiar Matlab so it will be easier for them to develop tools and agents in VOLTTRON.

A Simple Function

Python and Matlab are similar in many respects, syntactically and semantically. With the addition of the NumPy library in Python, almost all numerical operations in Matlab can be emulated or directly translated. Here are functions in each language that perform the same operation:

```
% Matlab
function [result] = times_two(number)
    result = number * 2;
end
```

```
# Python
def times_two(number):
    result = number * 2
    return result
```

Some notes about the previous functions:

1. Values are explicitly returned with the *return* statement. It is possible to return multiple values, as in Matlab, but doing this without a good reason can lead to overcomplicated functions.
2. Semicolons are not used to end statements in python, and white space is significant. After a block is started (if, for, while, functions, classes) subsequent lines should be indented with four spaces. The block ends when the programmer stops adding the extra level of indentation.

Translating

The following may be helpful if you already have a Matlab file or function that will be translated into Python. Many of the syntactic differences between Matlab and Python can be rectified with your text editor's find and replace feature.

Start by copying all of your Matlab code into a new file with a `.py` extension. I recommend commenting everything out and uncommenting the Matlab code in chunks. This way you can write valid Python and verify it as you translate, instead of waiting till the whole file is "translated". Editors designed to work with Python should be able to highlight syntax errors for you as well.

1. Comments are created with a `%`. Find and replace these with `#`.
2. Change *elseif* blocks to *elif* blocks.
3. Python indexes start at zero instead of one. Array slices and range operations, however, don't include the upper bound, so only the lower bound should decrease by one.
4. Semicolons in Matlab are used to suppress output at the end of lines and for organizing array literals. After arranging the arrays into nested lists, all semicolons can be removed.

5. The `end` keyword in Matlab is used both to access the last element in an array and to close blocks. The array use case can be replaced with `-1` and the others can be removed entirely.

A More Concrete Example

In the [Building Economic Dispatch](#) project, a sibling project to VOLTRON, a number of components written in Matlab would create a matrix out of some collection of columns and perform least squares regression using the *matrix division* operator. This is straightforward and very similar in both languages so long as all of the columns are defined and are the same length.

```
% Matlab
XX = [U, xbp, xbp2, xbp3, xbp4, xbp5];
AA = XX \ ybp;
```

```
# Python
import numpy as np

XX = np.column_stack((U, xbp, xbp2, xbp3, xbp4, xbp5))
AA, resid, rank, s = np.linalg.lstsq(XX, ybp)
```

This pattern also included the creation of the U column, a column of ones used as the bias term in the linear equation. In order to make the Python version more readable and more robust, the pattern was removed from each component and replaced with a single function call to `least_squares_regression`.

This function does some validation on the input parameters, automatically creates the bias column, and returns the least squares solution to the system. Now if we want to change how the solution is calculated we only have to change the one function, instead of each instance where the pattern was written originally.

```
def least_squares_regression(inputs=None, output=None):
    if inputs is None:
        raise ValueError("At least one input column is required")
    if output is None:
        raise ValueError("Output column is required")

    if type(inputs) != tuple:
        inputs = (inputs,)

    ones = np.ones(len(inputs[0]))
    x_columns = np.column_stack((ones,) + inputs)

    solution, resid, rank, s = np.linalg.lstsq(x_columns, output)
    return solution
```

Lessons Learned (sometimes the hard way)

Variable Names

Use descriptive function and variable names whenever possible. The most important things to consider here are reader comprehension and searching. Consider a variable called `hdr`. Is it *header* without any vowels, or is it short for *high-dynamic-range*? Spelling out full words in variable names can save someone else a lot of guesswork.

Searching comes in when we're looking for instances of a string or variable. Single letter variable names are impossible to search for. Variables with two or three characters are often not much better.

Matlab load/save

Matlab has built-in functions to automatically save and load variables from your programs to disk. Using these functions can lead to poor program design and should be avoided if possible. It would be best to refactor as you translate if they are being used. Few operations are so expensive that that cannot be redone every time the program is run. For part of the program that saves variables, consider making a function that simply returns them instead.

If your Matlab program is loading csv files then use the Pandas library when working in python. Pandas works well with NumPy and is the go-to library when using csv files that contain numeric data.

More Resources

[NumPy for Matlab Users](#) Has a nice list of common operations in Matlab and NumPy.

[NumPy Homepage](#)

[Pandas Homepage](#)

3.4.8 Development History and Roadmap

For information on updating to the latest version of the platform see *4.0 to 5.0 migration*.

Migration from 1.2 - 2.0

The most significant changes to the base VOLTTRON platform are the set of commands for controlling the platform and the way agents are managed.

Existing agent code does not need to be modified except in cases of hardcoded paths and some imports. The way agents are packaged has changed but that does not change the setup.py file or any configuration agents were using.

Summary of changes:

- “lite” has been removed from the code tree. For packages, “lite” has been replaced by “platform”.
- The agents are no longer built as eggs but are instead built as Python wheels
- There is a new package command instead of using a script to build an egg
- Agents are no longer installed with a 2 step process of “install-executable” and “load-agent”. Now the agent package is configured then installed.
- Agents are no longer distinguished by their configuration files but can by a platform provided uuid and/or a user supplied tag.
- The base topic for publishing data from devices is no longer “RTU” but “devices”
- Application configuration files no longer need to contain the “exec” information. For an example of launching a non-Python agent, please see [ProcessAgent](#)

Migration from 2.x to 3.x

If you are upgrading an existing 2.0 installation, there are a few manual steps. From the project directory in unactivated mode:

- `rm -r env`
- `rm -r volttron/platform/control`

- python bootstrap.py

An overview of changes can be found at: [VOLTRON Primer Overview](#)

Drivers

Drivers are no longer tied to smap. Please see the [drivers](#) page.

sMAP Driver INI File

Previously, all driver setup was done in an smap.ini file with sections for each device. Now, this setup is done in two parts: the Master Driver Agent and individual drivers. The sections from the smap ini are now contained in their own files. These files are tied together by the master-driver.config:

```
{
  "agentid": "master_driver",
  "driver_config_list": [
    "/home/volttron/git/config/bacnet-device1.config",
    "/home/volttron/git/config/bacnet-device2.config"
  ]
}
```

The following portion of the file is no longer needed for the driver but could be used to setup an [sMAP Historian](#)

```
[report 0]
ReportDeliveryLocation = http://\<IP\>/backend/add/\<KEY\>

[/datalogger]
type = volttron.drivers.data_logger.DataLogger
interval = 1
```

Setting up paths for the collection and devices are now handled in the driver config file:

```
[/]
type = Collection
Metadata/SourceName = MySource
uuid = <UUID>

[/Campus]
type = Collection
Metadata/Location/Campus = My Campus

[/Campus/Building]
type = Collection
Metadata/Location/Building = Building

[/Campus/Building/device]
type = volttron.drivers.bacnet.BACnet
target_address = IP
self_address = IP:PORT
interval = 60
Metadata/Instrument/Manufacturer = Manufacturer
Metadata/Instrument/ModelName = Model Name
register_config = /home/volttron/git/volttron/config/my-bacnet-config.csv
```

Becomes bacnet-device1.config:

```
{
  "driver_config": {
    "device_address": "13200:56"
  },
  "campus": "Campus",
  "building": "Building",
  "unit": "Device",
  "driver_type": "bacnet",
  "registry_config": "/home/volttron/git/volttron/config/my-bacnet-config.csv",
  "interval": 60,
  "timezone": "US/Pacific"
}
```

Register Files (CSV)

These files are almost unchanged from v2.0. The sole change is the renaming of “PNNL Point Name” to “Volttron Point Name” This was a legacy label from the initial version of the platform and has now been updated.

Point Name, PNNL Point Name, Units, Unit Details, BACnet Object Type, Property, Writable, Index, Notes

Becomes:

Point Name, Volttron Point Name, Units, Unit Details, BACnet Object Type, Property, Writable, Index, Notes

The rest of the file remains the same.

Historian

Please look through the page [\[\[Historian|VOLTTRON-Historians\]\]](#) to see the support storage solutions. [sMAP](#) can still be used but is now optional.

ActuatorAgent

The Actuator can now be accessed via RPC which greatly simplifies the code needed to work with devices. The following shows how the old SchedulerExample agent was [upgraded](#). The `use_rpc` method contains examples for replacing all the code for the pubsub interaction.

Agents

3.X drivers

Changes from v2.X

- PNNL Point Name is now: Volttron Point Name
- Drivers are now agents
- No more smap config file, now it is an Agent config file.
- MODBUS, add port argument to driver_config dictionary
- BACnet Change of Value services are supported by the Master Driver Agent starting with version 3.2.

- Agent config file has links to driver config files which have links to driver register file.

Edit the master driver config. This points to the configuration files for specific drivers. Each of these drivers uses a CSV file to specify their points (registry file).

Master Driver Config

- agentid - name of agent
- driver_config_list - list of configuration files for drivers under this master

```
{
  "agentid": "master_driver",
  "driver_config_list": [
    "/home/user/git/voltron/services/core/MasterDriverAgent/master_driver/test_modbus1.config" | ] ] }
```

Device Driver Config

- driver_config - driver specific information, modbus just needs the ip for the device being controlled
- campus/building/unit - path to the device
- driver_type - specify the type of driver (modbus, bacnet, custom)
- registry_config - the registry file specifying points to collect
- interval - how often to grab/publish data
- timezone - TZ of data being collected
- heart_beat_point - registry point to use as a heartbeat to indicate that VOLTRON is still controlling device

```
{
  "driver_config": { "device_address": "",
    "proxy_address": "9f18c8d7-ec4b-4674-ad49-e7d0d3328f99" },
  "campus": "campus",
  "building": "building",
  "unit": "bacnet1",
  "driver_type": "bacnet",
  "registry_config": "/home/user/git/voltron/voltron/drivers/bacnet_lab.csv", | "interval": 5, | "timezone":
  "UTC" | }
```

Migration from 3.0 to 3.5

Drivers

The BACnet driver configurations now require device ids.

3.0 configurations had the line:

```
"driver_config": {"device_address": address},
```

3.5 configs needs the following addition to the the *driver_config* dictionary:

```
"driver_config": {"device_address": address,  
                  "device_id": id},
```

Historian

The 3.5 MySQL historian will try adding rows to a metadata table but will not create the table automatically. It can be added to the database with

```
CREATE TABLE meta(topic_id INTEGER NOT NULL,  
                  metadata TEXT NOT NULL,  
                  PRIMARY KEY(topic_id));
```

ActuatorAgent

The Heartbeat agent has been removed in version 3.5, its job now being done from within the actuator. The period of the heartbeat toggle function can be set by adding

```
"heartbeat_period": 20
```

to the actuator's config file. This period defaults to 60 seconds if it is not specified.

Migration from 4.1 to 5.0

5.0 includes numerous changes (Tagging Service, Message Bus performance increase, Multi-platform pub/sub, etc.), but the majority of these should be invisible to most users.

Key issues to note are:

Operations Agents

Several agents have been moved from “services/core” to “services/ops” to highlight their use in monitoring a deployment. They are not necessary when developing against a single instance, but are essential for VOLTTRON(tm) in a deployed environment.

Agents affected:

- services/ops/AgentWatcher
- services/ops/AlertAgent,0.4
- services/ops/AlertMonitor

- services/ops/Alerter
- services/ops/EmailerAgent
- services/ops/FailoverAgent
- services/ops/FileWatchPublisher
- services/ops/LogStatisticsAgent
- services/ops/MessageDebuggerAgent
- services/ops/SysMonAgent
- services/ops/ThresholdDetectionAgent

Rebuild Agents

Rebuilding agents is **required** when upgrading to a new VOLTRON(tm) version to ensure that agents are operating with the latest code. Errors will occur if agents built in a previous version attempt to run with the latest version of the platform.

ForwardHistorian

The ForwardHistorian configuration has been changed. Please see: <https://github.com/VOLTRON/voltron/blob/develop/services/core/ForwardHistorian/README.rst> for the new options.

Note: NOTE If you have no entry for service_topic_list in your configuration, the new default will cause

ALL data to be forwarded. Please update your configuration if you are forwarding a subset of data.

VOLTRON Central Management UI

The url for VOLTRON Central Management is now <http://IP:port/vc/index.html>

Agent Versions

To get the versions of agents in the VOLTRON project, run “python scripts/get_versions.py”.

Agent Name	4.1	5.0
CAgent	1.0	1.0
CSVHistorian	N/A	1.0.1
ConfigActuation	0.1	0.1
DataPublisher	3.0.1	3.0.1
DataPuller	N/A	3.5
ExampleDrivenControlAgent	0.1	0.1
ExampleSubscriber	3.0	3.0
ListenerAgent	3.2	3.2
ProcessAgent	0.1	0.1
SchedulerExample	0.1	0.1
SimpleForwarder	3.0	3.0

Continued on next page

Table 4 – continued from previous page

Agent Name	4.1	5.0
SimpleWebAgent	0.1	0.1
WeatherForecastCSV_UW		
WebRPC		
WebSocketAgent	0.0.1	0.0.1
PrometheusScrapeAgent	N/A	0.0.1
WeatherAgent		
ActuatorAgent	1.0	1.0
BACnetProxy	0.2	0.3
CrateHistorian	1.0.1	1.0.2
DataMover	0.1	0.1
ExternalData	1.0	1.0
ForwardHistorian	3.7	4.0
MQTTHistorian	0.1	0.2
MasterDriverAgent	3.1.1	3.1.1
MongodbAggregateHistorian	1.0	1.0
MongodbHistorian	2.1	2.1
MongodbTaggingService	N/A	1.0
OpenEISHistorian	3.1	3.1
SEP2Agent	N/A	1.0
SEP2DriverTestAgent	N/A	1.0
SQLAggregateHistorian	1.0	1.0
SQLHistorian	3.6.1	3.6.1
SQLiteTaggingService	N/A	1.0
VolttronCentral	4.0.3	4.2
VolttronCentralPlatform	4.0	4.5.2
WeatherAgent	3.0	3.0
AgentWatcher	0.1	0.1
AlertAgent	0.4	0.4
AlertMonitor	0.1	0.1
Alerter	0.1	0.1
MailerAgent	1.3	1.3.1
FailoverAgent	0.2	0.2
FileWatchPublisher	3.6	3.6
LogStatisticsAgent	1.0	1.0
MessageDebuggerAgent	N/A	1.0
SysMonAgent	3.6	3.6
ThresholdDetectionAgent	3.7	3.7

Change Logs

1/31/2014

The VOLTRON(tm) 1.0 release includes the following features:

- Scheduler 2.0: The new ActuatorAgent scheduler allows applications to reserve devices ahead of time
- SchedulerExample: This simple agent provides an example of publishing a schedule request.

VOLTRON v1.0 also includes features in a beta stage. These features are slated for release in v1.1 but are included in 1.0 for those who wish to begin investigating them. These features are:

- Multi-node communication: Enables platforms to publish and subscribe to each other
- BACNet Driver: Enables reading/writing to devices using the BACNet protocol

Included are PNNL developed applications: AFDD and DR which are in the process of being modified to work with the new scheduler. DR will not currently function with Scheduler 2.0.

11/7/2013

- Renamed Catalyst driver to Modbus driver to reflect the generic nature of the driver.
- Changed the configuration for the driver to fully take advantage of the Python struct module.

9/9/2013

- Catalyst registry file update for 372s
- catalystreg.csv.371 contains the points for the 371

9/4/2013

- [Scheduling](#) implemented
- Logging implemented

8/21/2013

- Added libevent-dev to required software

8/6/2013

WeatherAgent updated and back into the repository.

7/22/2017

The agent module was split into multiple pieces.

- The BaseAgent and PublishMixin classes and the periodic decorator remain in the agent package.
- The matching module was moved under the agent package and is now available as `volttron.lite.agent.matching`.
- The utility functions, like `run_agent` (which is deprecated) and the base agent `ArgumentParser`, were moved to `volttron.lite.agent.utils`.

All low-level messaging that is not agent-specific was moved to `volttron.lite.messaging` and includes the following new submodules:

- `headers`: contains common messaging headers, like `CONTENT_TYPE`, and values as constants
- `topics`: provides topic templates; see the module documentation for details
- `utils`: includes the `Topic` class and other messaging/topic utilities

The listener, control, archiver, and actuator agents were updated to use and demonstrate the changes above. Some of them also show how to use agent factories to perform dynamic matching. Using mercurial to show the diffs between revisions is a good technique for others to use to investigate how to migrate their agents.

6/24/2013

- Initial version of ExampleControllerAgent committed. This agent monitors outdoor air temp and randomly sets the coolsupply fan if temp has risen since the last reading. Wiki explanation for agent coming soon.
- Updates to ActuatorAgent
- ListenerAgent updated to reflect latest BaseAgent
- Use -config option instead of -config_path when starting agents

6/21/2013

- Updated ArchiverAgent checked in.
- ActuatorAgent for sending commands to the controller checked in.

6/19/2013

- Fixed a command line arg problem in ListenerAgent and updated wiki.

Version 1.0

This is the initial release of the Volttron Lite platform. The features contained in it are:

- Scripts for building the platform from scratch as well as updating
- A BaseAgent which expresses the basic functionality for an agent in the platform as well as hooks for adding functionality
- Example agents which utilize the BaseAgent to illustrate more complex behavior

In addition, this wiki will be constantly updated with documentation for working with the platform and developing agents in it. We intend to document as much as possible but please submit TRAC tickets in cases where documentation does not exist yet or there is difficulty locating it. Also, this is a living document so feel free to add your own content to this wiki and even make changes to the documentation if you can improve on its clarity and usefulness.

Please subscribe to this page to receive notification when new changelogs are posted for future releases.

Transitioning from 1.x to 2.x

VOLTTRON(tm) 2.0 introduces new features such as agent packaging/verification, agent mobility, and agent resource monitoring. In addition, some existing features from 1.2 have been refactored. These changes are mostly confined to platform administration and should require minimal changes to existing agents aside from fixing imports and any hardcoded paths/topics in the code.

- “lite” has been removed from the code tree. For packages, “lite” has been replaced by “platform”.
- The agents are no longer built as eggs but are instead built as Python wheels
- There is a new package command instead of using a script to build an egg

- Agents are no longer installed with a 2 step process of “install-executable” and “load-agent”. Now the agent package is configured then installed.
- Agents are no longer distinguished by their configuration files but by a platform provided uuid and/or a user supplied tag.
- The base topic for publishing data from devices is no longer “RTU” but “devices”

The most visible changes have been to the platform commands for building and managing agents. Please see [PlatformCommands] (PlatformCommands “wikilink”) for these changes.

VOLTRON Versions

VOLTRON 1.0 - 1.2

- VOLTRON platform based on PNNL research and needs of the RTU Network project
- Open Source Reimplementation omitting patented features
- Integrates researcher applications, devices, and cloud applications and resources
- 1.0 Focused on building up the framework
- Agent execution environment
- Basic platform services
- Modbus driver
- 1.2 Expanded capabilities of platform
- BACnet support
- Multi-node communication
- Released on GitHub

VOLTRON 2.0

- 2.0 Incorporated PNNL IP from the original research
- Different license: Free for buildings domain
- Resource monitoring
- Agents must present an execution contract to the platform stating their resource requirements
- Platform rejects agents which it cannot support
- Expandable framework for specify additional resources
- Agent signing and verification
- Agent package contains multiple layers which can be signed by different entities
- Creator of code
- Administrator of ‘Scope of Influence’/Deployment
- Instantiator of agent
- Most recent platform (for mobile agents)
- Each level verified before agent is allowed to run

- Entities cannot change content of other layers
- Agent Mobility
- Admin can send an agent to another platform for deployment/updating
- Agent can request to move
- Agent can bring along working files as part of ‘mutable luggage’
- Receiving platform verifies agent package and examines resource contract before executing agent

VOLTTRON 3.0

- Modularized Historian
- Historians can be built for any storage solution
 - Previous versions did not have option for local storage
- BaseHistorian
 - Can be extended for any solution
 - Handles subscribing to Bus
 - Local cache
- Modularized Drivers
- Standardized creating custom drivers to scrape data and publish to the message bus
- Simplify developing drivers and contributing new capabilities back to VOLTTRON
- Abstracted out driver interfaces allowing Actuator Agent to handle controlling devices via any protocol
- VIP - VOLTTRON Interconnect Protocol
- Increase security of the message bus and allow direct communication where appropriate
- New communication model underneath VOLTTRON Message Bus
- Compatibility layer so changes are transparent to existing agents
- Platform Agent
- Provides point of contact for the platform
- Enables VOLTTRON Management Central control of platform
- VOLTTRON Management Central
- Web interface for administering VOLTTRON platforms in deployment

3.5 Core Services

Platform services provide underlying functionality used by applications to perform their tasks.

3.5.1 Service Agents

ActuatorAgent

This agent is used to manage write access to devices. Agents may request scheduled times, called Tasks, to interact with one or more devices.

Actuator Agent Communication

Scheduling and canceling a Task.

Interacting with a device via the ActuatorAgent.

ActuatorAgent responses to a schedule or cancel request.

Schedule state announcements.

What happens when a running Task is preempted.

Setup heartbeat signal for a device.

ActuatorAgent configuration.

Notes on programming agents to work with the ActuatorAgent

Notes on Working With the ActuatorAgent

- An agent can watch the window value from [device state updates](#) to perform scheduled actions within a timeslot.
 - If an Agent's Task is LOW_PREEMPT priority it can watch for [device state updates](#) where the window is less than or equal to the grace period (default 60.0).
- When considering if to schedule long or multiple short time slots on a single device:
 - Do we need to ensure the device state for the duration between slots?
 - Yes. Schedule one long time slot instead.
 - No. Is it all part of the same Task or can we break it up in case there is a conflict with one of our time slots?
- When considering time slots on multiple devices for a single Task:
 - Is the Task really dependent on all devices or is it actually multiple Tasks?
- When considering priority:
 - Does the Task have to happen **on an exact day**?
 - No. Consider LOW and reschedule if preempted.
 - Yes. Use HIGH.
 - Is it problematic to prematurely stop a Task once started?
 - No. Consider LOW_PREEMPT and watch the [device state updates](#) for a small window value.
 - Yes. Consider LOW or HIGH.
- If an agent is only observing but needs to assure that no another Task is going on while taking readings it can schedule the time to prevent other agents from messing with a devices state. The schedule updates can be used as a reminder as to when to start watching.

- **Any** device, existing or not, can be scheduled. This allows for agents to schedule fake devices to create reminders to start working later rather than setting up their own internal timers and schedules.

ActuatorAgent Configuration

`schedule_publish_interval::` Interval between ``\ ``published

`schedule` announcements <ActuatorScheduleState>‘_‘ in seconds. Defaults to 30. |
`preempt_grace_time::` Minimum time given to Tasks which have been preempted to clean up in s
Defaults to 60. |`schedule_state_file::` File used to save and restore Task states if the Actu
File will be created if it does not exist when it is needed.

Sample configuration file

```
{  
    "schedule_publish_interval": 30,  
    "schedule_state_file": "actuator_state.pickle"  
}
```

Heartbeat Signal

The ActuatorAgent can be configured to send a heartbeat message to the device to indicate the platform is running. Ideally, if the heartbeat signal is not sent the device should take over and resume normal operation.

The configuration has two parts, the interval (in seconds) for sending the heartbeat and the specific point that should be modified each iteration.

The heart beat interval is specified with a global “heartbeat_interval” setting. The ActuatorAgent will automatically set the heartbeat point to alternating “1” and “0” values. Changes to the heartbeat point will be published like any other value change on a device.

The heartbeat points are specified in the driver configuration file of individual devices

Task Preemption

Both LOW and LOW_PREEMPT priority Tasks can be preempted. LOW priority Tasks may be preempted by a conflicting HIGH priority Task before it starts. LOW_PREEMPT priority Tasks can be preempted by HIGH priority Tasks even after they start.

When a Task is preempted the ActuatorAgent will publish to “devices/actuators/schedule/response” with the following header:

```
#python  
{  
    'type': 'CANCEL_SCHEDULE',  
    'requesterID': <Agent VIP identity for the preempted Task>,  
    'taskID': <Task ID for the preempted Task>  
}
```

And the message (after parsing the json):

```
#python
{
  'result': 'PREEMPTED',
  'info': '',
  'data':
  {
    'agentID': <Agent VIP identity of preempting task>,
    'taskID': <Task ID of preempting task>
  }
}
```

Preemption Grace Time

If a LOW_PREEMPT priority Task is preempted while it is running the Task will be given a grace period to clean up before ending. For every device which has a current time slot the window of remaining time will be reduced to the grace time. At the end of the grace time the Task will finish. If the Task has no currently open time slots on any devices it will end immediately.

Requesting Schedule Changes

For information on responses see [AcuatorAgent](#) responses to a schedule or cancel requests.

For 2.0 Agents using the pubsub interface: The actuator agent expects all messages to be JSON and will parse them accordingly. Use `publish_json` to send messages where possible.

3.0 agents using pubsub for scheduling and setting point values should publish python objects like normal.

Scheduling a Task

An agent can request a task schedule by publishing to the “devices/actuators/schedule/request” topic with the following header:

```
#python
{
  'type': 'NEW_SCHEDULE',
  'requesterID': <Ignored, VIP Identity used internally>
  'taskID': <unique task ID>, #The desired task ID for this task. It must be unique,
  ↳ among all other scheduled tasks.
  'priority': <task priority>, #The desired task priority, must be 'HIGH', 'LOW',
  ↳ or 'LOW_PREEMPT'
}
```

with the following message:

```
#python
[
  ["campus/building/device1", #First time slot.
   "2013-12-06 16:00:00",    #Start of time slot.
   "2013-12-06 16:20:00"],  #End of time slot.
  ["campus/building/device1", #Second time slot.
   "2013-12-06 18:00:00",    #Start of time slot.
   "2013-12-06 18:20:00"],  #End of time slot.
  ["campus/building/device2", #Third time slot.
```

(continues on next page)

(continued from previous page)

```

    "2013-12-06 16:00:00",      #Start of time slot.
    "2013-12-06 16:20:00"],    #End of time slot.
    #etc...
]

```

Warning: If time zones are not included in schedule requests then the Actuator will interpret them as being in local time. This may cause remote interaction with the actuator to malfunction.

Points on Task Scheduling

- Everything in the header is required.
- Task id and requester id (agentid) should be a non empty value of type string
- A Task schedule must have at least one time slot.
- The start and end times are parsed with `dateutil's date/time parser`. **The default string representation of a python datetime object will parse without issue.**
- Two Tasks are considered conflicted if at least one time slot on a device from one task overlaps the time slot of the other on the same device.
- The end time of one time slot can be the same as the start time of another time slot for the same device. This will not be considered a conflict. For example, `time_slot1(device0, time1, time2)` and `time_slot2(device0,time2, time3)` are not considered a conflict
- A request must not conflict with itself.
- If something goes wrong see [this failure string list](#) for an explanation of the error.

Task Priorities

HIGH: This Task cannot be preempted under any circumstance. This task may preempt other conflicting preemptable Tasks.

LOW: This Task cannot be preempted **once it has started**. A Task is considered started once the earliest time slot on any device has been reached. This Task may **not** preempt other Tasks.

LOW_PREEMPT: This Task may be preempted at any time. If the Task is preempted once it has begun running any current time slots will be given a grace period (configurable in the ActuatorAgent configuration file, defaults to 60 seconds) before being revoked. This Task may **not** preempt other Tasks.

Canceling a Task

A task may be canceled by publishing to the “devices/actuators/schedule/request” topic with the following header:

```

#python
{
    'type': 'CANCEL_SCHEDULE',
    'requesterID': <Ignored, VIP Identity used internally>
    'taskID': <unique task ID>, #The desired task ID for this task. It must be unique_
↪among all other scheduled tasks.
}

```

Points on Task Canceling

- The requesterID and taskID must match the original values from the original request header.
- After a Tasks time has passed there is no need to cancel it. Doing so will result in a “TASK_ID_DOES_NOT_EXIST” error.
- If something goes wrong see [this failure string list](#) for an explanation of the error.

ActuatorAgent Response

In response to a `Task schedule` request the ActuatorAgent will respond on the topic “devices/actuators/schedule/result” with the header:

```
#python
{
  'type': <'NEW_SCHEDULE', 'CANCEL_SCHEDULE'>
  'requesterID': <Agent VIP identity from the request>,
  'taskID': <Task ID from the request>
}
```

And the message (after parsing the json):

```
#python
{
  'result': <'SUCCESS', 'FAILURE', 'PREEMPTED'>,
  'info': <Failure reason, if any>,
  'data': <Data about the failure or cancellation, if any>
}
```

The ActuatorAgent may publish cancellation notices for preempted Tasks using the “PREEMPTED” result.

Preemption Data

Preemption data takes the form:

```
#python
{
  'agentID': <Agent ID of preempting task>,
  'taskID': <Task ID of preempting task>
}
```

Failure Reasons

In many cases the ActuatorAgent will try to give good feedback as to why a request failed.

General Failures

INVALID_REQUEST_TYPE:: Request type was not "NEW_SCHEDULE" or "CANCEL_SCHEDULE".

MISSING_TASK_ID:: Failed to supply a taskID.

MISSING_AGENT_ID:: AgentID not supplied.

Task Schedule Failures

TASK_ID_ALREADY_EXISTS: The supplied taskID already belongs to an existing task.

MISSING_PRIORITY: Failed to supply a priority for a Task schedule request.

INVALID_PRIORITY: Priority not one of "HIGH", "LOW", or "LOW_PREEMPT".

MALFORMED_REQUEST_EMPTY: Request list is missing or empty.

REQUEST_CONFLICTS_WITH_SELF: Requested time slots on the same device overlap.

MALFORMED_REQUEST: Reported when the request parser raises an unhandled exception. The exception name and info are appended to this info string.

CONFLICTS_WITH_EXISTING_SCHEDULES: This schedule conflict with an existing schedules that it cannot preempt. The data item for the results will contain info about the conflicts in this form (after parsing json):

```
#python
{
  '<agentID1>':
  {
    '<taskID1>':
    [
      ["campus/building/device1",
       "2013-12-06 16:00:00",
       "2013-12-06 16:20:00"],
      ["campus/building/device1",
       "2013-12-06 18:00:00",
       "2013-12-06 18:20:00"]
    ]
    '<taskID2>':[...]
  }
  '<agentID2>': {...}
}
```

Task Cancel Failures

TASK_ID_DOES_NOT_EXIST:: Trying to cancel a Task which does not exist.

This error can also occur when trying to cancel a finished Task.

AGENT_ID_TASK_ID_MISMATCH:: A different agent ID is being used when trying to cancel a Task.

Schedule State Broadcast

Periodically the ActuatorAgent will publish the state of all currently used devices.

For each device the ActuatorAgent will publish to an associated topic:

```
#python
'devices/actuators/schedule/announce/<full device path>'
```

With the following header:

```
#python
{
  'requesterID': <VIP identity of agent with access>,
  'taskID': <Task associated with the time slot>
  'window': <Seconds remaining in the time slot>
}
```

The frequency of the updates is configurable with the “schedule_publish_interval” setting.

ActuatorAgent Interaction

Once an Task has been scheduled and the time slot for one or more of the devices has started an agent may interact with the device using the **get** and **set** topics.

Both **get** and **set** are responded to the same way. See [#ActuatorReply Actuator Reply] below.

Getting values

While the sMap driver for a device should always be setup to periodically broadcast the state of a device you may want an up to the moment value for an actuation point on a device.

To request a value publish a message to the following topic:

```
#python
'devices/actuators/get/<full device path>/<actuation point>'
```

Setting Values

Value are set in a similar manner:

To set a value publish a message to the following topic:

```
#python
'devices/actuators/set/<full device path>/<actuation point>'
```

With this header:

```
#python
{
  'requesterID': <Ignored, VIP Identity used internally>
}
```

And the message contents being the new value of the actuator.

The actuator agent expects all messages to be JSON and will parse them accordingly. Use publish_json to send messages where possible. This is significant for Boolean values especially.

Actuator Reply

#ActuatorReply The ActuatorAgent will reply to both **get** and **set** on the **value** topic for an actuator:

```
#python
'devices/actuators/value/<full device path>/<actuation point>'
```

With this header:

```
#python
{
  'requesterID': <Agent VIP identity>
}
```

With the message containing the value encoded in JSON.

Actuator Error Reply

If something goes wrong the ActuatorAgent will reply to both **get** and **set** on the **error** topic for an actuator:

```
#python
'devices/actuators/error/<full device path>/<actuation point>'
```

With this header:

```
#python
{
  'requesterID': <Agent VIP identity>
}
```

The message will be in the following form:

```
#python
{
  'type': <Error Type or name of the exception raised by the request>
  'value': <Specific info about the error>
}
```

Common Error Types

LockError:: Returned when a request is made when we do not have permission to use a device (Forgot to schedule, preempted and we did not handle the preemption message correctly, ran out of time in time slot, etc...)

ValueError:: Message missing or could not be parsed as JSON.

Other error types involve problem with communication between the ActuatorAgent and sMap.

Alert Agent

Alert Agent

The Alert Agent listens to a set of configured topics and publishes an alert if they are not published within some time limit. In addition to “standard” topics the Alert Agent supports inspecting device *all* topics. This can be useful when

a device contains volatile points that may not be published.

Configuration

Topics are organized by groups. Any alerts raised will summarize all missing topics in the group.

Individual topics have two configuration options. For standard topics configuration consists of a key value pair of the topic to its time limit.

The other option is for *all* publishes. The topic key is paired with a dictionary that has two keys, “seconds” and “points”. “seconds” is the topic’s time limit and “points” is a list of points to watch.

```
{
  "groupname": {
    "devices/fakedriver0/all": 10,

    "devices/fakedriver1/all": {
      "seconds": 10,
      "points": ["temperature", "PowerState"]
    }
  }
}
```

Emailer Agent

Emailer Agent

Emailer agent is responsible for sending emails for an instance. It has been written so that any agent on the instance can send emails through it via the “send_email” method or through the pubsub message bus using the topic “platform/send_email”.

By default any alerts will be sent through this agent. In addition all emails will be published to the “record/sent_email” topic for a historian to be able to capture that data.

Configuration

A typical configuration for this agent is as follows. We need to specify the SMTP server address, email address of the sender, email addresses of all the recipients and minimum time for duplicate emails based upon the key.

```
{
  "smtp-address": "smtp.foo.com",
  "from-address": "billy@foo.com",
  "to-addresses": ["ann@foo.com", "bob@gmail.com"],
  "allow-frequency-minutes": 10
}
```

Finally package, install and start the agent. For more details, see *Agent Creation Walkthrough*

Failover Agent

Failover Agent

- **agent_id** - primary **or** secondary
- **simple_behavior** - Switch to turn on or off simple behavior. Both instances should match.
- **remote_vip** - Address where *remote_id* can be reached.
- **remote_serverkey** - The public key of the platform where *remote_id* lives.
- **agent_vip_identity** - The vip identity of the agent that we want to manage.
- **heartbeat_period** - Send a message to *remote_id* with this period. Measured in seconds.
- **timeout** - Consider a platform inactive if a heartbeat has not been received for *timeout* seconds.

File Watch Publisher Agent

FileWatchPublisher Agent

Introduction

FileWatchPublisher agent watches files for changes and publishes those changes per line on the corresponding topics. Files and topics should be provided in the configuration.

Configuration

A simple configuration for FileWatchPublisher with two files to monitor is as follows:

```
[
  {
    "file": "/var/log/syslog",
    "topic": "platform/syslog"
  },
  {
    "file": "/home/volttron/tempfile.txt",
    "topic": "temp/filepublisher"
  }
]
```

Using this example configuration, FileWatchPublisher will watch syslog and tempFile.txt files and publish the changes per line on their respective topics.

Platform Agent

Platform Agent

Introduction

The Platform Agent allows communication from a VOLTRON Central instance. Each VOLTRON instance that is to be controlled through the VOLTRON Central agent should have one and only one Platform Agent. The Platform Agent must have the VIP identity of platform.agent.

Configuration

The minimal configuration (and most likely the only used) for a Platform Agent is as follows

```
{
  # Agent id is used in the display on volttron central.
  "agentid": "Platform 1",
}
```

The other options for the Platform Agent configuration can be found in the Platform Agent source directory.

Market Service Agent

Market Service Agent

Introduction

The MarketServiceAgent implements a variation of a double-blind auction, in which each market participant bids to buy or sell a commodity for a given price.

In contrast to other common implementations, participants do not bid single price-quantity pairs. Instead, they bid a price-quantity curve, or “flexibility curve” into their respective markets. Market participants may be both buyers in one market and sellers in another. Settling of the market is a “single shot” process that begins with bidding that progresses from the bottom up and concludes with a clearing of the markets from the top down. This is termed “single shot” because there is no iteration required to find the clearing price or quantity at any level of the market structure. Once the market has cleared, the process begins again for the next market interval, and new bids are submitted based on the updated states of the agents.

Market Timing

The MarketServiceAgent is driven by the Director. The Director drives the MarketServiceAgent through a timed loop. The Director has just a few parameters that are configured by default with adequate values. They are:

1. The `market_period` with a default value of 5 minutes
2. The `reservation_delay` with a default value of 0 minutes
3. The `offer_delay` with a default value of 2 minutes

The timing loop works as follows:

- The market period begins.
- A request for reservations is published after the reservation delay.
- A request for offers/bids is published after the offer delay.
- The aggregate demand curve is published as soon all the buy offers are completed for the market.
- The aggregate supply curve is published as soon all the sell offers are completed for the market.
- The cleared price is published as soon as all bids have been received.
- Error messages are published when discovered and usually occur at the end of one of the delays.
- The cycle repeats.

How to Use the MarketServiceAgent

A given agent participates in one or more markets by inheriting from the *base MarketAgent*. The base MarketAgent handles all of the communication between the agent and the MarketServiceAgent. The agent only needs to join each market with the `join_market` method and then respond to the appropriate callback methods. The callback methods are described at the *base MarketAgent*.

Threshold Detection Agent

Threshold Detection Agent

The ThresholdDetectionAgent will publish an alert when a value published to a topic exceeds or falls below a configured value. The agent can be configured to watch topics are associated with a single value or to watch devices' all topics.

Configuration

The ThresholdDetectionAgent supports the *configstore* and can be configured with a file named “config”.

The file must be in the following format:

- Topics and points in device publishes may have maximum and minimum thresholds but both are not required
- A device's point entries are configured the same way as standard topic entries

```
{
  "topic": {
    "threshold_max": 10
  },
  "devices/some/device/all": {
    "point0": {
      "threshold_max": 10,
      "threshold_min": 0
    },
    "point1": {
      "threshold_max": 42
    }
  }
}
```

VOLTRON Central Management

VOLTRON Central Management Agent

Agent Introduction

The VOLTRON Central Agent (VCM) is responsible for controlling multiple VOLTRON instances through a single interfaces. The VOLTRON instances can be either local or remote. VCM leverages an internal VOLTRON web server providing a interface to our JSON-RPC based web api. Both the web api and the interface are served through the VCM agent. There is a *VOLTRON Central Demo* that will allow you to quickly setup and see the current offerings of the interface. VOLTRON Central will allow you to

- See a list of platforms being managed.
- Add and remove platforms.
- Install, start and stop agents to the registered platforms.
- Create dynamic graphs from the historians based upon points.
- Execute functions on remote platforms.

Note: see *VCM json-rpc web api* for how the web interface works.

Instance Configuration

In order for any web agent to be enabled, there must be a port configured to serve the content. The easiest way to do this is to create a config file in the root of your VOLTRON_HOME directory. (to do this automatically see *VOLTRON Config*)

The following is an example of the configuration file

```
[volttron]
vip-address=tcp://127.0.0.1:22916
bind-web-address=http://127.0.0.1:8080/vc/
```

**** Note the above configuration will open a discoverable port for the volttron instance.** In addition, the opening of this web address allows you to serve both static as well as dynamic pages.

Verify that the instance is serving properly by pointing your web browser to

```
http://127.0.0.1:8080/discovery/
```

This is the required information for a VolttronCentralPlatform to be able to be managed.

VOLTRON Central Manager Configuration

The following is the default configuration file for VOLTRON Central

```
{
  # The agentid is used during display on the VOLTRON central platform
  # it does not need to be unique.
  "agentid": "volttron central",

  # Authentication for users is handled through a naive password algorithm
  # Note in the following example the user and password are both admin.

  # DO NOT USE IN PRODUCTION ENVIRONMENT!

  # import hashlib
  # hashlib.sha512(password).hexdigest() where password is the plain text password.
  "users" : {
    "reader" : {
      "password" :
↵ "2d7349c51a3914cd6f5dc28e23c417ace074400d7c3e176bcf5da72fdbeb6ce7ed767ca00c6c1fb754b8df5114fc0b903
↵",
      "groups" : [
```

(continues on next page)

(continued from previous page)

```

        "reader"
      ]
    },
    "writer" : {
      "password" :
↪ "f7c31a682a838bbe0957cfa0bb060daff83c488fa5646eb541d334f241418af3611ff621b5a1b0d327f1ee80da25e04099
↪ ",
      "groups" : [
        "writer"
      ]
    },
    "admin" : {
      "password" :
↪ "c7ad44cbad762a5da0a452f9e854fdc1e0e7a52a38015f23f3eab1d80b931dd472634dfac71cd34ebc35d16ab7fb8a90c
↪ ",
      "groups" : [
        "admin"
      ]
    },
    "dorothy" : {
      "password" :
↪ "cf1b67402d648f51ef6ff8805736d588ca07cbf018a5fba404d28532d839a1c046bfcd31558dff658678b3112502f4da9
↪ ",
      "groups" : [
        "reader, writer"
      ]
    }
  }
}

```

Agent Execution

To start VOLTRON Central first make sure the VOLTRON instance is running Next create/choose the config file to use. Finally from an activated shell in the root of the VOLTRON repository execute

```

# Arguments are package to execute, config file to use, tag to use as reference
./scripts/core/pack_install.sh services/core/VolttronCentral services/core/
↪ VolttronCentral/config vc

# Start the agent
volttron-ctl start --tag vc

```

VOLTRON Central Web Services Api Documentation

VOLTRON Central (VC) is meant to be the hub of communication within a cluster of VOLTRON instances. VC exposes a [JSON-RPC 2.0](#) based api that allows a user to control multiple instances of VOLTRON.

Why JSON-RPC

SOAP messaging is unfriendly to many developers, especially those wanting to make calls in a browser from AJAX environment. We have therefore have implemented a JSON-RPC API capability to VC, as a more JSON/JavaScript friendly mechanism.

How the API is Implemented

- All calls are made through a POST to `/vc/jsonrpc`
- All calls (not including the call to authenticate) will include an authorization token (a json-rpc extension).

JSON-RPC Request Payload

All posted JSON payloads will look like the following block:

```
{
  "jsonrpc": "2.0",
  "method": "method_to_invoke",
  "params": {
    "param1name": "param1value",
    "param2name": "param2value"
  },
  "id": "unique_message_id",
  "authorization": "server_authorization_token"
}
```

As an alternative, the params can be an array as illustrated by the following:

```
{
  "jsonrpc": "2.0",
  "method": "method_to_invoke",
  "params": [
    "param1value",
    "param2value"
  ],
  "id": "unique_message_id",
  "authorization": "server_authorization_token"
}
```

For full documentation of the Request object please see section 4 of the [JSON-RPC 2.0](#) specification.

JSON-RPC Response Payload

All responses shall have either an error response or a result response. The result key shown below can be a single instance of a json type, an array or a JSON object.

A result response will have the following format:

```
{
  "jsonrpc": "2.0",
  "result": "method_results",
  "id": "sent_in_unique_message_id"
}
```

An error response will have the following format:

```
{
  "jsonrpc": "2.0",
  "error": {
    "code": "standard_code_or_extended_code",

```

(continues on next page)

(continued from previous page)

```

    "message": "error message"
  }
  "id": "sent_in_unique_message_id_or_null"
}

```

For full documentation of the Response object please see section 5 of the [JSON-RPC 2.0](#) specification.

JSON-RPC Data Objects

Table 5: Platform

Key	Type	Value
uuid	string	A unique identifier for the platform.
name	string	A user defined string for the platform.
status	Status	A status object for the platform.

Table 6: PlatformDetails

Key	Type	Value
uuid	string	A unique identifier for the platform.
name	string	A user defined string for the platform.
status	Status	A status object for the platform.

Table 7: Agent

Key	Type	Value
uuid	string	A unique identifier for the agent.
name	string	Defaults to the agentid of the installed agent
tag	string	A shortcut that can be used for referencing the agent
priority	int	If this is set the agent will autostart on the instance.
process_id	int	The process id or null if not running.
status	string	A status string made by the status rpc call, on an agent.

Table 8: AdvancedRegistryEntry_TODO

Key	Type	Value
name		
vip_address		

Table 9: Agent_TODO

Key	Type	Value
uuid	string	A unique identifier for the platform.
name	string	A user defined string for the platform.
status	Status	A status object for the platform.

Table 10: Building_TODO

Key	Type	Value
uuid	string	A unique identifier for the platform.
name	string	A user defined string for the platform.
status	Status	A status object for the platform.

Table 11: Device_TODO

Key	Type	Value
uuid	string	A unique identifier for the platform.
name	string	A user defined string for the platform.
status	Status	A status object for the platform.

Table 12: Status

Key	Type	Value
status	string	A value of GOOD, BAD, UNKNOWN, SUCCESS, FAIL
context	string	Provides context about what the status means (optional)

JSON-RPC API Methods

header “method”, “parameters”, “returns” :widths: 10, 10, 40
“get_authentication”, “(username, password)”, “authentication token”

Messages

Retrieve Authorization Token

```
# POST /vc/jsonrpc
{
  "jsonrpc": "2.0",
  "method": "get_authentication",
  "params": {
    "username": "dorothy",
    "password": "toto123"
  },
  "id": "someID"
}
```

Response Success

```
# 200 OK
{
  "jsonrpc": "2.0",
  "result": "somAuthorizationToken",
  "id": "someID"
}
```

Failure HTTP Status Code 401

Register A Volttron Platform Instance (Using Discovery)

```
# POST /vc/jsonrpc
{
  "jsonrpc": "2.0",
  "method": "register_instance",
  "params": {
    "discovery_address": "http://127.0.0.2:8080",
    "display_name": "foo" # Optional
  }
  "authorization": "someAuthorizationToken",
  "id": "someID"
}
```

Success

```
# 200 OK
{
  "jsonrpc": "2.0",
  "result": {
    "status": {
      "code": "SUCCESS"
      "context": "Registered instance foo" # or the uri if not_
↪specified.
    }
  },
  "id": "someID"
}
```

TODO: Request Registration of an External Platform

```
# POST /vc/jsonrpc
{
  "jsonrpc": "2.0",
  "method": "register_platform",
  "params": {
    "uri": "127.0.0.2:8080?serverkey=...&publickey=...&secretkey=..."
  }
  "authorization": "someAuthorizationToken",
  "id": #
}
```

Unregister a Volttron Platform Instance

```
# POST /vc/jsonrpc
{
  "jsonrpc": "2.0",
  "method": "unregister_platform",
  "params": {
    "platform_uuid": "somePlatformUuid",
  }
  "authorization": "someAuthorizationToken",
  "id": "someID"
}
```

Retrieve Managed Instances

```
#POST /vc/jsonrpc
{
```

(continues on next page)

(continued from previous page)

```
"jsonrpc": "2.0",
"method": "list_platforms",
"authorization": "someAuthorizationToken",
"id": #
}
```

Response Success

```
200 OK
{
  "jsonrpc": "2.0",
  "result": [
    {
      "name": "platform1",
      "uuid": "abcd1234-ef56-ab78-cd90-efabcd123456",
      "health": {
        "status": "GOOD",
        "context": null,
        "last_updated": "2016-04-27T19:47:05.184997+00:00"
      }
    },
    {
      "name": "platform2",
      "uuid": "0987fedc-65ba-43fe-21dc-098765bafedc",
      "health": {
        "status": "BAD",
        "context": "Expected 9 agents running, but only 5 are",
        "last_updated": "2016-04-27T19:47:05.184997+00:00",
      }
    },
    {
      "name": "platform3",
      "uuid": "0000aaaa-1111-bbbb-2222-cccc3333dddd",
      "health": {
        "status": "GOOD",
        "context": "Currently scraping 20 devices",
        "last_updated": "2016-04-27T19:47:05.184997+00:00",
      }
    }
  ],
  "id": #
}
```

TODO: change response Retrieve Installed Agents From platform1

```
# POST /vc/jsonrpc
{
  "jsonrpc": "2.0",
  "method": "platforms.uuid.abcd1234-ef56-ab78-cd90-efabcd123456.list_agents",
  "authorization": "someAuthorizationToken",
  "id": #
}
```

Response Success

```

200 OK
{
  "jsonrpc": "2.0",
  "result": [
    {
      "name": "HelloAgent",
      "identity": "helloagent-0.0_1",
      "uuid": "alb2c3d4-e5f6-a7b8-c9d0-e1f2a3b4c5d6",
      "process_id": 3142,
      "error_code": null,
      "is_running": true,
      "permissions": {
        "can_start": true,
        "can_stop": true,
        "can_restart": true,
        "can_remove": true
      }
      "health": {
        "status": "GOOD",
        "context": null
      }
    },
    {
      "name": "Historian",
      "identity": "sqlhistorianagent-3.5.0_1",
      "uuid": "alb2c3d4-e5f6-a7b8-c9d0-e1f2a3b4c5d6",
      "process_id": 3143,
      "error_code": null,
      "is_running": true,
      "permissions": {
        "can_start": true,
        "can_stop": true,
        "can_restart": true,
        "can_remove": true
      }
      "health": {
        "status": "BAD",
        "context": "No publish in last 5 minutes"
      }
    },
    {
      "name": "VolltronCentralPlatform",
      "identity": "platform.agent",
      "uuid": "alb2c3d4-e5f6-a7b8-c9d0-e1f2a3b4c5d6",
      "process_id": 3144,
      "error_code": null,
      "is_running": true,
      "permissions": {
        "can_start": false,
        "can_stop": false,
        "can_restart": true,
        "can_remove": false
      }
      "health": {
        "status": "BAD",
        "context": "One agent has reported bad status"
      }
    }
  ]
}

```

(continues on next page)

(continued from previous page)

```

    }
  },
  {
    "name": "StoppedAgent-0.1",
    "identity": "stoppedagent-0.1_1",
    "uuid": "a1b2c3d4-e5f6-a7b8-c9d0-e1f2a3b4c5d6",
    "process_id": null,
    "error_code": 0,
    "is_running": false,
    "health": {
      "status": "UNKNOWN",
      "context": "Error code -15"
    }
    "permissions": {
      "can_start": true,
      "can_stop": false,
      "can_restart": true,
      "can_remove": true
    }
  }
],
  "id": #
}

```

TODO: Start An Agent

```

# POST /vc/jsonrpc
{
  "jsonrpc": "2.0",
  "method": "platforms.uuid.0987fedc-65ba-43fe-21dc-098765bafedc.start_agent",
  "params": ["a1b2c3d4-e5f6-a7b8-c9d0-e1f2a3b4c5d6"],
  "authorization": "someAuthorizationToken",
  "id": #
}

```

Response Success

```

200 OK
{
  "jsonrpc": "2.0",
  "result": {
    "process_id": 1000,
    "return_code": null
  },
  "id": #
}

```

TODO: Stop An Agent

```

# POST /vc/jsonrpc
{
  "jsonrpc": "2.0",
  "method": "platforms.uuid.0987fedc-65ba-43fe-21dc-098765bafedc.stop_agent",
  "params": ["a1b2c3d4-e5f6-a7b8-c9d0-e1f2a3b4c5d6"],
  "authorization": "someAuthorizationToken",
  "id": #
}

```

Response Success

```
200 OK
{
  "jsonrpc": "2.0",
  "result": {
    "process_id": 1000,
    "return_code": 0
  },
  "id": #
}
```

TODO: Remove An Agent

```
# POST /vc/jsonrpc
{
  "jsonrpc": "2.0",
  "method": "platforms.uuid.0987fedc-65ba-43fe-21dc-098765bafedc.remove_agent",
  "params": ["a1b2c3d4-e5f6-a7b8-c9d0-e1f2a3b4c5d6"],
  "authorization": "someAuthorizationToken",
  "id": #
}
```

Response Success

```
200 OK
{
  "jsonrpc": "2.0",
  "result": {
    "process_id": 1000,
    "return_code": 0
  },
  "id": #
}
```

TODO: Retrieve Running Agents

```
# POST /vc/jsonrpc
{
  "jsonrpc": "2.0",
  "method": "platforms.uuid.0987fedc-65ba-43fe-21dc-098765bafedc.status_agents",
  "authorization": "someAuthorizationToken",
  "id": #
}
```

Response Success

```
200 OK
{
  "jsonrpc": "2.0",
  "result": [
    {
      "name": "RunningAgent",
      "uuid": "a1b2c3d4-e5f6-a7b8-c9d0-e1f2a3b4c5d6"
      "process_id": 1234,
      "return_code": null
    },
    {
```

(continues on next page)

(continued from previous page)

```

        "name": "StoppedAgent",
        "uuid": "alb2c3d4-e5f6-a7b8-c9d0-e1f2a3b4c5d6"
        "process_id": 1000,
        "return_code": 0
    }
],
    "id": #
}

```

TODO: currently getting 500 error Retrieve An Agent's RPC Methods

```

# POST /vc/jsonrpc
{
    "jsonrpc": "2.0",
    "method": "platforms.uuid.0987fedc-65ba-43fe-21dc-098765bafedc.agents.
↪uuid.alb2c3d4-e5f6-a7b8-c9d0-e1f2a3b4c5d6.inspect",
    "authorization": "someAuthorizationToken",
    "id": #
}

```

Response Success

```

200 OK
{
    "jsonrpc": "2.0",
    "result": [
        {
            "method": "sayHello",
            "params": {
                "name": "string"
            }
        }
    ],
    "id": #
}

```

TODO: Perform Agent Action

```

# POST /vc/jsonrpc
{
    "jsonrpc": "2.0",
    "method": "platforms.uuid.0987fedc-65ba-43fe-21dc-098765bafedc.agents.uuid.
↪alb2c3d4-e5f6-a7b8-c9d0-e1f2a3b4c5d6.methods.say_hello",
    "params": {
        "name": "Dorothy"
    },
    "authorization": "someAuthorizationToken",
    "id": #
}

```

Success Response

```

200 OK
{
    "jsonrpc": "2.0",
    "result": "Hello, Dorothy!",
}

```

(continues on next page)

(continued from previous page)

```

    "id": #
  }

```

TODO: Install Agent

```

# POST /vc/jsonrpc
{
  "jsonrpc": "2.0",
  "method": "platforms.uuid.0987fedc-65ba-43fe-21dc-098765bafedc.install",
  "params": {
    "files": [
      {
        "file_name": "helloagent-0.1-py2-none-any.whl",
        "file": "data:application/octet-stream;base64,..."
      },
      {
        "file_name": "some-non-wheel-file.txt",
        "file": "data:application/octet-stream;base64,..."
      },
      ...
    ],
    "authorization": "someAuthorizationToken",
    "id": #
  }
}

```

Success Response

```

200 OK
{
  "jsonrpc": "2.0",
  "result": {
    [
      {
        "uuid": "a1b2c3d4-e5f6-a7b8-c9d0-ef2a3b4c5d6"
      },
      {
        "error": "Some error message"
      },
      ...
    ]
  },
  "id": #
}

```

Weather Service Agent

The Weather Agent provides weather data retrieved from Weather Underground to agents. It provides an example of accessing an external resource and publishing it to the VOLTRON Message Bus.

Weather Agent

The Weather Agent provides weather data retrieved from Weather Underground to agents. It provides an example of accessing an external resource and publishing it to the VOLTRON Message Bus. It also provides an example of

dividing data up into topics and for implementing the “all” topic at each level. The Weather agent retrieves data in two ways.

Usage

In the first, the agent periodically retrieves information about a specific area from weather underground and publishes current information periodically (default once an hour) to the message bus. The area is set in the agents configuration file. Agents that are interested in this information can subscribe to “*weather/<subtopic>/<field>*” where *<subtopic>* and *<field>* are described in the topics section below. Agents may also subscribe to “*weather/all*” which will give them a hierarchical json document laid out according to the weather agent topics (as below) or “*weather/<subtopic>/all*” which provides a json document with all of the fields for the weather topic specified.

The second mode of operation is on demand. An agent may specifically request weather data for an area by posting a message on the “*weather/request*” topic. The agents message includes the identifying information about the area (region/city, or zip code) which is used in the request to Weather Underground. The Weather Agent replies on the “*weather/response*” topic (to avoid confusion with agents subscribing to the periodic posts). As with the first mode of operation, the agent making the request can subscribe to the “*weather/response/all*” subtopic, “*weather/response/<subtopic>/all*”, or “*weather/response/<subtopic>/<field>*” topics to receive the response data from the Weather Agent. Agents making requests must have the requesterID header set.

3.0 Agent Example

To make a request of the weather agent.

```
Python 2.7.6 (default, Jun 22 2015, 17:58:13)
[GCC 4.8.2] on linux2
Type "help", "copyright", "credits" or "license" for more information.
>>> from volttron.platform.vip.agent import *
>>> a = Agent()
>>> import gevent
>>> gevent.spawn(a.core.run).join(0)
>>> a.vip.pubsub.publish('pubsub', 'weather/request', headers={'requesterID': 'agentid
↵'}, message={'zipcode': '99336'}).get(timeout=10)
```

```
headers = {}
headers[headers_mod.REQUESTER_ID] = agent_id
msg = {"zipcode": "99352"}
self.publish_json('weather/request', headers, msg)
```

Make sure you get your own API Key from Weather Underground before using the weather agent.

Sample Output

Following is the response message that would be returned on “*weather/all*” or “*weather/response/all*”.

```
{
  "cloud_cover": {
    "UV": "6",
    "solarradiation": "",
    "visibility_km": "16.1",
    "visibility_mi": "10.0",
    "weather": "Clear"
```

(continues on next page)

(continued from previous page)

```

    },
    "location": {
      "display_location": {
        "city": "Richland",
        "country": "US",
        "country_iso3166": "US",
        "elevation": "121.00000000",
        "full": "Richland, WA",
        "latitude": "46.28490067",
        "longitude": "-119.29721832",
        "state": "WA",
        "state_name": "Washington",
        "zip": "99352"
      },
      "local_tz_long": "America/Los_Angeles",
      "observation_location": {
        "city": "Richland, Richland",
        "country": "US",
        "country_iso3166": "US",
        "elevation": "397 ft",
        "full": "Richland, Richland, Washington",
        "latitude": "46.285866",
        "longitude": "-119.304375",
        "state": "Washington"
      },
      "station_id": "KWARICHL21"
    },
    "precipitation": {
      "dewpoint_c": 7,
      "dewpoint_f": 44,
      "dewpoint_string": "44 F (7 C)",
      "precip_1hr_in": "0.00",
      "precip_1hr_metric": "0",
      "precip_1hr_string": "0.00 in ( 0 mm)",
      "precip_today_in": "0.00",
      "precip_today_metric": "0",
      "precip_today_string": "0.00 in (0 mm)"
    },
    "pressure_humidity": {
      "pressure_mb": "1014",
      "pressure_trend": "-",
      "relative_humidity": "40%"
    },
    "temperature": {
      "feelslike_c": "20.6",
      "feelslike_f": "69.1",
      "feelslike_string": "69.1 F (20.6 C)",
      "heat_index_c": "NA",
      "heat_index_f": "NA",
      "heat_index_string": "NA",
      "temp_c": 20.6,
      "temp_f": 69.1,
      "temperature_string": "69.1 F (20.6 C)",
      "windchill_c": "NA",
      "windchill_f": "NA",
      "windchill_string": "NA"
    },
  },

```

(continues on next page)

(continued from previous page)

```
"time": {
  "local_epoch": "1368724778",
  "local_time_rfc822": "Thu, 16 May 2013 10:19:38 -0700",
  "local_tz_offset": "-0700",
  "local_tz_short": "PDT",
  "observation_epoch": "1368724692",
  "observation_time": "Last Updated on May 16, 10:18 AM PDT",
  "observation_time_rfc822": "Thu, 16 May 2013 10:18:12 -0700"
},
"wind": {
  "pressure_in": "29.94",
  "wind_degrees": 3,
  "wind_dir": "North",
  "wind_gust_kph": "4.8",
  "wind_gust_mph": "3.0",
  "wind_kph": 2.7,
  "wind_mph": 1.7,
  "wind_string": "From the North at 1.7 MPH Gusting to 3.0 MPH"
}
}
```

For a more comprehensive listing of Weather Agent subtopics see *WeatherAgentTopics*

Weather Agent Installation

Configuring and Launching the Weather Agent

The Weather agent, another VOLTRON service agent, retrieves weather information from the Weather Underground site and shares it with agents running on the platform. The first step to launching the Weather agent is to obtain a developer key from Weather Underground.

Obtaining a Developer Key from Weather Underground

Follow these steps to create a Weather Underground account and obtain a developer key.

- Go to Weather Underground site (Figure 1) the following URL <http://www.wunderground.com/weather/api/>
- Select, Sign Up for FREE

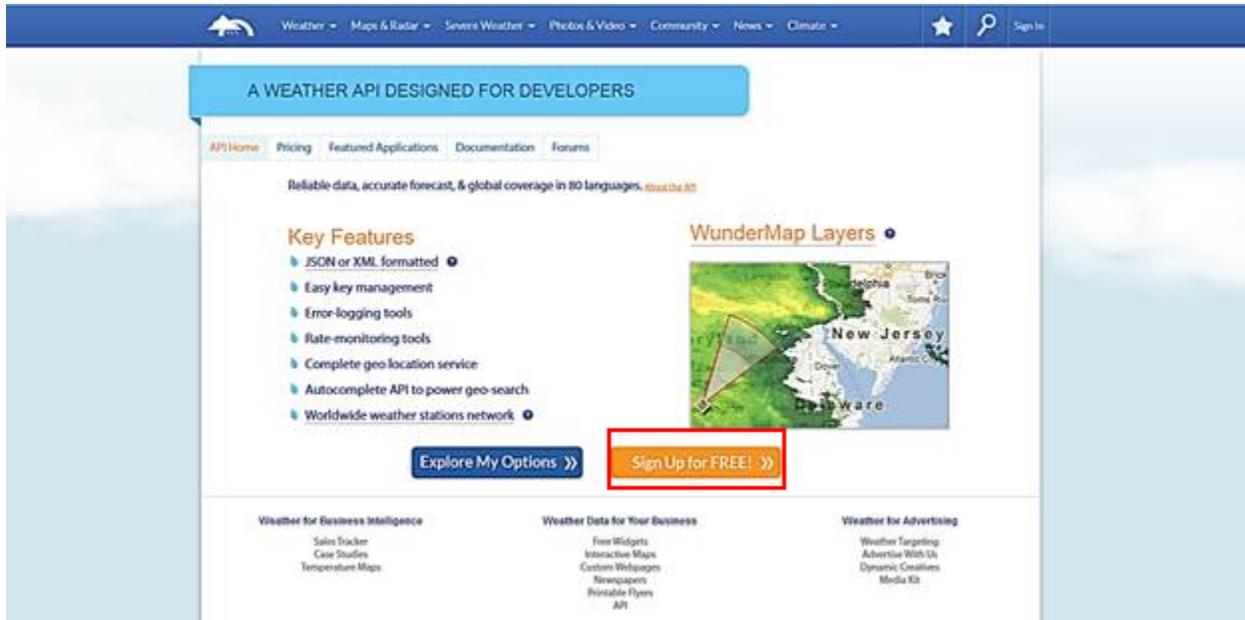


Figure 1: Weather Underground Website

- The window should now look similar to Figure 2. Enter your information to create an account.

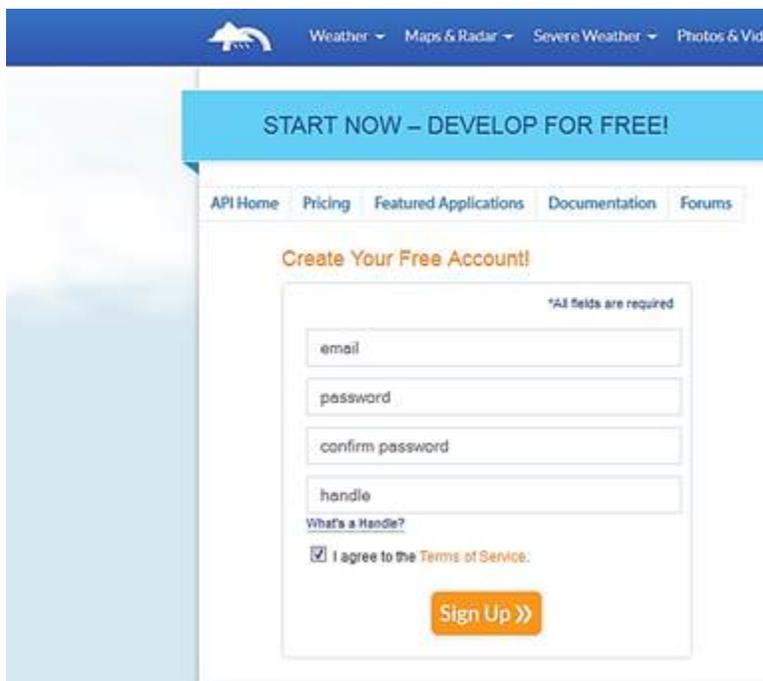


Figure 2: Setting up a developer account

- Select a plan that meets your needs. Login to with your username and password and click on “Explore my options button.” For most applications, the free plan will be adequate. The window should appear similar to Figure 3:

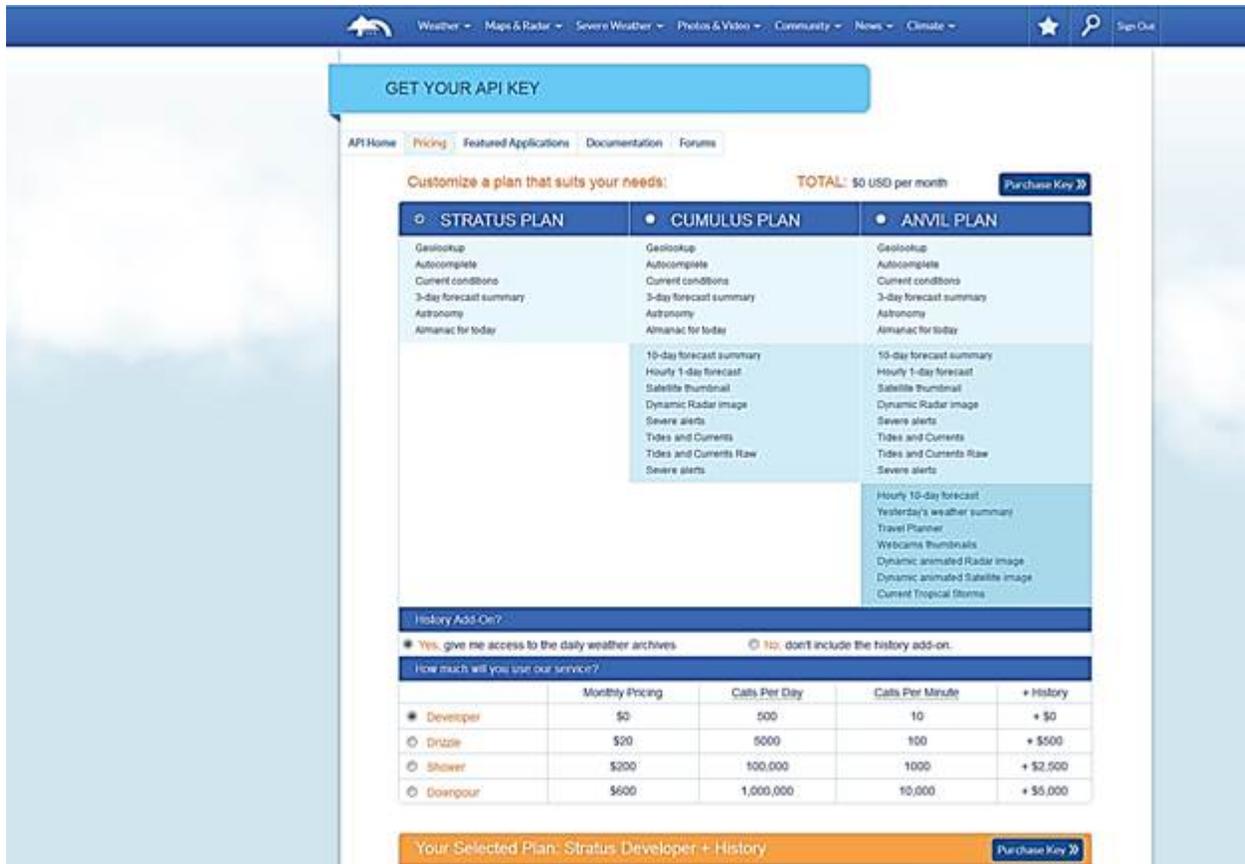


Figure 3: Creating a Weather Underground API key

- You now have access to your Weather Underground API key. An example API key is shown in the red box of Figure 4:

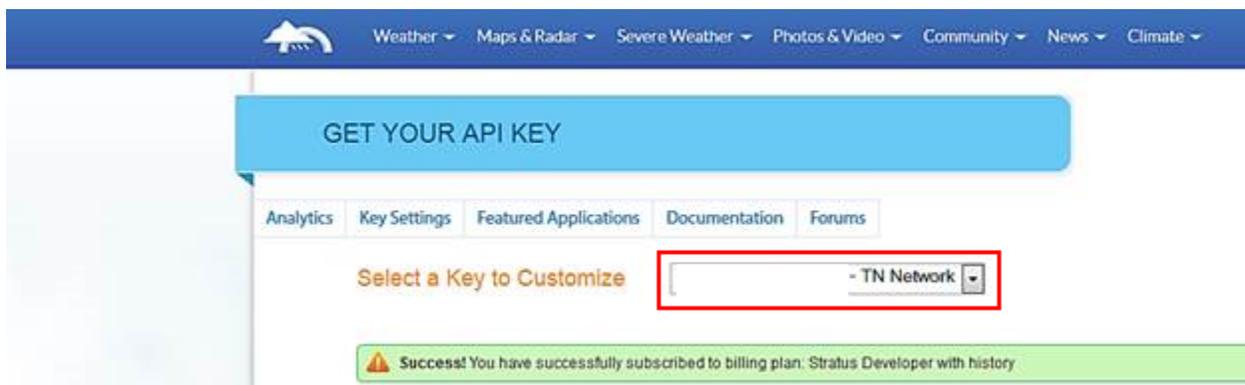


Figure 4: Weather Underground API key

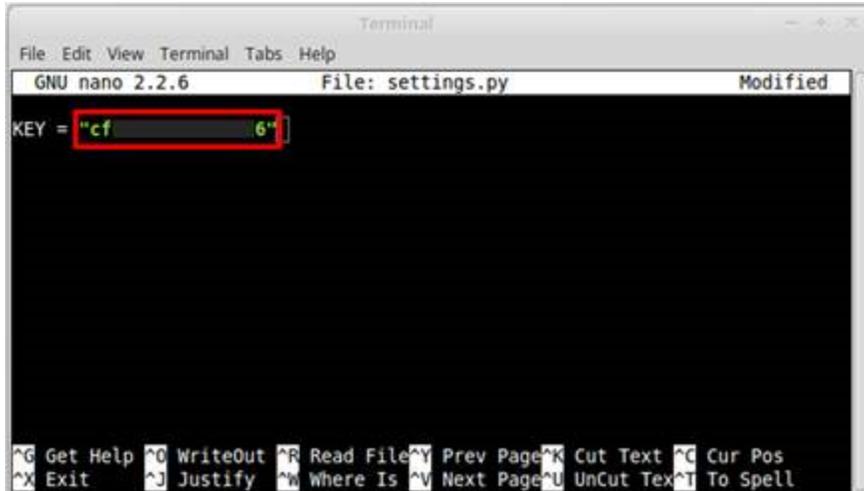
Configuring Weather Agent with API Key and Location

The following steps will show how to configure the Weather agent with the developer key from Weather Underground and how to enter a zip code to get weather data from that zip code. Edit (`<project directory>/services/core/WeatherAgent/weather/settings.py`) with your Weather Underground key. From the base VOLTTRON directory, enter the following terminal commands:

- Open settings.py at with a text editor or nano:

```
$ nano services/core/weather/settings.py
```

- Enter a Weather Underground Developer key, as shown in Figure 5:

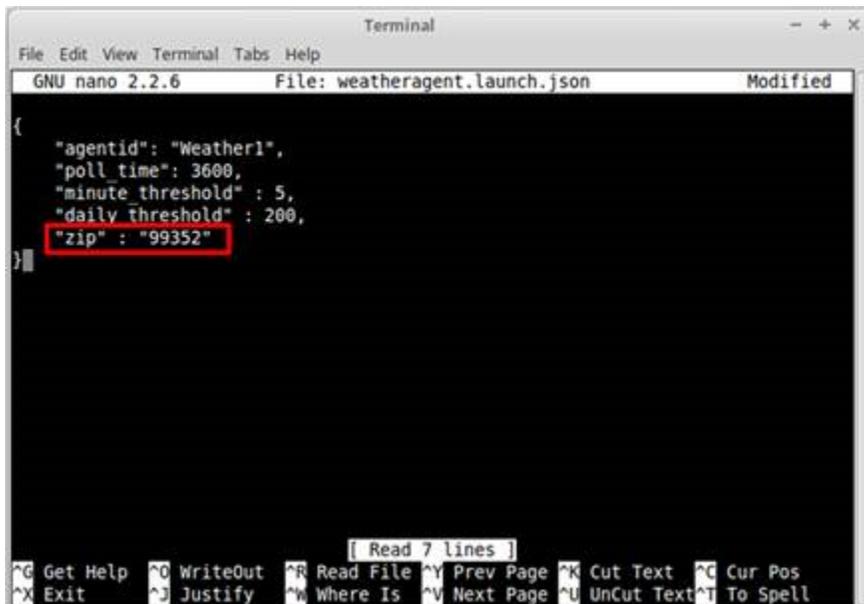


```
Terminal
File Edit View Terminal Tabs Help
GNU nano 2.2.6 File: settings.py Modified
KEY = "cf[redacted]6"
^G Get Help ^O WriteOut ^R Read File ^Y Prev Page ^K Cut Text ^C Cur Pos
^X Exit ^J Justify ^W Where Is ^V Next Page ^U UnCut Text ^T To Spell
```

Figure 5: Entering the Weather Underground Developer Key

- Open the Weather agent’s configuration file and edit the “zip” field, as shown in Figure 6:

```
$ nano WeatherAgent/weatheragent.config
```



```
Terminal
File Edit View Terminal Tabs Help
GNU nano 2.2.6 File: weatheragent.launch.json Modified
{
  "agentid": "Weather1",
  "poll time": 3600,
  "minute threshold" : 5,
  "daily threshold" : 200,
  "zip" : "99352"
}
[ Read 7 lines ]
^G Get Help ^O WriteOut ^R Read File ^Y Prev Page ^K Cut Text ^C Cur Pos
^X Exit ^J Justify ^W Where Is ^V Next Page ^U UnCut Text ^T To Spell
```

Figure 6: Entering Zip Code for Location

Launching the Weather Agent

Create a script modeled on scripts/core/make-listener called make-weather:

```
export SOURCE=services/core/WeatherAgent
export CONFIG=config/weatheragent.config
export TAG=weather
./scripts/core/make-agent.sh

# To set the agent to autostart with the platform, pass "enable"
# to make-agent.sh: ./scripts/core/make-agent.sh enable
```

Then:

```
chmod +x make-weather
```

Now you can run `./make-weather` to stop, remove, build, and reinstall in one script. To start the agent, run: `volttron-ctl start -tag weather`

```
2015-11-12 18:03:46,913 (weatheragent-3.0 2701) <stdout> INFO: {'display_location': {'city': 'West Richland', 'full': 'West Richland, WA', 'magic': '1', 'state_name': 'Washington', 'zip': '99353', 'country': 'US', 'longitude': '-119.36027527', 'state': 'WA', 'wmo': '99999', 'country_iso3166': 'US', 'latitude': '46.30444336', 'elevation': '126.00000000'}, 'local_tz_long': 'America/Los_Angeles', 'observation_location': {'city': 'Orchard Street, West Richland', 'full': 'Orchard Street, West Richland, Washington', 'elevation': '495 ft', 'country': 'US', 'longitude': '-119.335976', 'state': 'Washington', 'country_iso3166': 'US', 'latitude': '46.286118', 'station_id': 'KMAWESTR18'}}
2015-11-12 18:03:46,913 (weatheragent-3.0 2701) <stdout> INFO: {'city': 'West Richland', 'full': 'West Richland, WA', 'magic': '1', 'state_name': 'Washington', 'zip': '99353', 'country': 'US', 'longitude': '-119.36027527', 'state': 'WA', 'wmo': '99999', 'country_iso3166': 'US', 'latitude': '46.30444336', 'elevation': '126.00000000'}
2015-11-12 18:03:46,914 (weatheragent-3.0 2701) <stdout> INFO: {'city': 'Orchard Street, West Richland', 'full': 'Orchard Street, West Richland, Washington', 'elevation': '495 ft', 'country': 'US', 'longitude': '-119.335976', 'state': 'Washington', 'country_iso3166': 'US', 'latitude': '46.286118'}
2015-11-12 18:03:46,914 (weatheragent-3.0 2701) <stdout> INFO: {'local_tz_offset': '-0800', 'local_epoch': '1447380105', 'observation_time': 'Last Updated on November 12, 5:57 PM PST', 'local_tz_short': 'PST', 'observation_epoch': '1447379875', 'local_time_rfc822': 'Thu, 12 Nov 2015 18:01:45 -0800', 'observation_time_rfc822': 'Thu, 12 Nov 2015 17:57:55 -0800'}
2015-11-12 18:03:46,914 (weatheragent-3.0 2701) <stdout> INFO: {'relative_humidity': '59%', 'pressure_mb': '1021', 'pressure_trend': '-'}
2015-11-12 18:03:46,914 (weatheragent-3.0 2701) <stdout> INFO: {'dewpoint_string': '36 F (2 C)', 'precip_1hr_in': '-999.00', 'precip_today_in': '0.00', 'precip_today_metric': '0', 'precip_today_string': '0.00 in (0 mm)', 'dewpoint_f': '36', 'dewpoint_c': '2', 'precip_1hr_string': '-999.00 in (0 mm)', 'precip_1hr_metric': '0'}
2015-11-12 18:03:46,914 (weatheragent-3.0 2701) <stdout> INFO: {'wind_degrees': '270', 'wind_kph': '6.8', 'wind_gust_mph': '0', 'wind_mph': '4.2', 'wind_string': 'From the West at 4.2 MPH', 'pressure_in': '30.15', 'wind_dir': 'West', 'wind_gust_kph': '0'}
```

Figure 7: Example Output from the Weather Agent

Weather Agent Topics

Topics used by the WeatherAgent with example output.

```
[ 'weather/all', {
  "temperature": {
    "windchill_f": "NA",
    "temp_f": 69.1,
    "heat_index_f": "NA",
    "heat_index_string": "NA",
    "temp_c": 20.6,
    "feelslike_c": "20.6",
    "windchill_string": "NA",
    "feelslike_f": "69.1",
    "heat_index_c": "NA",
    "windchill_c": "NA",
    "feelslike_string": "69.1 F (20.6 C)",
    "temperature_string": "69.1 F (20.6 C)",
    "cloud_cover": {
      "visibility_mi": "10.0",
      "solarradiation": "",
      "weather": "Clear",
      "visibility_km": "16.1",
      "UV": "6",
      "location": {
        "display_location": {
          "city": "Richland",
          "full": "Richland, WA",
          "elevation": "121.00000000",
          "state_name": "Washington",
          "zip": "99352",
          "country": "US",
          "longitude": "-119.29721832",
          "state": "WA",
          "country_iso3166": "US",
          "latitude": "46.28490067",
          "local_tz_long": "America/Los_Angeles",
          "observation_location": {
            "city": "Richland, Richland",
            "full": "Richland, Richland, Washington",
            "elevation": "397 ft",
            "country": "US",
            "longitude": "-119.304375",
            "state": "Washington",
            "country_iso3166": "US",
            "latitude": "46.285866",
            "station_id": "KWARICHL21",
            "time": {
              "local_tz_offset": "-0700",
              "local_epoch": "1368724778",
              "observation_time": "Last Updated on May 16, 10:18 AM PDT",
              "local_tz_short": "PDT",
              "observation_epoch": "1368724692",
              "local_time_rfc822": "Thu, 16 May 2013 10:19:38 -0700",
              "observation_time_rfc822": "Thu, 16 May 2013 10:18:12 -0700",
              "pressure_humidity": {
                "relative_humidity": "40%",
                "pressure_mb": "1014",
                "pressure_trend": "-",
                "precipitation": {
                  "dewpoint_string": "44 F (7 C)",
                  "precip_1hr_in": "0.00",
                  "precip_today_in": "0.00",
                  "precip_today_metric": "0",
                  "precip_today_string": "0.00 in (0 mm)",
                  "dewpoint_f": 44,
                  "dewpoint_c": 7,
                  "precip_1hr_string": "0.00 in (0 mm)",
                  "precip_1hr_metric": "0",
                  "wind": {
                    "wind_degrees": 3,
                    "wind_kph": 2.7,
                    "wind_gust_mph": "3.0",
                    "wind_mph": 1.7,
                    "wind_string": "From the North at 1.7 MPH Gusting to 3.0 MPH",
                    "pressure_in": "29.94",
                    "wind_dir": "North",
                    "wind_gust_kph": "4.8"
                  }
                }
              }
            }
          }
        }
      }
    }
  }
}
```

```
[ 'weather/temperature/all', {
  "windchill_f": "NA",
  "temp_f": 69.1,
  "heat_index_f": "NA",
  "heat_index_string": "NA",
  "temp_c": 20.6,
  "feelslike_c": "20.6",
  "windchill_string": "NA",
  "feelslike_f": "69.1",
  "heat_index_c": "NA",
  "windchill_c": "NA",
  "feelslike_string": "69.1 F (20.6 C)",
  "temperature_string": "69.1 F (20.6 C)",
  "cloud_cover": {
    "visibility_mi": "10.0",
    "solarradiation": "",
    "weather": "Clear",
    "visibility_km": "16.1",
    "UV": "6",
    "location": {
      "display_location": {
        "city": "Richland",
        "full": "Richland, WA",
        "elevation": "121.00000000",
        "state_name": "Washington",
        "zip": "99352",
        "country": "US",
        "longitude": "-119.29721832",
        "state": "WA",
        "country_iso3166": "US",
        "latitude": "46.28490067",
        "local_tz_long": "America/Los_Angeles",
        "observation_location": {
          "city": "Richland, Richland",
          "full": "Richland, Richland, Washington",
          "elevation": "397 ft",
          "country": "US",
          "longitude": "-119.304375",
          "state": "Washington",
          "country_iso3166": "US",
          "latitude": "46.285866",
          "station_id": "KWARICHL21",
          "time": {
            "local_tz_offset": "-0700",
            "local_epoch": "1368724778",
            "observation_time": "Last Updated on May 16, 10:18 AM PDT",
            "local_tz_short": "PDT",
            "observation_epoch": "1368724692",
            "local_time_rfc822": "Thu, 16 May 2013 10:19:38 -0700",
            "observation_time_rfc822": "Thu, 16 May 2013 10:18:12 -0700",
            "pressure_humidity": {
              "relative_humidity": "40%",
              "pressure_mb": "1014",
              "pressure_trend": "-",
              "precipitation": {
                "dewpoint_string": "44 F (7 C)",
                "precip_1hr_in": "0.00",
                "precip_today_in": "0.00",
                "precip_today_metric": "0",
                "precip_today_string": "0.00 in (0 mm)",
                "dewpoint_f": 44,
                "dewpoint_c": 7,
                "precip_1hr_string": "0.00 in (0 mm)",
                "precip_1hr_metric": "0",
                "wind": {
                  "wind_degrees": 3,
                  "wind_kph": 2.7,
                  "wind_gust_mph": "3.0",
                  "wind_mph": 1.7,
                  "wind_string": "From the North at 1.7 MPH Gusting to 3.0 MPH",
                  "pressure_in": "29.94",
                  "wind_dir": "North",
                  "wind_gust_kph": "4.8"
                }
              }
            }
          }
        }
      }
    }
  }
}
```

```

“NA”, “windchill_c”: “NA”, “feelslike_string”: “69.1 F (20.6 C)”, “temperature_string”: “69.1 F (20.6 C)”}]
[‘weather/temperature/windchill_f’, ‘NA’]
[‘weather/temperature/temp_f’, ‘69.1’]
[‘weather/temperature/heat_index_f’, ‘NA’]
[‘weather/temperature/heat_index_string’, ‘NA’]
[‘weather/temperature/temp_c’, ‘20.6’]
[‘weather/temperature/feelslike_c’, ‘20.6’]
[‘weather/temperature/windchill_string’, ‘NA’]
[‘weather/temperature/feelslike_f’, ‘69.1’]
[‘weather/temperature/heat_index_c’, ‘NA’]
[‘weather/temperature/windchill_c’, ‘NA’]
[‘weather/temperature/feelslike_string’, ‘69.1 F (20.6 C)’]
[‘weather/temperature/temperature_string’, ‘69.1 F (20.6 C)’]
[‘weather/cloud_cover/all’, ‘{“visibility_mi”: “10.0”, “solarradiation”: “”, “weather”: “Clear”, “visibility_km”:
“16.1”, “UV”: “6”}’]
[‘weather/cloud_cover/visibility_mi’, ‘10.0’]
[‘weather/cloud_cover/solarradiation’, ‘’]
[‘weather/cloud_cover/weather’, ‘Clear’]
[‘weather/cloud_cover/visibility_km’, ‘16.1’]
[‘weather/cloud_cover/UV’, ‘6’]
[‘weather/location/all’, ‘{“display_location”: {“city”: “Richland”, “full”: “Richland, WA”, “elevation”:
“121.00000000”, “state_name”: “Washington”, “zip”: “99352”, “country”: “US”, “longitude”: “-119.29721832”,
“state”: “WA”, “country_iso3166”: “US”, “latitude”: “46.28490067”}, “local_tz_long”: “America/Los_Angeles”,
“observation_location”: {“city”: “Richland, Richland”, “full”: “Richland, Richland, Washington”, “elevation”: “397
ft”, “country”: “US”, “longitude”: “-119.304375”, “state”: “Washington”, “country_iso3166”: “US”, “latitude”:
“46.285866”}, “station_id”: “KWARICHL21”}’]
[‘weather/location/display_location/all’, ‘{“city”: “Richland”, “full”: “Richland, WA”, “elevation”: “121.00000000”,
“state_name”: “Washington”, “zip”: “99352”, “country”: “US”, “longitude”: “-119.29721832”, “state”: “WA”,
“country_iso3166”: “US”, “latitude”: “46.28490067”}’]
[‘weather/location/display_location/city’, ‘Richland’]
[‘weather/location/display_location/full’, ‘Richland, WA’]
[‘weather/location/display_location/elevation’, ‘121.00000000’]
[‘weather/location/display_location/state_name’, ‘Washington’]
[‘weather/location/display_location/zip’, ‘99352’]
[‘weather/location/display_location/country’, ‘US’]
[‘weather/location/display_location/longitude’, ‘-119.29721832’]
[‘weather/location/display_location/state’, ‘WA’]
[‘weather/location/display_location/country_iso3166’, ‘US’]
[‘weather/location/display_location/latitude’, ‘46.28490067’]

```

```
[‘weather/location/local_tz_long’, ‘America/Los_Angeles’]
[‘weather/location/observation_location/all’, ‘{“city”: “Richland, Richland”, “full”: “Richland, Richland, Washington”, “elevation”: “397 ft”, “country”: “US”, “longitude”: “-119.304375”, “state”: “Washington”, “country_iso3166”: “US”, “latitude”: “46.285866”}’]
[‘weather/location/observation_location/city’, ‘Richland, Richland’]
[‘weather/location/observation_location/full’, ‘Richland, Richland, Washington’]
[‘weather/location/observation_location/elevation’, ‘397 ft’]
[‘weather/location/observation_location/country’, ‘US’]
[‘weather/location/observation_location/longitude’, ‘-119.304375’]
[‘weather/location/observation_location/state’, ‘Washington’]
[‘weather/location/observation_location/country_iso3166’, ‘US’]
[‘weather/location/observation_location/latitude’, ‘46.285866’]
[‘weather/location/station_id’, ‘KWARICHL21’]
[‘weather/time/all’, ‘{“local_tz_offset”: “-0700”, “local_epoch”: “1368724778”, “observation_time”: “Last Updated on May 16, 10:18 AM PDT”, “local_tz_short”: “PDT”, “observation_epoch”: “1368724692”, “local_time_rfc822”: “Thu, 16 May 2013 10:19:38 -0700”, “observation_time_rfc822”: “Thu, 16 May 2013 10:18:12 -0700”}’]
[‘weather/time/local_tz_offset’, ‘-0700’]
[‘weather/time/local_epoch’, ‘1368724778’]
[‘weather/time/observation_time’, ‘Last Updated on May 16, 10:18 AM PDT’]
[‘weather/time/local_tz_short’, ‘PDT’]
[‘weather/time/observation_epoch’, ‘1368724692’]
[‘weather/time/local_time_rfc822’, ‘Thu, 16 May 2013 10:19:38 -0700’]
[‘weather/time/observation_time_rfc822’, ‘Thu, 16 May 2013 10:18:12 -0700’]
[‘weather/pressure_humidity/all’, ‘{“relative_humidity”: “40%”, “pressure_mb”: “1014”, “pressure_trend”: “-”}’]
[‘weather/pressure_humidity/relative_humidity’, ‘40%’]
[‘weather/pressure_humidity/pressure_mb’, ‘1014’]
[‘weather/pressure_humidity/pressure_trend’, ‘-’]
[‘weather/precipitation/all’, ‘{“dewpoint_string”: “44 F (7 C)”, “precip_1hr_in”: “0.00”, “precip_today_in”: “0.00”, “precip_today_metric”: “0”, “precip_today_string”: “0.00 in (0 mm)”, “dewpoint_f”: 44, “dewpoint_c”: 7, “precip_1hr_string”: “0.00 in (0 mm)”, “precip_1hr_metric”: ” 0”}’]
[‘weather/precipitation/dewpoint_string’, ‘44 F (7 C)’]
[‘weather/precipitation/precip_1hr_in’, ‘0.00’]
[‘weather/precipitation/precip_today_in’, ‘0.00’]
[‘weather/precipitation/precip_today_metric’, ‘0’]
[‘weather/precipitation/precip_today_string’, ‘0.00 in (0 mm)’]
[‘weather/precipitation/dewpoint_f’, ‘44’]
[‘weather/precipitation/dewpoint_c’, ‘7’]
[‘weather/precipitation/precip_1hr_string’, ‘0.00 in (0 mm)’]
```

```
['weather/precipitation/precip_1hr_metric', '0']
```

```
['weather/wind/all', {'wind_degrees': 3, "wind_kph": 2.7, "wind_gust_mph": "3.0", "wind_mph": 1.7, "wind_string": "From the North at 1.7 MPH Gusting to 3.0 MPH", "pressure_in": "29.94", "wind_dir": "North", "wind_gust_kph": "4.8"}]
```

```
['weather/wind/wind_degrees', '3']
```

```
['weather/wind/wind_kph', '2.7']
```

```
['weather/wind/wind_gust_mph', '3.0']
```

```
['weather/wind/wind_mph', '1.7']
```

```
['weather/wind/wind_string', 'From the North at 1.7 MPH Gusting to 3.0 MPH']
```

```
['weather/wind/pressure_in', '29.94']
```

```
['weather/wind/wind_dir', 'North']
```

```
['weather/wind/wind_gust_kph', '4.8']
```

3.5.2 Base Platform Functionality

The base platform functionality focuses on the agent lifecycle, management of the platform itself, and security.

Used Environmental Variables

- **AGENT_VIP_IDENTITY** - The router address an agent will attempt to connect to.
- **AGENT_CONFIG** - The path to a configuration file to use during agent launch.
- **VOLTRON_HOME** - The home directory where the voltron instances is located.

Agent Autostart

An agent can be setup to start when the platform is started with the “enable” command. This command also allows a priority to be set (0-100, default 50) so that agents can be started after any dependencies. This command can also be used with the `-tag` or `-name` options.

```
voltron-ctl enable <AGENT_UUID> <PRIORITY>
```

Agent Lifecycle Management

The VOLTRON platform has several commands for controlling the lifecycle of agents. This page discusses how to use them, for details of operation please see *PlatformConfiguration*

These examples assume the VOLTRON environment has been activated (. env/bin/activate). If not, add “bin/” to all commands.

Agent Packaging

The “voltron-pkg” command is used for packaging and configuring agents. It is not necessary to have the platform running to use this command. The platform uses [Python Wheel](#) for its packaging and follows the [Wheel naming convention](#).

To create an agent package, call `voltron-pkg <Agent Dir>`.

For instance: `voltttron-pkg package examples/ListenerAgent`

The `package` command uses the `setup.py` in the agent directory to create the package. The name and version number portion of the Wheel filename come from this. The resulting wheels are created at “`~/voltttron/packaged`”.

For example: `~/voltttron/packaged/listeneragent-3.0-py2-none-any.whl`.

Agent Configuration

Agent packages are configured with the `voltttron-pkg configure <AgentPackage> <ConfigFile>` command. It is suggested that this file use json formatting but the agent can be written to interpret any format it requires. The configuration of a particular agent is opaque to the VOLTTRON platform. The location of the agent config file is passed as an environmental variable “`AGENT_CONFIG`” which the provided utilities read in and pass to the agent.

An example config file passing in some parameters:

```
{
    "agentid": "listener1",
    "message": "hello"
}
```

Agent Installation and Removal

Agents are installed into the platform using:

`voltttron-ctl install <package>`. When agents are installed onto a platform, it creates a uuid for that instance of an agent. This allows multiple instances of the same agent package to be installed on the platform.

Agents can also be installed with a *tag* by using:

```
voltttron-ctl install <TAG>=<PACKAGE>
```

This allows the user to refer to the agent with “`--tag`” instead of the uuid when issuing commands. This tag can also distinguish instances of an agent from each other.

A stopped agent can be removed with:

- `voltttron-ctl remove <AGENT_UUID>`
- `voltttron-ctl remove --tag <AGENT_TAG>`
- `voltttron-ctl remove --name <AGENT_NAME>`

Removal by tag and name potentially allows multiple agents to be removed at once and should be used with caution. A “`-f`” option is required to delete more than one agent at a time.

Agent Control

Starting and Stopping an Agent

Agent that are installed in the platform can be launched with the “`start`” command. By default this operates off the agent’s UUID but can be used with “`--tag`” or “`--name`” to launch agents by those attributes. This can allow multiple agents to be started at once. For instance: `voltttron-ctl start --name myagent-0.1` would start all

instances of that agent regardless of their uuid, tag, or configuration information. After an agent is started, it will show up in *AgentStatus* as “running” with a process id.

Similarly, `voltttron-ctl stop <UUID>` can also operate off the tag and name of agent(s). After an agent is stopped, it will show an exit code of 0 in *AgentStatus*

Running an agent

For testing purposes, an agent package not installed in the platform can be run by using: `voltttron-ctl run <PACKAGE>`.

Agent Status

`voltttron-ctl list` lists the agents installed on the platform and

their priority | The `voltttron-ctl status` shows the list of installed agents and whether they are running or have exited. | See *AgentStatus* for more details.

Agent List Display

AGENT	IDENTITY	TAG	PRI
d listeneragent-3.0	listeneragent-3.0_1		30
2 testeragent-0.1	testeragent-0.1_1		

`voltttron-ctl list` shows the agents which have been installed on the platform along with their uuid, associated tag, and priority.

- **uuid** is the first column of the display and is displayed as the shorted unique portion. Using this portion, agents can be started, stopped, removed, etc.
- **AGENT** is the “name” of this agent based on the name of the wheel file which was installed. Agents can be controlled with this using “-name “. Note, if multiple instances of a wheel are installed they will all have the same name and can be controlled as a group.
- **TAG** is a user provided tag which makes it simpler to track and refer to agents. Using “-tag ” agents can be controlled using this
- **PRI** is the priority for agents which have been “enabled” using the `voltttron-ctl enable` command. When enabled, agents will be automatically started in priority order along with the platform.

Agent Status Display

AGENT	TAG	STATUS
d listeneragent-3.0	listener	running [3813]
2 testeragent-0.1		0

`voltttron-ctl statu` shows a list of all agents installed on the platform and their current status.

- **uuid** is the first column of the display and is displayed as the shorted unique portion. Using this portion, agents can be started, stopped, removed, etc.

- **AGENT** is the “name” of this agent based on the name of the wheel file which was installed. Agents can be controlled with this using “-name “. Note, if multiple instances of a wheel are installed they will all have the same name and can be controlled as a group.
- **TAG** is a user provided tag which makes it simpler to track and refer to agents. Using “-tag ” agents can be controlled using this
- **STATUS** is the current condition of the agent. If the agent is currently executing, it has “running” and the process id of the agent. If the agent is not running, the exit code is shown.

Tagging Agents

Agents can be tagged as they are installed with:

```
volttron-ctl install <TAG>=<AGENT_PACKAGE>
```

Agents can be tagged after installation with:

```
volttron-ctl tag <AGENT_UUID> <TAG>
```

Agents can be “tagged” to provide a meaningful user defined way to reference the agent instead of the uuid or the name. This allows users to differentiate between instances of agents which use the same codebase but are configured differently. For instance, the AFDDAgent can be configured to work against a single HVAC unit and can have any number of instances running on one platform. A tagging scheme for this could be by unit: afdd-rtu1, afdd-rtu2, etc.

Commands which operate off an agent’s UUID can optionally operate off the tag by using “-tag “. This can use wildcards to catch multiple agents at once.

Authentication Commands

All authentication sub-commands can be viewed by entering following command.

```
volttron-ctl auth --help
```

```
optional arguments:
-h, --help            show this help message and exit
-c FILE, --config FILE
                        read configuration from FILE
--debug                show tracebacks for errors rather than a brief message
-t SECS, --timeout SECS
                        timeout in seconds for remote calls (default: 30)
--vip-address ZMQADDR
                        ZeroMQ URL to bind for VIP connections
--keystore-file FILE  use keystore from FILE
--known-hosts-file FILE
                        get known-host server keys from FILE

subcommands:
  add                  add new authentication record
  add-group            associate a group name with a set of roles
  add-known-host       add server public key to known-hosts file
  add-role             associate a role name with a set of capabilities
  keypair              generate CurveMQ keys for encrypting VIP connections
  list                 list authentication records
  list-groups          show list of group names and their sets of roles
  list-known-hosts    list entries from known-hosts file
  list-roles           show list of role names and their sets of capabilities
```

(continues on next page)

(continued from previous page)

publickey	show public key for each agent
remove	removes one or more authentication records by indices
remove-group	disassociate a group name from a set of roles
remove-known-host	remove entry from known-hosts file
remove-role	disassociate a role name from a set of capabilities
serverkey	show the serverkey for the instance
update	updates one authentication record by index
update-group	update group to include (or remove) given roles
update-role	update role to include (or remove) given capabilities

Authentication record

An authentication record consist of following parameters

```
domain []:
address []: Either a single agent identity or an array of agents identities
user_id []: Arbitrary string to indentify the agent
capabilities (delimit multiple entries with comma) []: Array of strings referring to
→authorized capabilities defined by exported RPC methods
roles (delimit multiple entries with comma) []:
groups (delimit multiple entries with comma) []:
mechanism [CURVE]:
credentials []: Public key string for the agent
comments []:
enabled [True]:
```

For more details on how to create authentication record, please see section [Agent Authentication](#)

Platform Commands

VOLTRON files for a platform instance are stored under a single directory known as the VOLTRON home. This home directory is set via the VOLTRON_HOME environment variable and defaults to ~/.volttron. Multiple instances of the platform may exist under the same account on a system by setting the VOLTRON_HOME environment variable appropriately before executing VOLTRON commands.

Configuration files use a modified INI format where section names are command names for which the settings in the section apply. Settings before the first section are considered global and will be used by all commands for which the settings are valid. Settings keys are long options (with or without the opening -) and are followed by a colon (:) or equal (=) and then the value. Boolean options need not include the separator or value, but may specify a value of 1, yes, or true for true or 0, no, or false for false.

A default configuration file, \$VOLTRON_HOME/config, may be created to override default options. If it exists, it will be automatically parsed before all other command-line options. To skip parsing the default configuration file, either move the file out of the way or set the SKIP_VOLTRON_CONFIG environment variable.

All commands and sub-commands have help available with “-h” or “-help”. Additional configuration files may be specified with “-c” or “-config”. To specify a log file, use “-l” or “-log”.

```
env/bin/volttron -c config.ini -l volttron.log
```

Full options:

VOLTRON platform service

optional arguments:

```
-c FILE, --config FILE          read configuration from FILE
-l FILE, --log FILE             send log output to FILE instead of stderr
-L FILE, --log-config FILE     read logging configuration from FILE
--log-level LOGGER:LEVEL      override default logger logging level
--monitor                      monitor and log connections (implies -v)
--msgdebug                     publish all messages to a socket for debugging
                               purposes; used with MessageDebuggerAgent
-q, --quiet                    decrease logger verbosity; may be used multiple
                               times
-v, --verbose                  increase logger verbosity; may be used multiple
                               times
--verbosity LEVEL             set logger verbosity
-h, --help                    show this help message and exit
--version                      show program's version number and exit
```

agent options:

```
--autostart                    automatically start enabled agents and services
--publish-address ZMQADDR      ZeroMQ URL used for pre-3.x agent publishing
                               (deprecated)
--subscribe-address ZMQADDR   ZeroMQ URL used for pre-3.x agent subscriptions
                               (deprecated)
--vip-address ZMQADDR         ZeroMQ URL to bind for VIP connections
--vip-local-address ZMQADDR   ZeroMQ URL to bind for local agent VIP connections
--bind-web-address BINDWEBADDR Bind a web server to the specified ip:port passed
--volttron-central-address VOLTRON_CENTRAL_ADDRESS The web address of a volttron central install
                               instance.
--volttron-central-serverkey VOLTRON_CENTRAL_SERVERKEY The serverkey of volttron central.
--instance-name INSTANCE_NAME The name of the instance that will be reported to
                               VOLTRON central.
```

Boolean options, which take no argument, may be inverted by prefixing the option with no- (e.g. --autostart may be inverted using --no-autostart).

voltron-ctl Commands

voltron-ctl is used to issue commands to the platform from the command line. Through voltron-ctl it is possible to install and removed agents, start and stop agents, manage the configuration store, get the platform status, and shutdown the platform.

Warning: `voltron-ctl` creates a special temporary agent to communicate with the platform with a specific VIP IDENTITY, thus multiple instances of `voltron-ctl` cannot run at the same time. Attempting to do so will result in a conflicting identity error.

```
usage: voltron-ctl command [OPTIONS] ...
```

Manage and control VOLTRON agents.

commands:

<code>install</code>	install agent from wheel
<code>tag</code>	set, show, or remove agent tag
<code>remove</code>	remove agent
<code>list</code>	list installed agent
<code>status</code>	show status of agents
<code>clear</code>	clear status of defunct agents
<code>enable</code>	enable agent to start automatically
<code>disable</code>	prevent agent from start automatically
<code>start</code>	start installed agent
<code>stop</code>	stop agent
<code>restart</code>	restart agent
<code>run</code>	start any agent by path
<code>auth</code>	manage authorization entries and encryption keys
<code>config</code>	manage the platform configuration store
<code>shutdown</code>	stop all agents
<code>send</code>	send agent and start on a remote platform
<code>stats</code>	manage router message statistics tracking

voltron-ctl auth subcommands

subcommands:

<code>add</code>	add new authentication record
<code>add-known-host</code>	add server public key to known-hosts file
<code>keypair</code>	generate CurveMQ keys for encrypting VIP connections
<code>list</code>	list authentication records
<code>publickey</code>	show public key for each agent
<code>remove</code>	removes one or more authentication records by indices
<code>serverkey</code>	show the serverkey for the instance
<code>update</code>	updates one authentication record by index

voltron-ctl config subcommands

subcommands:

<code>store</code>	store a configuration
<code>delete</code>	delete a configuration
<code>list</code>	list stores or configurations in a store
<code>get</code>	get the contents of a configuration

voltron-pkg Commands

```
usage: voltron-pkg [-h] [-l FILE] [-L FILE] [-q] [-v] [--verbooseness LEVEL]
                  {package, repackage, configure} ...

optional arguments:
  -h, --help            show this help message and exit

subcommands:
  valid subcommands

  {package, repackage, configure}
                        additional help
  package              Create agent package (whl) from a directory or
                        installed agent name.
  repackage            Creates agent package from a currently installed
                        agent.
  configure            add a configuration file to an agent package
```

voltron-pkg commands (with Volttron Restricted package installed and enabled):

```
usage: voltron-pkg [-h] [-l FILE] [-L FILE] [-q] [-v] [--verbooseness LEVEL]
                  {package, repackage, configure, create_ca, create_cert, sign, verify}
                  ...

VOLTRON packaging and signing utility

optional arguments:
  -h, --help            show this help message and exit
  -l FILE, --log FILE  send log output to FILE instead of stderr
  -L FILE, --log-config FILE
                        read logging configuration from FILE
  -q, --quiet          decrease logger verbosity; may be used multiple
                        times
  -v, --verbose        increase logger verbosity; may be used multiple
                        times
  --verbooseness LEVEL
                        set logger verbosity

subcommands:
  valid subcommands

  {package, repackage, configure, create_ca, create_cert, sign, verify}
                        additional help
  package              Create agent package (whl) from a directory or
                        installed agent name.
  repackage            Creates agent package from a currently installed
                        agent.
  configure            add a configuration file to an agent package
  sign                sign a package
  verify              verify an agent package
```

voltron-cfg Commands

voltron-cfg is a tool aimed at making it easier to get up and running with Volttron and a handful of agents. Running the tool without any arguments will start a *wizard* with a walk through for setting up instance configuration options and available agents. If only individual agents need to be configured they can be listed at the command line.

```
usage: volttron-cfg [-h] [--list-agents | --agent AGENT [AGENT ...]]

optional arguments:
  -h, --help            show this help message and exit
  --list-agents         list configurable agents
                        listener
                        platform_historian
                        vc
                        vcp
  --agent AGENT [AGENT ...]
                        configure listed agents
```

VOLTRON Config File

The VOLTRON platform config file can contain any of the command line arguments for starting the platform...

```
-c FILE, --config FILE
                        read configuration from FILE
-l FILE, --log FILE    send log output to FILE instead of stderr
-L FILE, --log-config FILE
                        read logging configuration from FILE
-q, --quiet            decrease logger verbosity; may be used multiple
                        times
-v, --verbose          increase logger verbosity; may be used multiple
                        times
--verbosity LEVEL     set logger verbosity
--help                show this help message and exit
--version              show program's version number and exit
```

agent options:

```
--autostart            automatically start enabled agents and services
--publish-address ZMQADDR
                        ZeroMQ URL for used for agent publishing
--subscribe-address ZMQADDR
                        ZeroMQ URL for used for agent subscriptions
```

control options:

```
--control-socket FILE
                        path to socket used for control messages
--allow-root           allow root to connect to control socket
--allow-users LIST    users allowed to connect to control socket
--allow-groups LIST   user groups allowed to connect to control socket
```

Boolean options, which take no argument, may be inverted by prefixing

the `l` option with `no-` (e.g. `--autostart` may be inverted using `--no-autostart`).

VOLTRON Environment

By default, the VOLTRON projects bases its files out of `VOLTRON_HOME` which defaults to “`~/voltron`”.

- \$VOLTTTRON_HOME/agents contains the agents installed on the platform
- \$VOLTTTRON_HOME/certificates contains the certificates for use with the Licensed VOLTTRON code.
- \$VOLTTTRON_HOME/run contains files create by the platform during execution. The main ones are the 0MQ files created for publish and subscribe.
- \$VOLTTTRON_HOME/ssh keys used by agent mobility in the Licensed VOLTTRON code
- \$VOLTTTRON_HOME/config Default location to place a config file to override any platform settings.
- \$VOLTTTRON_HOME/packaged is where agent packages created with 'volttron-pkg package are created

VOLTTRON Config

The new volttron-cfg commands allows for the easy configuration of a VOLTTRON platform. This includes setting up the platform configuration, historian, VOLTTRON Central UI, and platform agent.

example volttron-cfg output:

Note:

- In this example, <user> represents the user's home directory, and <localhost> represents the machine's localhost.
 - The platform has been bootstrapped with rabbitmq enabled e.g. (python bootstrap.py -rabbitmq)
-

```
Your VOLTTRON_HOME currently set to: /home/<user>/.volttron

Is this the volttron you are attempting to setup? [Y]: y
What type of message bus (rmq/zmq)? [zmq]: rmq

The rmq message bus has a backward compatibility
layer with current zmq instances. What is the
zmq bus's vip address? [tcp://127.0.0.1]:
What is the port for the vip address? [22916]:
Is this instance web enabled? [N]: y
What is the hostname for this instance? (https) [https://<localhost>]:
What is the port for this instance? [8443]:
Is this an instance of volttron central? [N]: y
Configuring /home/<user>/volttron/services/core/VolttronCentral.
Enter volttron central admin user name: admin
Enter volttron central admin password:
Retype password:
Installing volttron central.
Should the agent autostart? [N]: y
Will this instance be controlled by volttron central? [Y]:
Configuring /home/<user>/volttron/services/core/VolttronCentralPlatform.
What is the name of this instance? [volttron1]:
What is the hostname for volttron central? [https://<localhost>]:
What is the port for volttron central? [8443]:
Should the agent autostart? [N]: y
Would you like to install a platform historian? [N]: y
Configuring /home/<user>/volttron/services/core/SQLHistorian.
Should the agent autostart? [N]: y
Would you like to install a master driver? [N]: y
Configuring /home/<user>/volttron/services/core/MasterDriverAgent.
Would you like to install a fake device on the master driver? [N]: y
Should the agent autostart? [N]: y
```

(continues on next page)

(continued from previous page)

```

Would you like to install a listener agent? [N]: y
Configuring examples/ListenerAgent.
Should the agent autostart? [N]: y
Finished configuration!

You can now start the volttron instance.

If you need to change the instance configuration you can edit
the config file is at /home/<user>/.volttron/config

```

3.5.3 VOLTRON Historian Framework

Historian Agents are the way by which device, actuator, datalogger, and analysis are captured and stored in some sort of data store. Historians exist for the following storage options:

- A general *SQL Historian* implemented for MySQL and SQLite
- *MongoDB Historian*
- *Crate Historian*
- *Forward Historian* for sending data to another VOLTRON instance
- *OpenEIS Historian*
- *MQTT Historian* Forwards data to an MQTT broker
- *InfluxDB Historian*

Other implementations of historians can be created by following the *developing historian agents* section of the wiki.

Historians are all built upon the BaseHistorian which provides general functionality the specific implementations is built upon.

In most cases the default settings are fine for all deployments.

All historians support the following settings:

```

{
  # Maximum amount of time to wait before retrying a failed publish in seconds.
  # Will try more frequently if new data arrives before this timeline expires.
  # Defaults to 300
  "retry_period": 300.0,

  # Maximum number of records to submit to the historian at a time.
  # Defaults to 1000
  "submit_size_limit": 1000,

  # In the case where a historian needs to catch up after a disconnect
  # the maximum amount of time to spend writing to the database before
  # checking for and caching new data.
  # Defaults to 30
  "max_time_publishing": 30.0,

  # Limit how far back the historian will keep data in days.
  # Partial days supported via floating point numbers.
  # A historian must implement this feature for it to be enforced.
  "history_limit_days": 366,

```

(continues on next page)

(continued from previous page)

```

# Limit the size of the historian data store in gigabytes.
# A historian must implement this feature for it to be enforced.
"storage_limit_gb": 2.5

# Size limit of the backup cache in Gigabytes.
# Defaults to no limit.
"backup_storage_limit_gb": 8.0,

# Do not actually gather any data. Historian is query only.
"readonly": false,

# capture_device_data
# Defaults to true. Capture data published on the `devices/` topic.
"capture_device_data": true,

# capture_analysis_data
# Defaults to true. Capture data published on the `analysis/` topic.
"capture_analysis_data": true,

# capture_log_data
# Defaults to true. Capture data published on the `datalogger/` topic.
"capture_log_data": true,

# capture_record_data
# Defaults to true. Capture data published on the `record/` topic.
"capture_record_data": true,

# Replace a one topic with another before saving to the database.
# Deprecated in favor of retrieving the list of
# replacements from the VCP on the current instance.
"topic_replace_list": [
  {"from": "FromString", "to": "ToString"}
],

# For historian developers. Adds benchmarking information to gathered data.
# Defaults to false and should be left that way.
"gather_timing_data": false

# Allow for the custom topics or for limiting topics picked up by a historian_
↪instance.
# the key for each entry in custom topics is the data handler. The topic and_
↪data must
# conform to the syntax the handler expects (e.g., the capture_device_data_
↪handler expects
# data the driver framework). Handlers that expect specific data format are
# capture_device_data, capture_log_data, and capture_analysis_data. All other_
↪handlers will be
# treated as record data. The list associated with the handler is a list of custom
# topics to be associated with that handler.
#
# To restrict collection to only the custom topics, set the following config_
↪variables to False
# capture_device_data
# capture_analysis_data
# capture_log_data
# capture_record_data
"custom_topics": {

```

(continues on next page)

(continued from previous page)

```

    "capture_device_data": ["devices/campus/building/device/all"],
    "capture_analysis_data": ["analysis/application_data/example"],
    "capture_record_data": ["example"]
  }
}

```

By default the base historian will listen to 4 separate root topics *datalogger/**, *record/**, *analysis/**, and *device/**.

Each of root topics has a *specific message syntax* that it is expecting for incoming data.

Messages published to *datalogger* will be assumed to be timepoint data that is composed of units and specific types with the assumption that they have the ability to be graphed easily.

Messages published to *devices* are data that comes directly from drivers.

Messages published to *analysis* are analysis data published by agnets in the form of key value pairs.

Finally Messages that are published to *record* will be handled as string data and can be customized to the user specific situation.

Please consult the *Historian Topic Syntax* page for a specific syntax.

This base historian will cache all received messages to a local database before publishing it to the historian. This allows recovery from unexpected happenings before the successful writing of data to the historian.

Analytics Historian

An Analytics Historian (OpenEIS) has been developed to integrate real time data ingestion into the OpenEIS platform. In order for the OpenEIS historian to be able to communicate with an OpenEIS server a datasource must be created on the OpenEIS server. The process of creating a dataset is documented in the [OpenEIS User's Guide](#) under *Creating a Dataset* heading. Once a dataset is created you will be able to add datasets through the configuration file. An example configuration for the historian is as follows:

```

{
  # The agent id is used for display in volttron central.
  "agentid": "openeishistorian",
  # The vip identity to use with this historian.
  # should not be a platform.historian!
  #
  # Default value is un referenced because it listens specifically to the bus.
  #"identity": "openeis.historian",

  # Require connection section for all historians. The openeis historian
  # requires a url for the openis server and login credentials for publishing
  # to the correct user's dataset.
  "connection": {
    "type": "openeis",
    "params": {
      # The server that is running openeis
      # the rest path for the dataset is dataset/append/{id}
      # and will be populated from the topic_dataset list below.
      "uri": "http://localhost:8000",

      # Openeis requires a username/password combination in order to
      # login to the site via rest or the ui.
      #
      "login": "volttron",

```

(continues on next page)

```

        "password": "volttron"
    },
    # All datasets that are going to be recorded by this historian need to be
    # defined here.
    #
    # A dataset definition consists of the following parts
    #   "ds1": {
    #
    #       The dataset id that was created in openeis.
    #       "dataset_id": 1,
    #
    #       Setting to 1 allows only the caching of data that actually meets
    #       the mapped point criteria for this dataset.
    #       Defaults to 0
    #       "ignore_unmapped_points": 0,
    #
    #       An ordered list of points that are to be posted to openeis. The
    #       points must contain a key specifying the incoming topic with the
    #       value an openeis schema point:
    #       [
    #           {"rtu4/OutsideAirTemp": "campus1/building1/rtu4/OutdoorAirTemperature
↪"}
    #       ]
    #   },
    "dataset_definitions": {
        "ds1": {
            "dataset_id": 1,
            "ignore_unmapped_points": 0,
            "points": [
                {"campus1/building1/OutsideAirTemp": "campus1/building1/
↪OutdoorAirTemperature"},
                {"campus1/building1/HVACStatus": "campus1/building1/HVACStatus"},
                {"campus1/building1/CompressorStatus": "campus1/building1/
↪LightingStatus"}
            ]
        }
    },
    #,
    # "ds2": {
    #   "id": 2,
    #   "points": [
    #     "rtu4/OutsideAirTemp",
    #     "rtu4/MixedAirTemp"
    #   ]
    # }
}

```

Crate Historian

Crate is an open source SQL database designed on top of a No-SQL design. It allows automatic data replication and self-healing clusters for high availability, automatic sharding, and fast joins, aggregations and sub-selects.

Find out more about crate from <https://crate.io/>.

Prerequisites

1. Crate Database

For Arch Linux, Debian, RedHat Enterprise Linux and Ubuntu distributions there is a dead simple installer to get crate up and running on your system.

```
sudo bash -c "$(curl -L install.crate.io)"
```

This command will download and install all of the requirements for running crate, create a crate user and install a crate service. After the installation the service will be available for viewing at <http://localhost:4200> by default.

Note: There is no authentication support within crate.

2. Crate Driver

There is a python library for crate that must be installed in the volttron python environment in order to access crate. From an activated environment, in the root of the volttron folder, execute the following command:

```
python bootstrap.py --crate
```

or

```
pip install crate
```

Configuration

Because there is no authorization to access a crate database the configuration for the CrateHistorian is very easy.

```
{
  "connection": {
    "type": "crate",
    # Optional table prefix defaults to historian
    "schema": "testing",
    "params": {
      "host": "localhost:4200"
    }
  }
}
```

Finally package, install and start the CrateHistorian agent.

See also:

Agent Development Walkthrough

DataMover Historian

DataMover sends data from its platform to a remote platform in cases where there is not sufficient resources to store locally. It shares this functionality with the Forward Historian, However DataMover does not have a goal of data appearing “live” on the remote platform. This allows DataMover to be more efficient by both batching data and

by sending an RCP call to a remote historian instead of publishing data on the remote message bus. This allows DataMover to be more robust by ensuring that the receiving historian is running. If the target is unreachable, DataMover will cache data until it is available.

Configuration

The default configuration file is `services/core/DataMover/config`. Change **destination-vip** to point towards the foreign Volttron instance.

```
{
  "destination-vip": "ipc://@/home/volttron/.volttron/run/vip.socket",
  "destination-serverkey": null,
  "required_target_agents": [],
  "custom_topic_list": [],
  "services_topic_list": [
    "devices", "analysis", "record", "datalogger", "actuators"
  ],
  "topic_replace_list": [
    #{"from": "FromString", "to": "ToString"}
  ]
}
```

The **services_topic_list** allows you to specify which of the main data topics to forward. If there is no entry, the historian defaults to sending all.

topic_replace_list allows you to replace portions of topics if needed. This could be used to correct or standardize topics or to replace building/device names with an anonymized version. The receiving platform will only see the replaced values.

Adding the configuration option below will limit the backup cache to *n* gigabytes. This will keep your hard drive from filling up if the agent is disconnected from its target for a long time.

```
"backup_storage_limit_gb": n
```

See Also

[Historians Historians](#)

Forward Historian

The primary use case for the ForwardHistorian is to send data to another instance of VOLTTRON as if the data were live. This allows agents running on a more secure and/or more powerful machine to run analysis on data being collected on a potentially less secure/powerful board.

Given this use case, it is not optimized for batching large amounts of data when liveness is not needed. For this use case, please see the *DataMover Historian*.

The forward historian can be found in the `services/core` directory.

Configuration

The default configuration file is `services/core/ForwardHistorian/config`. Change **destination-vip** to point towards the foreign Volttron instance.

```
{
  "agentid": "forwarder",
  "destination-vip": "ipc://@/home/voltron/.voltron/run/vip.socket"
}
```

In order to send to a remote platform, you will need its VIP address and server key. The server key can be found by running

```
voltron-ctl auth serverkey
```

Put the result into the following example (Note the example uses a local IP address)

```
{
  "agentid": "forwarder",
  "destination-vip": "tcp://127.0.0.1:22916",
  "destination-serverkey": "<SOME_KEY>"
}
```

Adding the configuration option below will limit the backup cache to *n* gigabytes. This will keep your hard drive from filling up if the agent is disconnected from its target for a long time.

```
"backup_storage_limit_gb": n
```

See Also

Historians Historians

Historian Topic Syntax

Each historian will subscribe to the following message bus topics (*datalogger/*, *analysis/*, *record/** and *devices/**). For each of these topics there is a different message syntax that must be adhered to in order for the correct interpretation of the data being specified.

record/*

The record topic is the most flexible of all of the topics. This topic allows any serializable message to be published to any topic under the root topic `/record`.

Note: this topic is not recommended to plot, as the structure of the messages are not necessarily numeric.

```
# Example messages that can be published

# Dictionary data
{'foo': 'world'}

# Numerical data
52

# Time data (note not a datetime object)
'2015-12-02T11:06:32.252626'
```

devices/*

The devices topic is meant to be data structured from a scraping of a ModBus or BacNet device. Currently drivers for both of these protocols write data to the message bus in the proper format. VOLTRON drivers also publish an aggregation of points in an “all” topic. **Only the “all” topic messages are read and published to a historian.** Both the all topic and point topic have the same header information, but the message body for each is slightly different. For a complete working example of these messages please see

- `examples.ExampleSubscriber.subscriber.subscriber_agent`

Format of header and message for device topics (i.e. messages published to topics with pattern “devices/*/all”):

```
# Header contains the data associated with the message.
{
  # python code to get this is
  # from datetime import datetime
  # from volttron.platform.messaging import headers as header_mod
  # from volttron.platform.agent import utils
  # now = utils.format_timestamp( datetime.utcnow())
  # {
  #   headers_mod.DATE: now,
  #   headers_mod.TIMESTAMP: now
  # }
  "Date": "2015-11-17 21:24:10.189393+00:00",
  "TimeStamp": "2015-11-17 21:24:10.189393+00:00"
}

# Message Format:

# WITH METADATA
# Messages contains a two element list. The first element contains a
# dictionary of all points under a specific parent. While the second
# element contains a dictionary of meta data for each of the specified
# points. For example devices/pnml/building/OutsideAirTemperature and
# devices/pnml/building/MixedAirTemperature ALL message would be created as:
[
  {"OutsideAirTemperature ": 52.5, "MixedAirTemperature ": 58.5},
  {
    "OutsideAirTemperature ": {'units': 'F', 'tz': 'UTC', 'type': 'float'},
    "MixedAirTemperature ": {'units': 'F', 'tz': 'UTC', 'type': 'float'}
  }
]

#WITHOUT METADATA
# Message contains a dictionary of all points under a specific parent
{"OutsideAirTemperature ": 52.5, "MixedAirTemperature ": 58.5}
```

analysis/*

Data sent to analysis/* topics is result of analysis done by applications. The format of data sent to analysis/* topics is similar to data sent to devices/*/all topics.

datalogger/*

Messages published to datalogger will be assumed to be time point data that is composed of units and specific types with the assumption that they have the ability to be graphed easily.

```
{'MixedAirTemperature': {'Readings': ['2015-12-02T00:00:00',
                                     mixed_reading],
                          'Units': 'F',
                          'tz': 'UTC',
                          'data_type': 'float'}}
```

If no datetime value is specified as a part of the reading, current time is used. Message can be published without any header. In the above message 'Readings' and 'Units' are mandatory

Influxdb Historian

InfluxDB is an open source time series database with a fast, scalable engine and high availability. It's often used to build DevOps Monitoring (Infrastructure Monitoring, Application Monitoring, Cloud Monitoring), IoT Monitoring, and Real-Time Analytics solutions.

More information about InfluxDB is available from <https://www.influxdata.com/>.

Prerequisites

InfluxDB Installation

To install InfluxDB on an Ubuntu or Debian operating system, run the script:

```
services/core/InfluxdbHistorian/scripts/install-influx.sh
```

For installation on other operating systems, see <https://docs.influxdata.com/influxdb/v1.4/introduction/installation/>.

Authentication in InfluxDB

By default, the InfluxDB *Authentication* option is disabled, and no user authentication is required to access any InfluxDB database. You can enable authentication by updating the InfluxDB configuration file. For detailed information on enabling authentication, see: https://docs.influxdata.com/influxdb/v1.4/query_language/authentication_and_authorization/.

If *Authentication* is enabled, authorization privileges are enforced. There must be at least one defined admin user with access to administrative queries as outlined in the linked document above. Additionally, you must pre-create the user and database that are specified in the configuration file (the default configuration file for InfluxDB is `services/core/InfluxdbHistorian/config`). If your user is a non-admin user, they must be granted a full set of privileges on the desired database.

InfluxDB Driver

In order to connect to an InfluxDb client, the Python library for InfluxDB must be installed in VOLTRON's virtual environment. From the command line, after enabling the virtual environment, install the InfluxDB library as follows:

```
pip install influxdb
```

Configuration

The default configuration file for VOLTRON's InfluxDBHistorian agent should be in the format:

```
{
  "connection": {
    "params": {
      "host": "localhost",
      "port": 8086,          # Don't change this unless default bind port
                          # in influxdb config is changed
      "database": "historian",
      "user": "historian", # user is optional if authentication is turned off
      "passwd": "historian" # passwd is optional if authentication is turned off
    }
  },
  "aggregations": {
    "use_calendar_time_periods": true
  }
}
```

The InfluxDBHistorian agent can be packaged, installed and started according to the standard VOLTRON agent creation procedure. A sample VOLTRON configuration file has been provided: `services/core/InfluxdbHistorian/config`.

See also:

[Agent Development Walkthrough](#)

Connection

The `host`, `database`, `user` and `passwd` values in the VOLTRON configuration file can be modified. `user` and `passwd` are optional if InfluxDB *Authentication* is disabled.

Note: Be sure to initialize or pre-create the database and user that you defined in the configuration file, and if user is a non-admin user, be make sure to grant privileges for the user on the specified database. For more information, see [Authentication in InfluxDB](#).

Aggregations

In order to use aggregations, the VOLTRON configuration file must also specify a value, either `true` or `false`, for `use_calendar_time_periods`, indicating whether the aggregation period should align to calendar time periods. If this value is omitted from the configuration file, aggregations cannot be used.

For more information on historian aggregations, see: [Aggregate Historian Agent Specification](#).

Supported Influxdb aggregation functions:

Aggregations: `COUNT()`, `DISTINCT()`, `INTEGRAL()`, `MEAN()`, `MEDIAN()`, `MODE()`, `SPREAD()`, `STDDEV()`, `SUM()`

Selectors: `FIRST()`, `LAST()`, `MAX()`, `MIN()`

Transformations: `CEILING()`, `CUMULATIVE_SUM()`, `DERIVATIVE()`, `DIFFERENCE()`, `ELAPSED()`, `NON_NEGATIVE_DERIVATIVE()`, `NON_NEGATIVE_DIFFERENCE()`

More information how to use those functions: https://docs.influxdata.com/influxdb/v1.4/query_language/functions/

Note: Historian aggregations in InfluxDB are different from aggregations employed by other historian agents in VOLTRON. InfluxDB doesn't have a separate agent for aggregations. Instead, aggregation is supported through the `query_historian` function. Other agents can execute an aggregation query directly in InfluxDB by calling the `RPC.export` method `query`. For an example, see [Aggregate Historian Agent Specification](#)

Database Schema

Each InfluxDB database has a `meta` table as well as other tables for different measurements, e.g. one table for "power_kw", one table for "energy", one table for "voltage", etc. (An InfluxDB measurement is similar to a relational table, so for easier understanding, InfluxDB measurements will be referred to below as tables.)

Measurement Table

Example: If a topic name is "*CampusA/Building1/Device1/Power_KW*", the `power_kw` table might look as follows:

time	building	campus	device	source	value
2017-12-28T20:41:00.004260096Z	building1	campusa	device1	scrape	123.4
2017-12-30T01:05:00.004435616Z	building1	campusa	device1	scrape	567.8
2018-01-15T18:08:00.126345Z	building1	campusa	device1	scrape	10

`building`, `campus`, `device`, and `source` are InfluxDB *tags*. `value` is an InfluxDB *field*.

Note: The topic is converted to all lowercase before being stored in the table. In other words, a set of *tag* names, as well as a table name, are created by splitting `topic_id` into substrings (see [meta table](#) below).

So in this example, where the typical format of a topic name is `<campus>/<building>/<device>/<measurement>`, `campus`, `building` and `device` are each stored as tags in the database.

A topic name might not confirm to that convention:

1. The topic name might contain additional substrings, e.g. *CampusA/Building1/LAB/Device/OutsideAirTemperature*. In this case, `campus` will be *campusa/building*, `building` will be *lab*, and `device` will be *device*.
2. The topic name might contain fewer substrings, e.g. *LAB/Device/OutsideAirTemperature*. In this case, the `campus` tag will be empty, `building` will be *lab*, and `device` will be *device*.

Meta Table

The meta table will be structured as in the following example:

time	last_updated	meta_dict	topic	topic_id
1970-01-01T00:00:00Z	2017-12-28T20:47:00.003051+00:00	{u'units': u'kw', u'tz': u'Pacific', u'type': u'float' }	CampusA/Building1/Device1/Power_KW	cam-Power_KW/Building1/device1/power_kw
1970-01-01T00:00:00Z	2017-12-28T20:47:00.003051+00:00	{u'units': u'kwh', u'tz': u'Pacific', u'type': u'float' }	CampusA/Building1/Device1/Energy_KWH	cam-Energy_KWH/Building1/device1/energy_kwh

In the InfluxDB, last_updated, meta_dict and topic are *fields* and topic_id is a *tag*.

Since InfluxDB is a time series database, the time column is required, and a dummy value (time=0, which is 1970-01-01T00:00:00Z based on epoch unix time) is assigned to all topics for easier metadata updating. Hence, if the contents of meta_dict change for a specific topic, both last_updated and meta_dict values for that topic will be replaced in the table.

MQTT Historian

Overview

The MQTT Historian agent publishes data to an MQTT broker.

The mqttlistener.py script will connect to the broker and print all messages.

Dependencies

The Paho MQTT library from Eclipse is needed for the agent and can be installed with:

```
pip install paho-mqtt
```

The Mosquitto MQTT broker may be useful for testing and can be installed with

```
apt-get install mosquitto
```

Mongo Historian

Prerequisites

1. MongoDB

Setup mongodb based on using one of the three below scripts.

1. Install as root on Redhat or Cent OS

```
sudo scripts/historian-scripts/root_install_mongo_rhel.sh
```

The above script will prompt user for os version, db user name, password and database name Once installed you can start and stop the service using the command:

```
sudo service mongod [start|stop|service]
```

2. Install as root on Ubuntu

```
sudo scripts/historian-scripts/root_install_mongo_ubuntu.sh
```

The above script will prompt user for os version, db user name, password and database name Once installed you can start and stop the service using the command:

sudo service mongod [start|stop|service]

3. Install as non root user on any Linux machine

```
scripts/historian-scripts/install_mongodb.sh
```

Usage: install_mongodb.sh [-h] [-d download_url] [-i install_dir] [-c config_file] [-s]

Optional arguments: -s setup admin user and test collection after install and startup

-d download url. defaults to https://fastdl.mongodb.org/linux/mongodb-linux-x86_64-3.2.4.tgz

-i install_dir. defaults to current_dir/mongo_install

-c config file to be used for mongodb startup. Defaults to default_mongodb.conf in the same directory as this script. Any datapath mentioned in the config file should already exist and should have write access to the current user

-h print this help message

2. Mongodb connector

This historian requires a mongodb connector installed in your activated volttron environment to talk to mongodb. Please execute the following from an activated shell in order to install it.

```
pip install pymongo
```

3. Configuration Options

The historian configuration file can specify

```
"history_limit_days": <n days>
```

which will remove entries from the data and rollup collections older than n days. Timestamps passed to the manage_db_size method are truncated to the day.

Platform Historian

A platform historian is a “friendly named” historian on a VOLTRON instance. It always has the identity (see vip) of platform.historian. A platform historian is made available to a volttron central agent for monitoring of the VOLTRON instances health and plotting topics from the platform historian. In order for one of the (historians)[Historians] to be turned into a platform historian the identity keyword must be added to it’s configuration with the value of platform.historian. The following configuration file shows a sqlite based platform historian configuration.

```
{
  "agentid": "sqlhistorian-sqlite",
  "identity": "platform.historian",
  "connection": {
    "type": "sqlite",
```

(continues on next page)

(continued from previous page)

```
    "params": {
      "database": "~/voltron/data/platform.historian.sqlite"
    }
  }
}
```

SQL Historian

An SQL Historian is available as a core service. The sql historian has been programmed to allow for inconsistent network connectivity (automatic re-connection to tcp based databases). All additions to the historian are batched and wrapped within a transaction with commit and rollback functions properly implemented. This allows the maximum throughput of data with the most protection. The following example configurations show the different options available for configuring the SQL Historian Agent.

MySQL Specifics

MySQL requires a third party driver (mysql-connector) to be installed in order for it to work. Please execute the following from an activated shell in order to install it.

```
pip install --allow-external mysql-connector-python mysql-connector-python
```

In addition, the mysql database must be created and permissions

granted for select, insert and update before the agent is started. In order to support timestamp with microseconds you need at least MySQL 5.6.4. Please see this [MySQL documentation](#) for more details | The following is a minimal configuration file for using a MySQL based historian. Other options are available and are documented <http://dev.mysql.com/doc/connector-python/en/connector-python-connectargs.html>. **Not all parameters have been tested, use at your own risk.**

```
{
  "agentid": "sqlhistorian-mysql",
  "connection": {
    "type": "mysql",
    "params": {
      "host": "localhost",
      "port": 3306,
      "database": "voltron",
      "user": "user",
      "passwd": "pass"
    }
  }
}
```

Sqlite3 Specifics

An Sqlite Historian provides a convenient solution for under powered systems. The database is parameter is a location on the file system. By default it is relative to the agents installation directory, however it will respect a rooted or relative path to the database.

```
{
  "agentid": "sqlhistorian-sqlite",
  "connection": {
    "type": "sqlite",
    "params": {
      "database": "data/historian.sqlite"
    }
  }
}
```

sMAP Historian

This historian allows VOLTRON data to be published to an sMAP server. This replaces the DataLogger functionality of V2.0 as well as the capability of 2.0 drivers to publish directly to sMAP.

To configure this historian the following must be in the config file. This file is setup to point at the [available Test Instance of sMAP](#). For reliable storage, please setup your own instance.

```
{
  "agentid": "smap_historian",
  "source": "MyTestSource",
  "archiver_url": "http://smap-test.cloudapp.net",
  "key": "LEq1cEGc04RtcKX6riiX7eaML8Z82xEgQrp7"
}
```

That's it! With this configuration, data will be pulled off the message bus and published to your sMAP server.

3.5.4 VOLTRON Driver Framework

All Voltton drivers are implemented through the Master Driver Agent and are technically sub-agents running in the same process as the Master Driver Agent. Each of these driver sub-agents is responsible for creating an interface to a single device. Creating that interface is facilitated by an instance of an interface class. There are a variety of interface classes included. The most commonly used interfaces are BACnet and Modbus.

Automatically Generating BACnet Configuration Files

Included with the platform are two scripts for finding and configuring BACnet devices. These scripts are located in `scripts/bacnet`. `bacnet_scan.py` will scan the network for devices. `grab_bacnet_config.py` creates a CSV file for the BACnet driver that can be used as a starting point for creating your own register configuration.

Both scripts are configured with the file `BACpypes.ini`.

Configuring the Utilities

While running both scripts create a temporary virtual BACnet device using the `bacpypes` library. The virtual device must be configured properly in order to work. This configuration is stored in `scripts/bacnet/BACpypes.ini` and will be read automatically when the utility is run.

The only value that (usually) needs to be changed is the **address** field. ****This is the address bound to the port on the machine you are running the script from, NOT A TARGET DEVICE! **** This value should be set to the IP address of the network interface used to communicate with the remote device. If there is more than one network interface you must use the address of the interface connected to the network that can reach the device.

In Linux you can usually get the addresses bound to all interfaces by running `ifconfig` from the command line.

If a different outgoing port other than the default 47808 must be used, it can be specified as part of the address in the form

```
<ADDRESS>:<PORT>
```

In some cases, the netmask of the network will be needed for proper configuration. This can be done following this format

```
<ADDRESS>/<NETMASK>:<PORT>
```

where `<NETMASK>` is the netmask length. The most common value is 24. See <http://www.computerhope.com/jargon/n/netmask.htm>

In some cases, you may also need to specify a different device ID by changing the value of `objectIdentifier` so the virtual BACnet device does not conflict with any devices on the network. `objectIdentifier` defaults to 599.

Sample BACpypes.ini

```
[BACpypes]
objectName: Betelgeuse
address: 10.0.2.15/24
objectIdentifier: 599
maxApduLengthAccepted: 1024
segmentationSupported: segmentedBoth
vendorIdentifier: 15
```

Scanning for BACnet Devices

If the addresses for BACnet devices are unknown they can be discovered using the `bacnet_scan.py` utility.

To run the utility simply execute the following command:

```
python bacnet_scan.py
```

and expect output similar to this:

```
Device Address      = <Address 192.168.1.42>
Device Id           = 699
maxAPDULengthAccepted = 1024
segmentationSupported = segmentedBoth
vendorID            = 15

Device Address      = <RemoteStation 1002:11>
Device Id           = 540011
maxAPDULengthAccepted = 480
segmentationSupported = segmentedBoth
vendorID            = 5
```

Reading Output

The address where the device can be reached is listed on the **Device Address** line. The BACnet device ID is listed on the **Device Id** line. The remaining lines are informational and not needed to configure the BACnet driver.

For the first example, the IP address 192.168.1.42 can be used to reach the device. The second device is behind a BACnet router and can be reached at 1002:11. See RouterAddressing Remote Station addressing.

Options

- `--address ADDRESS` Send the WhoIs request only to a specific address. Useful as a way to ping devices on a network that blocks broadcast traffic.
- `--range LOW HIGH` Specify the device ID range for the results. Useful for filtering.
- `--timeout SECONDS` Specify how long to wait for responses to the original broadcast. This defaults to 5 which should be sufficient for most networks.
- `--csv-out CSV_OUT` Write the discovered devices to a CSV file. This can be used as input for `grab_multiple_configs.py`. See *Scraping Multiple Devices*.

Automatically Generating a BACnet Registry Configuration File

A CSV registry configuration file for the BACnet driver can be generated with the `grab_bacnet_config.py` script. **This configuration will need to be edited before it can be used.**

The utility is invoked with the command:

```
python grab_bacnet_config.py <device id>
```

This will query the device with the matching device ID for configuration information and print the resulting CSV file to the console.

Note: Previous to VOLTRON 3.5 `grab_bacnet_config.py` took the device address as an argument instead of the device ID.

In order to save the configuration to a file use the `--out-file` option to specify the output file name.

Optionally the `--address` option can be used to specify the address of the target. In some cases, this is needed to help establish a route to the device.

Output and Assumptions

Attempts at determining if a point is writable proved too unreliable. Therefore all points are considered to be read-only in the output.

The only property for which a point is setup for an object is **presentValue**.

By default, the **Volttron Point Name** is set to the value of the **name** property of the BACnet object on the device. In most cases this name is vague. No attempt is made at choosing a better name. A duplicate of “Volttron Point Name” column called “Reference Point Name” is created so that once “Volttron Point Name” is changed a reference remains to the actual BACnet device object name.

Meta data from the objects on the device is used to attempt to put useful info in the **Units Unit Details**, and **Notes** columns. Information such as the range of valid values, defaults, the resolution or sensor input, and enumeration or state names are scraped from the device.

With a few exceptions “Units” is pulled from the object’s “units” property and given the name used by the `bacpypes` library to describe it. If a value in the **Units** column takes the form

```
UNKNOWN UNIT ENUM VALUE: <value>
```

then the device is using a nonstandard value for the units on that object.

Scraping Multiple Devices

The `grab_multiple_configs.py` script will use the CSV output of `bacnet_scan.py` to automatically run `grab_bacnet_config.py` on every device listed in the CSV file.

The output is put in two directories. `devices/` contains basic driver configurations for the scrapped devices. `registry_configs/` contains the registry file generated by `grab_bacnet_config.py`.

`grab_multiple_configs.py` makes no assumptions about device names or topics, however the output is appropriate for the `install_master_driver_configs.py` script.

Options

- `--out-directory OUT_DIRECTORY` Specify the output directory.
- `--use-proxy` Use `proxy_grab_bacnet_config.py` to gather configuration data.

BACnet Proxy Alternative Scripts

Both `grab_bacnet_config.py` and `bacnet_scan.py` have alternative versions called `proxy_grab_bacnet_config.py` and `proxy_bacnet_scan.py` respectively. These versions require that the VOLTRON platform is running and BACnet Proxy agent is running. Both of these agents use the same command line arguments as their independent counterparts.

Warning: These versions of the BACnet scripts are intended as a proof of concept and have not been optimized for performance. `proxy_grab_bacnet_config.py` takes about 10 times longer to grab a configuration than `grab_bacnet_config.py`

Problems and Debugging

Both `grab_bacnet_config.py` and `bacnet_scan.py` creates a virtual device that open up a port for communication with devices. If BACnet Proxy is running on the VOLTRON platform it will cause both of these scripts to fail at startup. Stopping the BACnet Proxy will resolve the problem.

Typically the utility should run quickly and finish in 30 seconds or less. In our testing, we have never seen a successful scrape take more than 15 seconds on a very slow device with many points. Many devices will scrape in less that 3 seconds.

If the utility has not finished after about 60 seconds it is probably having trouble communicating with the device and should be stopped. Rerunning with debug output can help diagnose the problem.

To output debug messages to the console add the `--debug` switch to the **end** of the command line arguments.

```
python grab_bacnet_config.py <device ID> --out-file test.csv --debug
```

On a successful run you will see output similar to this:

```
DEBUG:<u>main</u>:initialization
DEBUG:<u>main</u>:  - args: Namespace(address='10.0.2.20', buggers=False, debug=[],
↪ini=<class 'bacpyes.consolelogging.ini'>, max_range_report=1e+20, out_file=<open_
↪file 'out.csv', mode 'wb' at 0x901b0d0>)
DEBUG:<u>main</u>.SynchronousApplication:<u>init</u> (<bacpyes.app.LocalDeviceObject_
↪object at 0x901de6c>, '10.0.2.15')
```

(continues on next page)

(continued from previous page)

```

DEBUG:<u>main</u>:starting build
DEBUG:<u>main</u>:pduSource = <Address 10.0.2.20>
DEBUG:<u>main</u>:iAmDeviceIdentifier = ('device', 500)
DEBUG:<u>main</u>:maxAPDULengthAccepted = 1024
DEBUG:<u>main</u>:segmentationSupported = segmentedBoth
DEBUG:<u>main</u>:vendorID = 5
DEBUG:<u>main</u>:device_name = MS-NCE2560-0
DEBUG:<u>main</u>:description =
DEBUG:<u>main</u>:objectCount = 32
DEBUG:<u>main</u>:object name = Building/FCB.Local Application.Room Real Temp 2
DEBUG:<u>main</u>: object type = analogInput
DEBUG:<u>main</u>: object index = 3000274
DEBUG:<u>main</u>: object units = degreesFahrenheit
DEBUG:<u>main</u>: object units details = -50.00 to 250.00
DEBUG:<u>main</u>: object notes = Resolution: 0.1
DEBUG:<u>main</u>:object name = Building/FCB.Local Application.Room Real Temp 1
DEBUG:<u>main</u>: object type = analogInput
DEBUG:<u>main</u>: object index = 3000275
DEBUG:<u>main</u>: object units = degreesFahrenheit
DEBUG:<u>main</u>: object units details = -50.00 to 250.00
DEBUG:<u>main</u>: object notes = Resolution: 0.1
DEBUG:<u>main</u>:object name = Building/FCB.Local Application.OSA
DEBUG:<u>main</u>: object type = analogInput
DEBUG:<u>main</u>: object index = 3000276
DEBUG:<u>main</u>: object units = degreesFahrenheit
DEBUG:<u>main</u>: object units details = -50.00 to 250.00
DEBUG:<u>main</u>: object notes = Resolution: 0.1
...

```

and will finish something like this:

```

...
DEBUG:<u>main</u>:object name = Building/FCB.Local Application.MOTOR1-C
DEBUG:<u>main</u>: object type = binaryOutput
DEBUG:<u>main</u>: object index = 3000263
DEBUG:<u>main</u>: object units = Enum
DEBUG:<u>main</u>: object units details = 0-1 (default 0)
DEBUG:<u>main</u>: object notes = BinaryPV: 0=inactive, 1=active
DEBUG:<u>main</u>:finally

```

Typically if the BACnet device is unreachable for any reason (wrong IP, network down/unreachable, wrong interface specified, device failure, etc) the scraper will stall at this message:

```

DEBUG:<u>main</u>:starting build

```

If you have not specified a valid interface in BACpypes.ini you will see the following error with a stack trace:

```

ERROR:<u>main</u>:an error has occurred: [Errno 99] Cannot assign requested address
<Python stack trace cut>

```

BACnet Proxy Agent

Introduction

Communication with BACnet device on a network happens via a single virtual BACnet device. Previous versions of Volttron used one virtual device per device on the network. This only worked in a limited number of circumstances. (This problem is fixed in the legacy sMap drivers in Volttron 3.0 only) In the new driver architecture, we have a separate agent specifically for communicating with BACnet devices and managing the virtual BACnet device.

Configuration

The agent configuration sets up the virtual BACnet device.

```
{
  "device_address": "10.0.2.15",
  "max_apdu_length": 1024,
  "object_id": 599,
  "object_name": "Volttron BACnet driver",
  "vendor_id": 15,
  "segmentation_supported": "segmentedBoth"
}
```

BACnet device settings

- **device_address** - Address bound to the network port over which BACnet communication will happen on the computer running VOLTTRON. This is **NOT** the address of any target device. See *Device Addressing*.
- **object_id** - ID of the Device object of the virtual BACnet device. Defaults to 599. Only needs to be changed if there is a conflicting BACnet device ID on your network.

These settings determine the capabilities of the virtual BACnet device. BACnet communication happens at the lowest common denominator between two devices. For instance, if the BACnet proxy supports segmentation and the target device does not communication will happen without segmentation support and will be subject to those limitations. Consequently, there is little reason to change the default settings outside of the **max_apdu_length** (the default is not the largest possible value).

- **max_apdu_length** - (From bacpypes documentation) BACnet works on lots of different types of networks, from high-speed Ethernet to “slower” and “cheaper” ARCNET or MS/TP (a serial bus protocol used for a field bus defined by BACnet). For devices to exchange messages they have to know the maximum size message the device can handle. (End BACpypes docs)

This setting determines the largest APDU accepted by the BACnet virtual device. Valid options are 50, 128, 206, 480, 1024, and 1476. Defaults to 1024.(Optional)

- **object_name** - Name of the object. Defaults to “Volttron BACnet driver”. (Optional)
- **vendor_id** - Vendor ID of the virtual BACnet device. Defaults to 15. (Optional)
- **segmentation_supported** - (From bacpypes documentation) A vast majority of BACnet communications traffic fits into one message, but there can be times when larger messages are convenient and more efficient. Segmentation allows larger messages to be broken up into segments and spliced back together. It is not unusual for “low power” field equipment to not support segmentation. (End BACpypes docs)

Possible setting are “segmentedBoth” (default), “segmentedTransmit”, “segmentedReceive”, or “noSegmentation” (Optional)

Device Addressing

In some cases, it will be needed to specify the subnet mask of the virtual device or a different port number to listen on. The full format of the BACnet device address is

```
<ADDRESS>/<NETMASK>:<PORT>
```

where <PORT> is the port to use and <NETMASK> is the netmask length. The most common value is 24. See <http://www.computerhope.com/jargon/n/netmask.htm>

For instance, if you need to specify a subnet mask of 255.255.255.0 and the IP address bound to the network port is 192.168.1.2 you would use the address

```
192.168.1.2/24
```

If your BACnet network is on a different port (47809) besides the default (47808) you would use the address

```
192.168.1.2:47809
```

If you need to do both

```
192.168.1.2/24:47809
```

Communicating With Multiple BACnet Networks

If two BACnet devices are connected to different ports they are considered to be on different BACnet networks. In order to communicate with both devices, you will need to run one BACnet Proxy Agent per network.

Each proxy will need to be bound to different ports appropriate for each BACnet network and will need a different VIP identity specified. When configuring drivers you will need to specify which proxy to use by *specifying the VIP identity*.

TODO: Add link to docs showing how to specify the VIP IDENTITY when installing an agent.

For example, a proxy connected to the default BACnet network

```
{
  "device_address": "192.168.1.2/24"
}
```

and another on port 47809

```
{
  "device_address": "192.168.1.2/24:47809"
}
```

a device on the first network

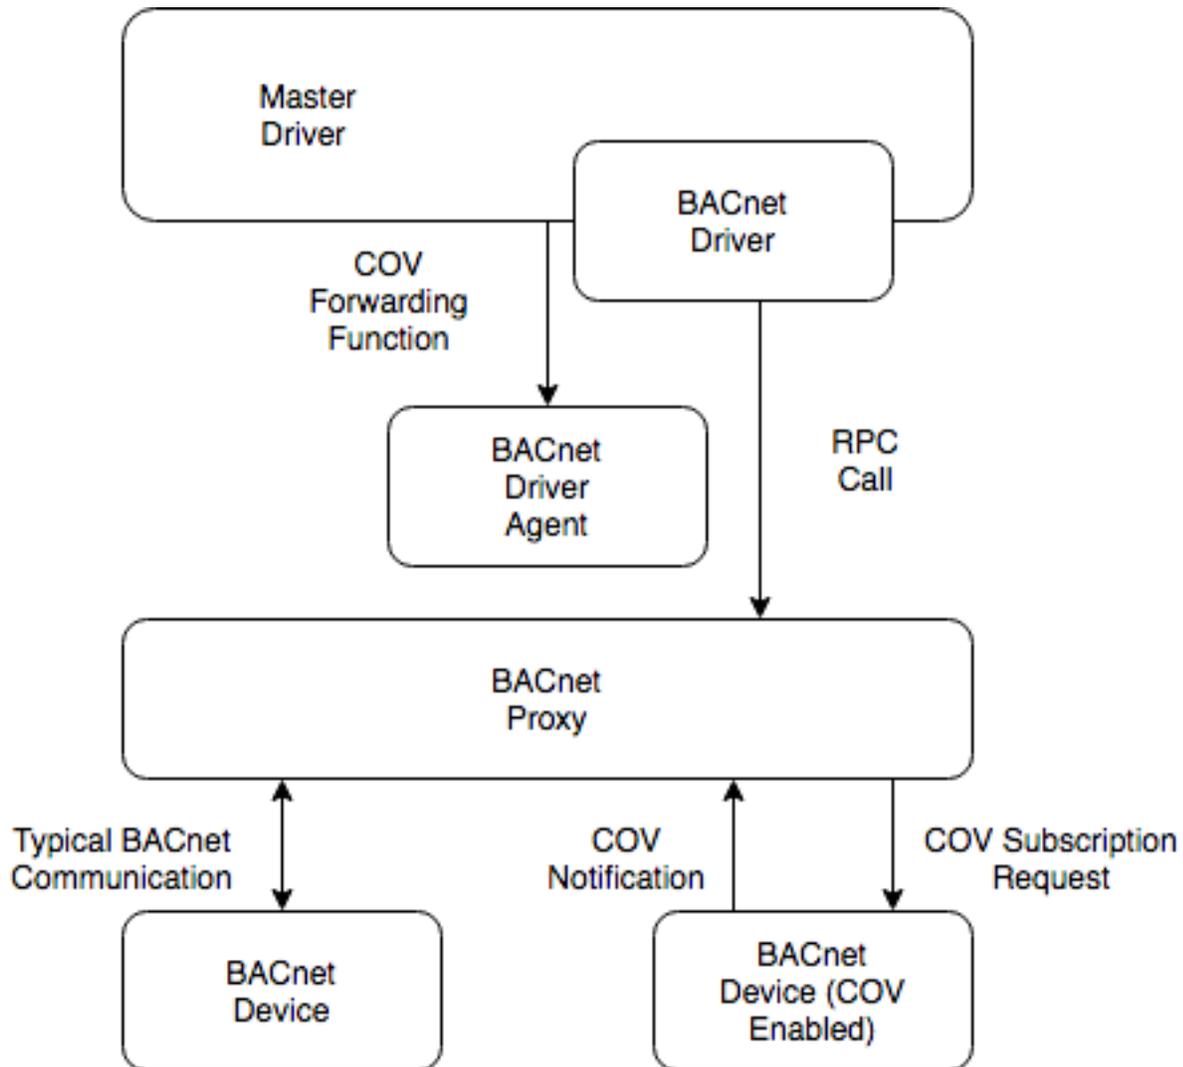
```
{
  "driver_config": {"device_address": "1002:12",
                    "proxy_address": "platform.bacnet_proxy_47808",
                    "timeout": 10},
  "driver_type": "bacnet",
  "registry_config": "config://registry_configs/bacnet.csv",
  "interval": 60,
  "timezone": "UTC",
  "heart_beat_point": "Heartbeat"
}
```

and a device on the second network

```
{
  "driver_config": {"device_address": "12000:5",
    "proxy_address": "platform.bacnet_proxy_47809",
    "timeout": 10},
  "driver_type": "bacnet",
  "registry_config": "config://registry_configs/bacnet.csv",
  "interval": 60,
  "timezone": "UTC",
  "heart_beat_point": "Heartbeat"
}
```

Notice that both configs use the same registry configuration (config://registry_configs/bacnet.csv). This is perfectly fine as long as the registry configuration is appropriate for both devices. For scraping large numbers of points from a single BACnet device, there is an optional timeout parameter provided, to prevent the master driver timing out while the BACnet Proxy Agent is collecting points.

BACnet Change of Value Services



Change of Value Services added in version 0.5 of the BACnet Proxy and version 3.2 of the Master Driver.

There are a variety of scenarios in which a user may desire data from some BACnet device point values to be published independently of the regular scrape interval. Bacpyes provides a “ChangeOfValueServices” (hereby referred to as ‘COV’) module, which enables a device to push updates to the platform.

The BACnet COV requires that points on the device be properly configured for COV. A point on the BACnet device can be configured with the ‘covIncrement’ property, which determines the threshold for a COV notification (note: this property must be configured by the device operator - VOLTRON does not provide the ability to set or modify this property).

Based on configuration options for BACnet drivers, the driver will instruct the BACnet Proxy to establish a COV subscription with the device. The subscription will last for an amount of time specified in the driver configuration, and will auto-renew the subscription. If the proxy loses communication with the device or the device driver is stopped the subscription will be removed when the lifetime expires. While the subscription exists, the device will send (confirmed) notifications to which will be published, with the topic based on the driver’s configured publish topics.

<https://bacpyes.readthedocs.io/en/latest/modules/service/cov.html>

BACnet Router Addressing

The underlying library that Volttron uses for BACnet supports IP to MS/TP routers. Devices behind the router use a Remote Station address in the form

```
<network>:<address>
```

where **<network>** is the configured network ID of the router and **<address>** is the address of the device behind the router.

For example to access the device at **<address>** 12 for a router configured for **<network>** 1002 can be accessed with this address:

```
1002:12
```

<network> must be number from 0 to 65534 and **<address>** must be a number from 0 to 255.

This type of address can be used anywhere an address is required in configuration of the Volttron BACnet driver.

Caveats

VOLTRON uses a UDP broadcast mechanism to establish the route to the device. If the route cannot be established it will fall back to a UDP broadcast for all communication with the device. If the IP network where the router is connected blocks UDP broadcast traffic then these addresses will not work.

Obix History Agent

The Obix History Agent captures data history data from an Obix RESTful interface and publishes it to the message bus like a driver for capture by agents and historians. The Agent will setup its queries to ensure that data is only publishes once. For points queried for the first time it will go back in time and publish old data as configured.

The data will be colated into device all publishes automatically and will use a timestamp in the header based on the timestamps reported by the Obix interface. The publishes will be made in chronological order.

Units data is automatically read from the device.

For sending commands to devices see Obix-config.

Agent Configuration

There are three arguments for the **driver_config** section of the device configuration file:

- **url** - URL of the interface.
- **username** - User name for site..
- **password** - Password for username.
- **check_interval** - How often to check for new data on each point.
- **path_prefix** - Path prefix for all publishes.
- **register_config** - Registry configuration file.
- **default_last_read** - Time, in hours, to go back and retrieve data for a point for the first time.

Here is an example device configuration file:

```
{
  "url": "http://example.com/obix/histories/EXAMPLE/",
  "username": "username",
  "password": "password",
  # Interval to query interface for updates in minutes.
  # History points are only published if new data is available
  # config points are gathered and published at this interval.
  "check_interval": 15,
  # Path prefix for all publishes
  "path_prefix": "devices/obix/history/",
  "register_config": "config://registry_config.csv",
  "default_last_read": 12
}
```

A sample Obix configuration file can be found in the VOLTTRON repository in `services/core/ObixHistoryPublish/config`

Registry Configuration File

Similar to a driver the Obix History Agent requires a registry file to select the points to publish.

The registry configuration file is a **CSV** file. Each row configures a point on the device.

The following columns are required for each row:

- **Device Name** - Name of the device to associate with this point.
- **Volttron Point Name** - The Volttron Point name to use when publishing this value.
- **Obix Name** - Name of the point on the obix interface. Escaping of spaces and dashes for use with the interface is handled internally.

Any additional columns will be ignored. It is common practice to include a **Notes** or **Unit Details** for additional information about a point.

The following is an example of a Obix History Agent registry configuration file:

Table 13: Obix

Device Name	Volttron Point Name	Obix Name
device1	Local Outside Dry Bulb	Local Outside Dry Bulb
device2	CG-1 Gas Flow F-2	CG-1 Gas Flow F-2
device2	Cog Plant Gas Flow F-1	Cog Plant Gas Flow F-1
device2	Boiler Plant Hourly Gas Usage	Boiler Plant Hourly Gas Usage
device3	CG-1 Water Flow H-1	CG-1 Water Flow H-1

A sample Obix History Agent configuration can be found in the VOLTRON repository in `services/core/ObixHistoryPublish/registry_config.csv`

Automatic Obix Configuration File Creation

A script that will automatically create both a device and register configuration file for a site is located in the repository at `scripts/obix/get_obix_history_config.py`.

The utility is invoked with the command:

```
python get_obix_history_config.py <url> <registry_file> <driver_file>
-u <username> -p <password> -d <device name>
```

If either the `registry_file` or `driver_file` is omitted the script will output those files to stdout.

If either the `username` or `password` options are left out the script will ask for them on the command line before proceeding.

The `device name` option specifies a default device for every point in the configuration.

The registry file produced by this script assumes that the *Volttron Point Name* and the *Obix Name* have the same value. Also, it is assumed that all points should be read only. Users are expected to fix this as appropriate.

Using Third Party Drivers

In some cases you will need to use a driver provided by a third-party to interact with a device. While the interface file can be copied into `services/core/MasterDriverAgent/master_driver/interfaces` this does not work well with third-party code that is under source control.

The recommended method is to create a symbolic link to the interface file in `services/core/MasterDriverAgent/master_driver/interfaces`. This will work in both a development environment and in production. When packing the agent for installation a copy of the linked file will be put in the resulting wheel file.

```
#A copy of the interface file lives in ~/my_driver/my_driver.py
#Create the link
ln -s ~/my_driver/my_driver.py services/core/MasterDriverAgent/master_driver/
↪interfaces/my_driver.py

#remove the link
rm services/core/MasterDriverAgent/master_driver/interfaces/my_driver.py
```

VOLTRON Drivers

Driver Configuration

The Master Driver Agent manages all device communication. To communicate with devices you must setup and deploy the Master Driver Agent.

Configuration for each device consists of 3 parts:

- Master Driver Agent configuration file - lists all driver configuration files to load
- Driver configuration file - contains the general driver configuration and device settings
- Device Register configuration file - contains the settings for each individual data point on the device

For each device, you must create a driver configuration file, device register configuration file, and an entry in the Master Driver Agent configuration file.

Once configured, the Master Driver Agent is configured and deployed in a manner similar to any other agent.

The Master Driver Agent along with Historian Agents replace the functionality of sMap from VOLTTRON 2.0 and thus sMap is no longer a requirement for VOLTTRON.

Master Driver Agent Configuration

The Master Driver Agent configuration consists of general settings for all devices. The default values of the master driver should be sufficient for most users. The user may optionally change the interval between device scrapes with the `driver_scrape_interval`.

The following example sets the `driver_scrape_interval` to 0.05 seconds or 20 devices per second:

```
{
  "driver_scrape_interval": 0.05,
  "publish_breadth_first_all": false,
  "publish_depth_first": false,
  "publish_breadth_first": false,
  "publish_depth_first_all": true,
  "group_offset_interval": 0.0
}
```

- **driver_scrape_interval** - Sets the interval between devices scrapes. Defaults to 0.02 or 50 devices per second. Useful for when the platform scrapes too many devices at once resulting in failed scrapes.
- **group_offset_interval** - Sets the interval between when groups of devices are scraped. Has no effect if all devices are in the same group.

In order to improve the scalability of the platform unneeded device state publishes for all devices can be turned off. All of the following setting are optional and default to *True*.

- **publish_depth_first_all** - Enable “depth first” publish of all points to a single topic for all devices.
- **publish_breadth_first_all** - Enable “breadth first” publish of all points to a single topic for all devices.
- **publish_depth_first** - Enable “depth first” device state publishes for each register on the device for all devices.
- **publish_breadth_first** - Enable “breadth first” device state publishes for each register on the device for all devices.

An example master driver configuration file can be found in the VOLTTRON repository in `examples/configurations/drivers/master-driver.agent`.

Driver Configuration File

Note: The terms *register* and *point* are used interchangeably in the documentation and in the configuration setting names. They have the same meaning.

Each device configuration has the following form:

```
{
  "driver_config": {"device_address": "10.1.1.5",
                   "device_id": 500},
  "driver_type": "bacnet",
  "registry_config": "config://registry_configs/vav.csv",
  "interval": 60,
  "heart_beat_point": "heartbeat",
  "group": 0
}
```

The following settings are required for all device configurations:

- **driver_config** - Driver specific setting go here. See below for driver specific settings.
- **driver_type** - Type of driver to use for this device: bacnet, modbus, fake, etc.
- **registry_config** - Reference to a configuration file in the configuration store for registers on the device. See the [Registry Configuration File](#) and [Adding Device Configurations to the Configuration Store](#) sections below.

These settings are optional:

- **interval** - Period which to scrape the device and publish the results in seconds. Defaults to 60 seconds.
- **heart_beat_point** - A Point which to toggle to indicate a heartbeat to the device. A point with this Volttron Point Name must exist in the registry. If this setting is missing the driver will not send a heart beat signal to the device. Heart beats are triggered by the Actuator Agent which must be running to use this feature.
- **group** - Group this device belongs to. Defaults to 0

These settings are used to create the topic that this device will be referenced by following the VOLTRON convention of {campus}/{building}/{unit}. This will also be the topic published on, when the device is periodically scraped for it's current state.

The topic used to reference the device is derived from the name of the device configuration in the store. See the [Adding Device Configurations to the Configuration Store](#) section.

Device Grouping

Devices may be placed into groups to separate them logically when they are scraped. This is done by setting the *group* in the device configuration. *group* is a number greater than or equal to 0. Only number of devices in the same group and the *group_offset_interval* are considered when determining when to scrape a device.

This is useful in two cases. First, if you need to ensure that certain devices are scraped in close proximity to each other you can put them in their own group. If this causes devices to be scraped too quickly the groups can be separated out time wise using the *group_offset_interval* setting. Second, you may scrape devices on different networks in parallel for performance. For instance BACnet devices behind a single MSTP router need to be scraped slowly and serially, but devices behind different routers may be scraped in parallel. Grouping devices by router will do this automatically.

The *group_offset_interval* is applied by multiplying it by the *group* number. If you intent to use *group_offset_interval* only use consecutive *group* values that start with 0.

Registry Configuration File

Registry configuration files setup each individual point on a device. Typically this file will be in CSV format, but the exact format is driver specific. See the section for a particular driver for the registry configuration format.

The following is a simple example of a MODBUS registry configuration file:

Table 14: Catalyst 371

Reference Point Name	Volttron Point Name	Units	Units Details	Modbus Register	Writable	Point Address	Default Value	Notes
CO2Sensor	ReturnAirCO2	PPM	0.00-2000.00	>f	FALSE	1001		CO2 Reading 0.00-2000.0 ppm
CO2Stpt	ReturnAirCO2Stpt	PPM	1000.00 (default)	>f	TRUE	1011	1000	Setpoint to enable demand control ventilation
HeatCall2	HeatCall2	On / Off	on/off	BOOL	FALSE	1114		Status indicator of heating stage 2 need

Adding Device Configurations to the Configuration Store

Configurations are added to the Configuration Store using the command line `volttron-ctl config store platform.driver <name> <file name> <file type>`.

- **name** - The name used to refer to the file from the store.
- **file name** - A file containing the contents of the configuration.
- **file type** - `--raw`, `--json`, or `--csv`. Indicates the type of the file. Defaults to `--json`.

The main configuration must have the name `config`

Device configuration but **not** registry configurations must have a name prefixed with `devices/`. Scripts that automate the process will prefix registry configurations with `registry_configs/`, but that is not a requirement for registry files.

The name of the device's configuration in the store is used to create the topic used to reference the device. For instance, a configuration named `devices/PNNL/ISB1/vav1` will publish scrape results to `devices/PNNL/ISB1/vav1` and is accessible with the Actuator Agent via `PNNL/ISB1/vav1`.

The name of a registry configuration must match the name used to refer to it in the driver configuration. The reference is not case sensitive.

If the Master Driver Agent is running any changes to the configuration store will immediately affect the running devices according to the changes.

Consider the following three configuration files:

A master driver configuration called `master-driver.agent`:

```
{
  "driver_scrape_interval": 0.05
}
```

A MODBUS device configuration file called `modbus1.config`:

```
{
  "driver_config": {"device_address": "10.1.1.2",
```

(continues on next page)

(continued from previous page)

```

        "port": 502,
        "slave_id": 5},
    "driver_type": "modbus",
    "registry_config": "config://registry_configs/hvac.csv",
    "interval": 60,
    "timezone": "UTC",
    "heart_beat_point": "heartbeat"
}

```

A MODBUS registry configuration file called *catalyst371.csv*:

Table 15: catalyst371.csv

Reference Point Name	Volttron Point Name	Units	Units Details	Modbus Register	Writable	Point Address	Default Value	Notes
CO2Sensor	ReturnAirCO2	PPM	0.00-2000.00	>f	FALSE	1001		CO2 Reading 0.00-2000.0 ppm
CO2Stpt	ReturnAirCO2Stpt	PPM	1000.00 (default)	>f	TRUE	1011	1000	Setpoint to enable demand control ventilation
HeatCall2	HeatCall2	On / Off	on/off	BOOL	FALSE	1114		Status indicator of heating stage 2 need

To store the master driver configuration run the command

```
volttron-ctl config store platform.driver config master-driver.agent
```

To store the registry configuration run the command (note the `-csv` option)

```
volttron-ctl config store platform.driver registry_configs/hvac.csv
catalyst371.csv --csv
```

Note the name `registry_configs/hvac.csv` matches the configuration reference in the file `modbus1.config`.

To store the driver configuration run the command

```
volttron-ctl config store platform.driver devices/my_campus/my_building/hvac1
modbus1.config
```

Converting Old Style Configuration

The new Master Driver no longer supports the old style of device configuration. The old `device_list` setting is ignored.

To simplify updating to the new format `scripts/update_master_driver_config.py` is provide to automatically update to the new configuration format.

With the platform running run:

```
python scripts/update_master_driver_config.py <old configuration> <output>
```

old_configuration is the main configuration file in the old format. The script automatically modifies the driver files to create references to CSV files and adds the CSV files with the appropriate name.

output is the target output directory.

If the `--keep-old` switch is used the old configurations in the output directory (if any) will not be deleted before new configurations are created. Matching names will still be overwritten.

The output from `scripts/update_master_driver_config.py` can be automatically added to the configuration store for the Master Driver agent with `scripts/install_master_driver_configs.py`.

Creating and naming configuration files in the form needed by `scripts/install_master_driver_configs.py` can speed up the process of changing and updating a large number of configurations. See the `--help` message for `scripts/install_master_driver_configs.py` for more details.

Device State Publishes

By default, the value of each register on a device is published 4 different ways when the device state is published. Consider the following settings in a driver configuration stored under the name `devices/pnnl/isb1/vav1`:

```
{
  "driver_config": {"device_address": "10.1.1.5",
                   "device_id": 500},
  "driver_type": "bacnet",
  "registry_config": "config://registry_configs/vav.csv",
}
```

In the `vav.csv` file is a register with the name `temperature`. For these examples the current value of the register on the device happens to be 75.2 and the meta data is

```
{"units": "F"}
```

When the driver publishes the device state the following 2 things will be published for this register:

A “depth first” publish to the topic `devices/pnnl/isb1/vav1/temperature` with the following message:

```
[75.2, {"units": "F"}]
```

A “breadth first” publish to the topic `devices/temperature/vav1/isb1/pnnl` with the following message:

```
[75.2, {"units": "F"}]
```

These publishes can be turned off by setting `publish_depth_first` and `publish_breadth_first` to `false` respectively.

Also these two publishes happen once for all registers:

A “depth first” publish to the topic `devices/pnnl/isb1/vav1/all` with the following message:

```
[{"temperature": 75.2, ...}, {"temperature":{"units": "F"}, ...}]
```

A “breadth first” publish to the topic `devices/all/vav1/isb1/pnnl` with the following message:

```
[{"temperature": 75.2, ...}, {"temperature":{"units": "F"}, ...}]
```

These publishes can be turned off by setting `publish_depth_first_all` and `publish_breadth_first_all` to `false` respectively.

Device Scalability Settings

In order to improve the scalability of the platform unneeded device state publishes for a device can be turned off. All of the following settings are optional and will override the value set in the main master driver configuration.

- **publish_depth_first_all** - Enable “depth first” publish of all points to a single topic.
- **publish_breadth_first_all** - Enable “breadth first” publish of all points to a single topic.
- **publish_depth_first** - Enable “depth first” device state publishes for each register on the device.
- **publish_breadth_first** - Enable “breadth first” device state publishes for each register on the device.

It is common practice to set **publish_breadth_first_all**, **publish_depth_first**, and **publish_breadth_first** to *False* unless they are specifically needed by an agent running on the platform.

Note: All Historian Agents require **publish_depth_first_all** to be set to *True* in order to capture data.

BACnet Driver Configuration

Communicating with BACnet devices requires that the BACnet Proxy Agent is configured and running. All device communication happens through this agent.

driver_config

There are nine arguments for the “driver_config” section of the device configuration file:

- **device_address** - Address of the device. If the target device is behind an IP to MS/TP router then Remote Station addressing will probably be needed for the driver to find the device.
- **device_id** - BACnet ID of the device. Used to establish a route to the device at startup.
- **min_priority** - (Optional) Minimum priority value allowed for this device whether specifying the priority manually or via the registry config. Violating this parameter either in the configuration or when writing to the point will result in an error. Defaults to 8.
- **max_per_request** - (Optional) Configure driver to manually segment read requests. The driver will only grab up to the number of objects specified in this setting at most per request. This setting is primarily for scraping many points off of low resource devices that do not support segmentation. Defaults to 10000.
- **proxy_address** - (Optional) VIP address of the BACnet proxy. Defaults to “platform.bacnet_proxy”. See *Communicating With Multiple BACnet Networks* for details. Unless your BACnet network has special needs you should not change this value.
- **ping_retry_interval** - (Optional) The driver will ping the device to establish a route at startup. If the BACnet proxy is not available the driver will retry the ping at this interval until it succeeds. Defaults to 5.
- **use_read_multiple** - (Optional) During a scrape the driver will tell the proxy to use a ReadPropertyMultipleRequest to get data from the device. Otherwise the proxy will use multiple ReadPropertyRequest calls. If the BACnet proxy is reporting a device is rejecting requests try changing this to false for that device. Be aware that setting this to false will cause scrapes for that device to take much longer. Only change if needed. Defaults to true.
- **cov_lifetime** - (Optional) When a device establishes a change of value subscription for a point, this argument will be used to determine the lifetime and renewal period for the subscription, in seconds. Defaults to 180. (Added to Master Driver version 3.2)

Here is an example device configuration file:

```
{
  "driver_config": {"device_address": "10.1.1.3",
                    "device_id": 500,
                    "min_priority": 10,
                    "max_per_request": 24
                   },
  "driver_type": "bacnet",
  "registry_config": "config://registry_configs/vav.csv",
  "interval": 5,
  "timezone": "UTC",
  "heart_beat_point": "heartbeat"
}
```

A sample BACnet configuration file can be found [here](#) or in the VOLTRON repository in `examples/configurations/drivers/bacnet1.config`

BACnet Registry Configuration File

The registry configuration file is a CSV file. Each row configures a point on the device.

Most of the configuration file can be generated with the `grab_bacnet_config.py` utility in `scripts/bacnet`. See BACnet-Auto-Configuration.

Currently, the driver provides no method to access array type properties even if the members of the array are of a supported type.

The following columns are required for each row:

- **Volttron Point Name** - The name by which the platform and agents running on the platform will refer to this point. For instance, if the Volttron Point Name is HeatCall1 (and using the example device configuration above) then an agent would use “pnnl/isb2/hvac1/HeatCall1” to refer to the point when using the RPC interface of the actuator agent.
- **Units** - Used for meta data when creating point information on the historian.
- **BACnet Object Type** - A string representing what kind of BACnet standard object the point belongs to. Examples include:
 - analogInput
 - analogOutput
 - analogValue
 - binaryInput
 - binaryOutput
 - binaryValue
 - multiStateValue
- **Property** - A string representing the name of the property belonging to the object. Usually, this will be “presentValue”.
- **Writable** - Either “TRUE” or “FALSE”. Determines if the point can be written to. Only points labeled TRUE can be written to through the ActuatorAgent. Points labeled “TRUE” incorrectly will cause an error to be returned when an agent attempts to write to the point.
- **Index** - Object ID of the BACnet object.

The following columns are optional:

- **Write Priority** - BACnet priority for writing to this point. Valid values are 1-16. Missing this column or leaving the column blank will use the default priority of 16.
- **COV Flag** - Either “True” or False”. Determines if a BACnet Change of Value subscription should be established for this point. Missing this column or leaving the column blank will result in no change of value subscriptions being established. (Added to Master Driver version 3.2)

Any additional columns will be ignored. It is common practice to include a **Point Name** or **Reference Point Name** to include the device documentation’s name for the point and **Notes** and **Unit Details**” for additional information about a point.

Table 16: BACnet

Point Name	Volttron Point Name	Units	Unit Details	BACnet Object Type	Property	Writable	Index	Notes
Building/FCB.Local Application.PH-T	PreheatTemperature	degreesFahrenheit	-50.00 to 250.00	analogInput	present-Value	FALSE	30001	Resolution: 0.1
Building/FCB.Local Application.RA-T	ReturnAirTemperature	degreesFahrenheit	-50.00 to 250.00	analogInput	present-Value	FALSE	30001	Resolution: 0.1
Building/FCB.Local Application.RA-H	ReturnAirHumidity	percentRelativeHumidity	0.00 to 100.00	analogInput	present-Value	FALSE	30001	Resolution: 0.1
Building/FCB.Local Application.CLG-O	CoolingValveOutputCommand	percent	0.00 to 100.00 (default 0.0)	analogOutput	present-Value	TRUE	30001	Resolution: 0.1
Building/FCB.Local Application.MAD-O	MixedAirDamperOutputCommand	percent	0.00 to 100.00 (default 0.0)	analogOutput	present-Value	TRUE	30001	Resolution: 0.1
Building/FCB.Local Application.PH-O	PreheatValveOutputCommand	percent	0.00 to 100.00 (default 0.0)	analogOutput	present-Value	TRUE	30001	Resolution: 0.1
Building/FCB.Local Application.RH-O	ReheatValveOutputCommand	percent	0.00 to 100.00 (default 0.0)	analogOutput	present-Value	TRUE	30001	Resolution: 0.1
Building/FCB.Local Application.SF-O	SupplyFanSpeedOutputCommand	percent	0.00 to 100.00 (default 0.0)	analogOutput	present-Value	TRUE	30001	Resolution: 0.1

A sample BACnet registry file can be found [here](#) or in the VOLTRON repository in `examples/configurations/drivers/bacnet.csv`

Chargepoint Driver Configuration

The chargepoint driver requires at least one additional python library and has its own `requirements.txt`. Make sure to run `pip install -r <chargepoint driver path>/requirements.txt` before using this driver.

driver_config

There are three arguments for the **driver_config** section of the device configuration file:

- **stationID** - Chargepoint ID of the station. This format is usually '1:00001'
- **username** - Login credentials for the Chargepoint API
- **password** - Login credentials for the Chargepoint API

The Chargepoint login credentials are generated in the Chargepoint web portal and require a chargepoint account with sufficient privileges. Station IDs are also available on the web portal.

Here is an example device configuration file:

```
{
  "driver_config": {"stationID": "3:12345",
                    "username":
↪ "4b90fc0ae5fe8b6628e50af1215d4fcf5743a6f3c63ee1464012875",
                    "password": "ebaf1a3cdfb80baf5b274bdf831e2648"},
  "driver_type": "chargepoint",
  "registry_config": "config://chargepoint.csv",
  "interval": 60,
  "timezone": "UTC",
  "heart_beat_point": "heartbeat"
}
```

A sample Chargepoint configuration file can be found in the VOLTTRON repository in `examples/configurations/drivers/chargepoint1.config`

Chargepoint Registry Configuration File

The registry configuration file is a **CSV** file. Each row configures a point on the device.

The following columns are required for each row:

- **Volttron Point Name** - The name by which the platform and agents running on the platform will refer to this point.
- **Attribute Name** - Chargepoint API attribute name. This determines the field that will be read from the API response and must be one of the allowed values.
- **Port #** - If the point describes a specific port on the Chargestation, it is defined here. (Note 0 and an empty value are equivalent.)
- **Type** - Python type of the point value.
- **Units** - Used for meta data when creating point information on the historian.
- **Writable** - Either "TRUE" or "FALSE". Determines if the point can be written to. Only points labeled TRUE can be written.
- **Notes** - Miscellaneous notes field.
- **Register Name** - A string representing how to interpret the data register. Acceptable values are:
 - StationRegister
 - StationStatusRegister
 - LoadRegister
 - AlarmRegister

– StationRightsRegister

- **Starting Value** - Default value for writeable points. Read-only points should not have a value in this column.

Detailed descriptions for all available chargepoint registers may be found in the `README.rst` in the chargepoint driver directory.

A sample Chargepoint registry file can be found in the VOLTRON repository in `examples/configurations/drivers/chargepoint.csv`

DNP3 Driver Configuration

VOLTRON's DNP3 driver enables the use of **DNP3** (Distributed Network Protocol) communications, reading and writing points via a DNP3 Outstation.

In order to use a DNP3 driver to read and write point data, VOLTRON's DNP3Agent must also be configured and running. All communication between the VOLTRON Outstation and a DNP3 Master happens through this DNP3Agent. For information about the DNP3Agent, please see the DNP3 Platform Specification.

driver_config

There is one argument for the "driver_config" section of the DNP3 driver configuration file:

- **dnp3_agent_id** - ID of VOLTRON's DNP3Agent.

Here is a sample DNP3 driver configuration file:

```
{
  "driver_config": {
    "dnp3_agent_id": "dnp3agent"
  },
  "campus": "campus",
  "building": "building",
  "unit": "dnp3",
  "driver_type": "dnp3",
  "registry_config": "config://dnp3.csv",
  "interval": 15,
  "timezone": "US/Pacific",
  "heart_beat_point": "Heartbeat"
}
```

A sample DNP3 driver configuration file can be found in the VOLTRON repository in `services/core/MasterDriverAgent/example_configurations/test_dnp3.config`.

DNP3 Registry Configuration File

The driver's registry configuration file, a **CSV** file, specifies which DNP3 points the driver will read and/or write. Each row configures a single DNP3 point.

The following columns are required for each row:

- **Volttron Point Name** - The name used by the VOLTRON platform and agents to refer to the point.
- **Group** - The point's DNP3 group number.
- **Index** - The point's index number within its DNP3 data type (which is derived from its DNP3 group number).
- **Scaling** - A factor by which to multiply point values.

- **Units** - Point value units.
- **Writable** - TRUE or FALSE, indicating whether the point can be written by the driver (FALSE = read-only).

Consult the **DNP3 data dictionary** for a point's Group and Index values. Point definitions in the data dictionary are by agreement between the DNP3 Outstation and Master. The VOLTTRON DNP3Agent loads the data dictionary of point definitions from the JSON file at "point_definitions_path" in the DNP3Agent's config file.

A sample data dictionary is available in `services/core/DNP3Agent/dnp3/mesa_points.config`.

Point definitions in the DNP3 driver's registry should look something like this:

```
Voltron Point Name,Group,Index,Scaling,Units,Writable
DCHD.WTgt,41,65,1.0,NA,FALSE
DCHD.WTgt-In,30,90,1.0,NA,TRUE
DCHD.WinTms,41,66,1.0,NA,FALSE
DCHD.RmpTms,41,67,1.0,NA,FALSE
```

A sample DNP3 driver registry configuration file is available in `services/core/MasterDriverAgent/example_configurations/dnp3.csv`.

Fake Device Driver Configuration

This driver does not connect to any actual device and instead produces random and or pre-configured values.

driver_config

There are no arguments for the "driver_config" section of the device configuration file. The driver_config entry must still be present and should be left blank

Here is an example device configuration file:

```
{
  "driver_config": {},
  "driver_type": "bacnet",
  "registry_config": "config://registry_configs/vav.csv",
  "interval": 5,
  "timezone": "UTC",
  "heart_beat_point": "heartbeat"
}
```

A sample fake device configuration file can be found in the VOLTTRON repository in `examples/configurations/drivers/fake.config`

Fake Device Registry Configuration File

The registry configuration file is a **CSV** file. Each row configures a point on the device.

The following columns are required for each row:

- **Volttron Point Name** - The name by which the platform and agents running on the platform will refer to this point. For instance, if the Volttron Point Name is HeatCall1 (and using the example device configuration above) then an agent would use `pnnl/isb2/hvac1/HeatCall1` to refer to the point when using the RPC interface of the actuator agent.
- **Units** - Used for meta data when creating point information on the historian.

- **Writable** - Either “TRUE” or “FALSE”. Determines if the point can be written to. Only points labeled TRUE can be written to through the ActuatorAgent. Points labeled “TRUE” incorrectly will cause an error to be returned when an agent attempts to write to the point.

The following columns are optional:

- **Starting Value** - Initial value for the point. If the point is reverted it will change back to this value. By default, points will start with a random value (1-100).
- **Type** - Value type for the point. Defaults to “string”. Valid types are:
 - string
 - integer
 - float
 - boolean

Any additional columns will be ignored. It is common practice to include a **Point Name** or **Reference Point Name** to include the device documentation’s name for the point and **Notes** and **Unit Details** for additional information about a point. Please note that there is nothing in the driver that will enforce anything specified in the **Unit Details** column.

Table 17: BACnet

Volttron Point Name	Units	Units Details	Writable	Starting Value	Type	Notes
Heartbeat	On/Off	On/Off	TRUE	0	boolean	Point for heartbeat toggle
Out-sideAirTemperature1	F	-100 to 300	FALSE	50	float	CO2 Reading 0.00-2000.0 ppm
SampleWritable-Float1	PPM	10.00 (default)	TRUE	10	float	Setpoint to enable demand control ventilation
SampleLong1	Enumeration	1 through 13	FALSE	50	int	Status indicator of service switch
SampleWritableShort1	%	0.00 to 100.00 (20 default)	TRUE	20	int	Minimum damper position during the standard mode
SampleBool1	On / Off	on/off	FALSE	TRUE	boolean	Status indicator of cooling stage 1
SampleWritable-Bool1	On / Off	on/off	TRUE	TRUE	boolean	Status indicator

A sample fake registry configuration file can be found [here](#) or in the VOLTRON repository in `examples/configurations/drivers/fake.csv`

Modbus Driver Configuration

VOLTRON’s modbus driver supports the Modbus over TCP/IP protocol only. For Modbus RTU support, see VOLTRON’s modbus-tk driver.

driver_config

There are three arguments for the **driver_config** section of the device configuration file:

- **device_address** - IP Address of the device.

- **port** - Port the device is listening on. Defaults to 502 which is the standard port for MODBUS devices.
- **slave_id** - Slave ID of the device. Defaults to 0. Use 0 for no slave.

Here is an example device configuration file:

```
{
  "driver_config": {"device_address": "10.1.1.2",
                   "port": 502,
                   "slave_id": 5},
  "driver_type": "modbus",
  "registry_config": "config://registry_configs/hvac.csv",
  "interval": 60,
  "timezone": "UTC",
  "heart_beat_point": "heartbeat"
}
```

A sample MODBUS configuration file can be found in the VOLTTRON repository in `examples/configurations/drivers/modbus1.config`

Modbus Registry Configuration File

The registry configuration file is a **CSV** file. Each row configures a point on the device.

The following columns are required for each row:

- **Volttron Point Name** - The name by which the platform and agents running on the platform will refer to this point. For instance, if the Volttron Point Name is HeatCall1 (and using the example device configuration above) then an agent would use `pnnl/isb2/hvac1/HeatCall1` to refer to the point when using the RPC interface of the actuator agent.
- **Units** - Used for meta data when creating point information on the historian.
- **Modbus Register** - A string representing how to interpret the data register and how to read it from the device. The string takes two forms:
 - “**BOOL**” for coils and discrete inputs.
 - A format string for the Python struct module. See <http://docs.python.org/2/library/struct.html> for full documentation. The supplied format string must only represent one value. See the documentation of your device to determine how to interpret the registers. Some Examples:
 - * “>f” - A big endian 32-bit floating point number.
 - * “<H” - A little endian 16-bit unsigned integer.
 - * “>I” - A big endian 32-bit integer.
- **Writable** - Either “**TRUE**” or “**FALSE**”. Determines if the point can be written to. Only points labeled **TRUE** can be written to through the ActuatorAgent.
- **Point Address** - Modbus address of the point. Cannot include any offset value, it must be the exact value of the address.
- **Mixed Endian** - (Optional) Either “**TRUE**” or “**FALSE**”. For mixed endian values. This will reverse the order of the MODBUS registers that make up this point before parsing the value or writing it out to the device. Has no effect on bit values.

The following column is optional:

- **Default Value** - The default value for the point. When the point is reverted by an agent it will change back to this value. If this value is missing it will revert to the last known value not set by an agent.

Any additional columns will be ignored. It is common practice to include a **Point Name** or **Reference Point Name** to include the device documentation's name for the point and **Notes** and **Unit Details** for additional information about a point.

The following is an example of a MODBUS registry configuration file:

Table 18: Catalyst 371

Reference Point Name	Volttron Point Name	Units	Units Details	Modbus Register	Writable	Point Address	Default Value	Notes
CO2Sensor	ReturnAirCO2	PPM	0.00-2000.00	>f	FALSE	E1001		CO2 Reading 0.00-2000.0 ppm
CO2Stpt	ReturnAirCO2Stpt	PPM	1000.00 (default)	>f	TRUE	E1011	1000	Setpoint to enable demand control ventilation
Cool1Spd	Cool-Supply-FanSpeed1	%	0.00 to 100.00 (75 default)	>f	TRUE	E1005	75	Fan speed on cool 1 call
Cool2Spd	Cool-Supply-FanSpeed2	%	0.00 to 100.00 (90 default)	>f	TRUE	E1007	90	Fan speed on Cool2 Call
Damper	DamperSignal	%	0.00 - 100.00	>f	FALSE	E1023		Output to the economizer damper
DaTemp	DischargeAirTemperature	F	(-)39.99 to 248.00	>f	FALSE	E1009		Discharge air reading
ESME-conMin	ESM-Damper-MinPosition	%	0.00 to 100.00 (5 default)	>f	TRUE	E1013	5	Minimum damper position during the energy savings mode
FanPower	Supply-FanPower	kW	0.00 to 100.00	>f	FALSE	E1015		Fan power from drive
FanSpeed	Supply-FanSpeed	%	0.00 to 100.00	>f	FALSE	E1003		Fan speed from drive
HeatCall1	HeatCall1	On / Off	on/off	BOOL	FALSE	E1113		Status indicator of heating stage 1 need
HeartBeat	heartbeat	On / Off	on/off	BOOL	FALSE	E1114		Status indicator of heating stage 2 need

A sample MODBUS registry file can be found [here](#) or in the VOLTRON repository in `examples/configurations/drivers/catalyst371.csv`

Modbus-TK Driver Configuration

VOLTRON's Modbus-TK driver, built on the Python Modbus-TK library, is an alternative to the original VOLTRON modbus driver. Unlike the original modbus driver, the Modbus-TK driver supports Modbus RTU as well as Modbus over TCP/IP.

The Modbus-TK driver introduces a map library and configuration builder, intended as a way to streamline configuration file creation and maintenance.

The Modbus-TK driver is mostly backward-compatible with the parameter definitions in the original Modbus driver's configuration (.config and .csv files). If the config file's parameter names use the Modbus driver's name conventions, they are translated to the Modbus-TK name conventions, e.g. a Modbus CSV file's "Point Address" is interpreted as a Modbus-TK "Address". Backward-compatibility exceptions are:

- If the config file has no **port**, the default is 0, not 502.
- If the config file has no **slave_id**, the default is 1, not 0.

driver_config

The **driver_config** section of a Modbus-TK device configuration file supports a variety of parameter definitions, but only **device_address** is required:

- **name** (Optional) - Name of the device. Defaults to "UNKNOWN".
- **device_type** (Optional) - Name of the device type. Defaults to "UNKNOWN".
- **device_address** (Required) - IP Address of the device.
- **port** (Optional) - Port the device is listening on. Defaults to 0 (no port). Use port 0 for RTU transport.
- **slave_id** (Optional) - Slave ID of the device. Defaults to 1. Use ID 0 for no slave.
- **baudrate** (Optional) - Serial (RTU) baud rate. Defaults to 9600.
- **bytesize** (Optional) - Serial (RTU) byte size: 5, 6, 7, or 8. Defaults to 8.
- **parity** (Optional) - Serial (RTU) parity: none, even, odd, mark, or space. Defaults to none.
- **stopbits** (Optional) - Serial (RTU) stop bits: 1, 1.5, or 2. Defaults to 1.
- **xonxoff** (Optional) - Serial (RTU) flow control: 0 or 1. Defaults to 0.
- **addressing (Optional) - Data address table: offset, offset_plus, or address. Defaults to offset.**
 - address: The exact value of the address without any offset value.
 - offset: The value of the address plus the offset value.
 - offset_plus: The value of the address plus the offset value plus one.
 - : **If an offset value is to be added, it is determined based on a point's properties in the CSV file:**
 - * Type=bool, Writable=TRUE: 0
 - * Type=bool, Writable=FALSE: 10000
 - * Type!=bool, Writable=TRUE: 30000
 - * Type!=bool, Writable=FALSE: 40000
- **endian** (Optional) - Byte order: big or little. Defaults to big.
- **write_multiple_registers (Optional) - Write multiple coils or registers at a time. Defaults to true.**
 - : If write_multiple_registers is set to false, only register types unsigned short (uint16) and boolean (bool) are supported. The exception raised during the configure process.
- **register_map** (Optional) - Register map csv of unchanged register variables. Defaults to registry_config csv.

Sample Modbus-TK configuration files are checked into the VOLTRON repository in `services/core/MasterDriverAgent/master_driver/interfaces/modbus_tk/maps`.

Here is a sample TCP/IP Modbus-TK device configuration:

```
{
  "driver_config": {
    "device_address": "10.1.1.2",
    "port": "5020",
    "register_map": "config://modbus_tk_test_map.csv"
  },
  "driver_type": "modbus_tk",
  "registry_config": "config://modbus_tk_test.csv",
  "interval": 60,
  "timezone": "UTC",
  "heart_beat_point": "heartbeat"
}
```

Here is a sample RTU Modbus-TK device configuration, using all default settings:

```
{
  "driver_config": {
    "device_address": "/dev/tty.usbserial-AL00IEEY",
    "register_map": "config://modbus_tk_test_map.csv"
  },
  "driver_type": "modbus_tk",
  "registry_config": "config://modbus_tk_test.csv",
  "interval": 60,
  "timezone": "UTC",
  "heart_beat_point": "heartbeat"
}
```

Here is a sample RTU Modbus-TK device configuration, with completely-specified settings:

```
{
  "driver_config": {
    "device_address": "/dev/tty.usbserial-AL00IEEY",
    "port": 0,
    "slave_id": 2,
    "name": "watts_on",
    "baudrate": 115200,
    "bytesize": 8,
    "parity": "none",
    "stopbits": 1,
    "xonxoff": 0,
    "addressing": "offset",
    "endian": "big",
    "write_multiple_registers": true,
    "register_map": "config://watts_on_map.csv"
  },
  "driver_type": "modbus_tk",
  "registry_config": "config://watts_on.csv",
  "interval": 120,
  "timezone": "UTC"
}
```

Modbus-TK Register Map CSV File

The registry configuration file is a **CSV** file. Each row configures a register definition on the device.

- **Register Name** (Required) - The field name in the modbus client. This field is distinct and unchangeable.
- **Address** (Required) - The point's modbus address. The **addressing** option in the driver configuration controls whether this is interpreted as an exact address or an offset.
- **Type** (Required) - The point's data type: bool, string[length], float, int16, int32, int64, uint16, uint32, or uint64.
- **Units** (Optional) - Used for metadata when creating point information on a historian. Default is an empty string.
- **Writable** (Optional) - TRUE/FALSE. Only points for which Writable=TRUE can be updated by a VOLTTRON agent. Default is FALSE.
- **Default Value** (Optional) - The point's default value. If it is reverted by an agent, it changes back to this value. If this value is missing, it will revert to the last known value not set by an agent.
- **Transform** (Optional) - Scaling algorithm: scale(multiplier), scale_int(multiplier), mod10k(reverse), or none. Default is an empty string.
- **Table** (Optional) - Standard modbus table name defining how information is stored in slave device. There are 4 different tables:
 - discrete_output_coils: read/write coil numbers 1-9999
 - discrete_input_contacts: read only coil numbers 10001-19999
 - analog_input_registers: read only register numbers 30001-39999
 - analog_output_holding_registers: read/write register numbers 40001-49999

If this field is empty, the modbus table will be defined by **type** and **writable** fields. By that, when user sets read only for read/write coils/registers or sets read/write for read only coils/registers, it will select wrong table, and therefore raise exception.

- **Mixed Endian** (Optional) - TRUE/FALSE. If Mixed Endian is set to TRUE, the order of the MODBUS registers will be reversed before parsing the value or writing it out to the device. By setting mixed endian, transform must be None (no op). Defaults to FALSE.
- **Description** (Optional) - Additional information about the point. Default is an empty string.

Any additional columns are ignored.

Sample Modbus-TK registry CSV files are checked into the VOLTTRON repository in `services/core/MasterDriverAgent/master_driver/interfaces/modbus_tk/maps`.

Here is a sample Modbus-TK registry configuration:

Register Name	Address	Type	Units	Writable	Default Value	Transform	Table
unsigned_short	0	uint16	None	TRUE	0	scale(10)	ana-log_output_holding_registers
unsigned_int	1	uint32	None	TRUE	0	scale(10)	ana-log_output_holding_registers
unsigned_long	3	uint64	None	TRUE	0	scale(10)	ana-log_output_holding_registers
sample_short	7	int16	None	TRUE	0	scale(10)	ana-log_output_holding_registers
sample_int	8	int32	None	TRUE	0	scale(10)	ana-log_output_holding_registers
sample_float	10	float	None	TRUE	0.0	scale(10)	ana-log_output_holding_registers
sample_long	12	int64	None	TRUE	0	scale(10)	ana-log_output_holding_registers
sample_bool	16	bool	None	TRUE	False		ana-log_output_holding_registers
sample_str	17	string[12]	None	TRUE	hello world!		ana-log_output_holding_registers

Modbus-TK Registry Configuration CSV File

The registry configuration file is a [CSV](#) file. Each row configures a point on the device.

- **Volttron Point Name** (Required) - The name by which the platform and agents refer to the point. For instance, if the Volttron Point Name is HeatCall1, then an agent would use `my_campus/building2/hvac1/HeatCall1` to refer to the point when using the RPC interface of the actuator agent.
- **Register Name** (Required) - The field name in the modbus client. It must be matched with the field name from `register_map`.

Any additional columns will override the existed fields from `register_map`.

Sample Modbus-TK registry CSV files are checked into the VOLTRON repository in `services/core/MasterDriverAgent/master_driver/interfaces/modbus_tk/maps`.

Here is a sample Modbus-TK registry configuration with defined `register_map`:

Volttron Point Name	Register Name
unsigned short	unsigned_short
unsigned int	unsigned_int
unsigned long	unsigned_long
sample short	sample_short
sample int	sample_int
sample float	sample_float
sample long	sample_long
sample bool	sample_bool
sample str	sample_str

Modbus-TK Driver Maps

To help facilitate the creation of VOLTRON device configuration entries (.config files) for Modbus-TK devices, a library of device type definitions is now maintained in `services/core/MasterDriverAgent/master_driver/interfaces/modbus_tk/maps/maps.yaml`. A command-line tool (described below under **MODBUS TK Config Command Tool**) uses the contents of `maps.yaml` while generating .config files.

Each device type definition in `maps.yaml` consists of the following properties:

- **name** (Required) - Name of the device type (see the `driver_config` parameters).
- **file** (Required) - The name of the CSV file that defines all of the device type's supported points, e.g. `watts_on.csv`.
- **description** (Optional) - A description of the device type.
- **addressing** (Optional) - Data address type: `offset`, `offset_plus`, or `address` (see the `driver_config` parameters).
- **endian** (Optional) - Byte order: `big` or `little` (see the `driver_config` parameters).
- **write_multiple_registers** (Optional) - Write multiple registers at a time. Defaults to `true`.

A device type definition is a template for a device configuration. Some additional data must be supplied when a specific device's configuration is generated. In particular, the `device_address` must be supplied.

A sample `maps.yaml` file is checked into the VOLTRON repository in `services/core/MasterDriverAgent/master_driver/interfaces/modbus_tk/maps/maps.yaml`.

Here is a sample `maps.yaml` file:

```
- name: modbus_tk_test
  description: Example of reading selected points for Modbus-TK driver testing
  file: modbus_tk_test_map.csv
  addressing: offset
  endian: little
  write_multiple_registers: true
- name: watts_on
  description: Read selected points from Elkor WattsOn meter
  file: watts_on_map.csv
  addressing: offset
- name: ion6200
  description: ION 6200 meter
  file: ion6200_map.csv
- name: ion8600
  description: ION 8600 meter
  file: ion8600_map.csv
```

Modbus-TK Config Command Tool

`config_cmd.py` is a command-line tool for creating and maintaining VOLTRON driver configurations. The tool runs from the command line:

```
$ cd services/core/MasterDriverAgent/master_driver/interfaces/modbus_tk/maps
$ python config_cmd.py
```

`config_cmd.py` supports the following commands:

- **help** - List all commands.
- **quit** - Quit the command-line tool.

- **list_directories** - List all setup directories, with an option to edit their paths.
 - By default, all directories are in the VOLTRON repository in `services/core/MasterDriverAgent/master_driver/interfaces/modbus_tk/maps`.
 - It is important to use the correct directories when adding/editing device types and driver configs, and when loading configurations into VOLTRON.
 - * `map_dir`: directory in which `maps.yaml` is stored.
 - * `config_dir`: directory in which driver config files are stored.
 - * `csv_dir`: directory in which registry config CSV files are stored.
- **edit_directories** - Add/Edit map directory, driver config directory, and/or CSV config directory. Press <Enter> if no change is needed. Exits if the directory does not exist.
- **list_device_type_description** - List all device type descriptions in `maps.yaml`. Option to edit device type descriptions.
- **list_all_device_types** - List all device type information in `maps.yaml`. Option to add more device types.
- **device_type** - List information for a selected device type. Option to select another device type.
- **add_device_type** - Add a device type to `maps.yaml`. Option to add more than one device type. Each device type includes its name, CSV file, description, addressing, and endian, as explained in **MODBUS-TK Driver Maps**. If an invalid value is entered for addressing or endian, the default value is used instead.
- **edit_device_type** - Edit an existing device type. If an invalid value is entered for addressing or endian, the previous value is left unchanged.
- **list_drivers** - List all driver config names in `config_dir`.
- **driver_config <driver_name>** - Get a driver config from `config_dir`. Option to select the driver if no driver is found with that name.
- **add_driver_config <driver_name>** - Add/Edit `<config_dir>/<driver name>.config`. Option to select the driver if no driver is found with that name. Press <Enter> to exit.
- **load_volttron** - Load a driver config and CSV into VOLTRON. Option to add the config or CSV file to `config_dir` or to `csv_dir`. VOLTRON must be running when this command is used.
- **delete_volttron_config** - Delete a driver config from VOLTRON. VOLTRON must be running when this command is used.
- **delete_volttron_csv** - Delete a registry csv config from VOLTRON. VOLTRON must be running when this command is used.

The `config_cmd.py` module is checked into the VOLTRON repository as `services/core/MasterDriverAgent/master_driver/interfaces/modbus_tk/config_cmd.py`.

Obix Driver Configuration

VOLTRON's uses Obix's restful interface to facilitate communication.

This driver does *not* handle reading data from the history section of the interface. If the user wants data published from the management systems historical data use the *Obix History Agent* agent.

driver_config

There are three arguments for the **driver_config** section of the device configuration file:

- **url** - URL of the interface.
- **username** - User name for site..
- **password** - Password for username.

Here is an example device configuration file:

```
{
  "driver_config": {"url": "http://example.com/obix/config/Drivers/Obix/exports/",
                   "username": "username",
                   "password": "password"},
  "driver_type": "obix",
  "registry_config": "config://registry_configs/obix.csv",
  "interval": 30,
  "timezone": "UTC"
}
```

A sample Obix configuration file can be found in the VOLTTRON repository in `examples/configurations/drivers/obix.config`

Obix Registry Configuration File

The registry configuration file is a **CSV** file. Each row configures a point on the device.

The following columns are required for each row:

- **Volttron Point Name** - The name by which the platform and agents running on the platform will refer to this point. For instance, if the Volttron Point Name is HeatCall1 then an agent would use `<device topic>/HeatCall1` to refer to the point when using the RPC interface of the actuator agent.
- **Obix Point Name** - Name of the point on the obix interface. Escaping of spaces and dashes for use with the interface is handled internally.
- **Obix Type** - One of “bool”, “int”, or “real” without quotes.
- **Units** - Used for meta data when creating point information on the historian.
- **Writable** - Either “TRUE” or “FALSE”. Determines if the point can be written to. Only points labeled TRUE can be written to through the ActuatorAgent. This can be used to protect points that should not be accessed by the platform.

The following column is optional:

- **Default Value** - The default value for the point. When the point is reverted by an agent it will change back to this value. If this value is missing it will revert to the last known value not set by an agent.

Any additional columns will be ignored. It is common practice to include a **Point Name** or **Reference Point Name** to include the device documentation’s name for the point and **Notes** and **Unit Details** for additional information about a point.

The following is an example of a Obix registry configuration file:

Table 19: Obix

Voltron Point Name	Obix Point Name	Obix Type	Units	Writable	Notes
CostEL	CostEL	real	dollar	FALSE	Precision: 2
CostELBB	CostELBB	real	dollar	FALSE	Precision: 2
CDHEnergyHeartbeat	CDHEnergyHeartbeat	real	null	FALSE	
ThermalFollowing	ThermalFollowing	bool		FALSE	
CDHTestThermFollow	CDHTestThermFollow	bool		FALSE	
CollegeModeFromCDH	CollegeModeFromCDH	real	null	FALSE	Precision: 0, Min: 3.0, Max: 3.0
HospitalModeFromCDH	HospitalModeFrom-CDH	real	null	FALSE	Precision: 0, Min: 3.0, Max: 3.0
HomeModeFromCDH	HomeModeFromCDH	real	null	FALSE	Precision: 0, Min: 3.0, Max: 3.0
CostNG	CostNG	real	null	FALSE	Precision: 2
CollegeBaseload-SPFromCDH	CollegeBaseload-SPFromCDH	real	kilo-watt	FALSE	Precision: 0
CollegeImportSPFrom-CDH	CollegeImportSPFrom-CDH	real	kilo-watt	FALSE	Precision: 0
HospitalImportSPFrom-CDH	HospitalImportSPFrom-CDH	real	kilo-watt	FALSE	Precision: 0
HospitalBaseload-SPFromCDH	HospitalBaseload-SPFromCDH	real	kilo-watt	FALSE	Precision: 0
HomeImportSPFrom-CDH	HomeImportSPFrom-CDH	real	kilo-watt	FALSE	Precision: 0
ThermalFollowingAlarm	ThermalFollowingAlarm	bool		FALSE	

A sample Obix configuration can be found in the VOLTRON repository in `examples/configurations/drivers/obix.csv`

Automatic Obix Configuration File Creation

A script that will automatically create both a device and register configuration file for a site is located in the repository at `scripts/obix/get_obix_driver_config.py`.

The utility is invoked with the command:

```
“python get_obix_driver_config.py <url> <registry_file> <driver_file> -u <username> -p <password> “
```

If either the `registry_file` or `driver_file` is omitted the script will output those files to stdout.

If either the `username` or `password` arguments are left out the script will ask for them on the command line before proceeding.

The registry file produced by this script assumes that the *Voltron Point Name* and the *Obix Point Name* have the same value. Also, it is assumed that all points should be read only. Users are expected to fix this as appropriate.

Rainforest Emu2 Driver Configuration

The Emu2 is a device for connecting to and reading data from smart power meters. We have an experimental driver to talk to this device. It requires cloning the Rainforest Automation library which can be found [here](#).

Note: The Emu Serial Api library has its own dependencies which should be installed with pip while the VOLTTRON environment is activated.

The Emu2 device interface is configured as follows. Set *emu_library_path* to the location of the cloned library. *tty* should be set to the name of the Emu2's character special file. One way to find this is to run *dmesg* before and after plugging in the Emu2, and checking the new output.

```
{
  "driver_config": {
    "tty": "ttyACM0",
    "emu_library_path": "/home/volttron/Emu-Serial-Api"
  },
  "driver_type": "rainforestemu2",
  "interval": 30,
  "registry_config": "config://emu2.json",
  "timezone": "UTC"
}
```

The registry config file referred to in the first configuration must be an array of strings. This tells the interface which data points should be retrieved from the device every interval. If the NetworkInfo point is omitted it will be included automatically.

```
[
  "NetworkInfo",
  "InstantaneousDemand",
  "PriceCluster"
]
```

SEP2 Driver Configuration

Communicating with SEP2 devices requires that the SEP2 Agent is configured and running. All device communication happens through this agent. For information about the SEP2 Agent, please see *SEP 2.0 DER Support*.

driver_config

There are two arguments for the “driver_config” section of the SEP2 device configuration file:

- **sfdi** - Short-form device ID of the SEP2 device.
- **sep2_agent_id** - ID of VOLTTRON's SEP2 agent.

Here is a sample SEP2 device configuration file:

```
{
  "driver_config": {
    "sfdi": "097935300833",
    "sep2_agent_id": "sep2agent"
  },
  "campus": "campus",
  "building": "building",
  "unit": "sep2",
  "driver_type": "sep2",
  "registry_config": "config://sep2.csv",
}
```

(continues on next page)

(continued from previous page)

```

"interval": 15,
"timezone": "US/Pacific",
"heart_beat_point": "Heartbeat"
}

```

A sample SEP2 driver configuration file can be found in the VOLTRON repository in `services/core/MasterDriverAgent/example_configurations/test_sep2_1.config`.

SEP2 Registry Configuration File

For a description of SEP2 registry values, see *SEP 2.0 DER Support*.

A sample SEP2 registry configuration file can be found in the VOLTRON repository in `services/core/MasterDriverAgent/example_configurations/sep2.csv`.

3.5.5 VOLTRON Message Bus

VOLTRON Interconnect Protocol

VIP Authentication

VIP (VOLTRON Interconnect Protocol) authentication is implemented in the `auth` module and extends the ZeroMQ Authentication Protocol *ZAP* to *VIP* by including the *ZAP* User-Id in the *VIP* payload, thus allowing peers to authorize access based on *ZAP* credentials. This document does not cover *ZAP* in any detail, but its understanding is fundamental to securely configuring ZeroMQ. While this document will attempt to instruct on securely configuring VOLTRON for use on the Internet, it is recommended that the *ZAP* documentation also be consulted.

Default Encryption

By default, ZeroMQ operates in plain-text mode, without any sort of encryption. While this is okay for in-process and interprocess communications, via UNIX domain sockets, it is insecure for any kind of inter-network communications, especially when traffic must traverse the Internet. Therefore, VOLTRON automatically generates an encryption key and enables *CurveMQ* by default on all TCP connections.

To see VOLTRON's public key run the `voltttron-ctl auth serverkey` command. For example:

```

(voltttron)[user@home]$ voltttron-ctl auth serverkey
FSG7LHhy3v8tdNz3gK35G6-oxUcyln54pYRKu5fBJzU

```

Peer Authentication

ZAP defines a method for verifying credentials exchanged when a connection is initially established. The authentication mechanism provides three main pieces of information useful for authentication:

- **domain:** a name assigned to a locally bound address (to which peers connect)
- **address:** the remote address of the peer
- **credentials:** includes the authentication method and any associated credentials

During authentication, VOLTRON checks these pieces against a list of accepted peers defined in a file, called the “auth file” in this document. This JSON-formatted file is located at `$VOLTRON_HOME/auth.json` and must have a matching entry in the allow list for remote connections to be accepted.

The auth file should not be modified directly. To change the auth file, use `voltttron-ctl auth` subcommands: `add`, `list`, `remove`, and `update`. (Run `voltttron-ctl auth --help` for more details and see the authentication commands documentation.)

Here are some example entries:

```
(voltttron)[user@home]$ voltttron-ctl auth list

INDEX: 0
{
  "domain": null,
  "user_id": "platform",
  "roles": [],
  "enabled": true,
  "mechanism": "CURVE",
  "capabilities": [],
  "groups": [],
  "address": null,
  "credentials": "k1C9-FPRAVjL-cHliQqAJaCHUNVXaAlkVc7EqK0u9mI",
  "comments": "Automatically added by platform on start"
}

INDEX: 2
{
  "domain": null,
  "user_id": "platform.sysmon",
  "roles": [],
  "enabled": true,
  "mechanism": "CURVE",
  "capabilities": [],
  "groups": [],
  "address": null,
  "credentials": "5UD_GTk5dM2g4pk8d1-wM-BYgt4RAKiHf4SnT_YU6jY",
  "comments": "Automatically added on agent install"
}
```

Note: If using regular expressions in the “address” portion, denote this with “/”. Backslashes must be escaped “\”.

This is a valid regular expression: `"/192\\.168\\.1\\.\\.*/"`

These are invalid: `"/192\\.168\\.1\\.\\.*/"`, `"/192\\.168\\.1\\.\\.*/"`, `"192\\.168\\.1\\.\\.*/"`

When authenticating, the credentials are checked. If they don’t exist or don’t match, authentication fails. Otherwise, if domain and address are not present (or are null), authentication succeeds. If address and/or domain exist, they must match as well for authentication to succeed.

CURVE credentials include the remote peer’s public key. Watching the **INFO** level log output of the auth module can help determine the required values for a specific peer.

Configuring Agents

A remote agent must know the platform’s public key (also called the server key) to successfully authenticate. This server key can be passed to the agent’s `__init__` method in the `serverkey` parameter, but in most scenarios it is preferable to add the server key to the *known-hosts* file.

URL-style Parameters

VOLTRON extends ZeroMQ’s address scheme by supporting URL-style parameters for configuration. The following parameters are supported when connecting:

- `serverkey`: encoded public key of remote server
- `secretkey`: agent’s own private/secret key
- `publickey`: agent’s own public key
- `ipv6`: instructs ZeroMQ to attempt to use IPv6

Note: Although these parameters are still supported they should rarely need to be specified in the VIP-address URL. Agent *key stores* and the *known-hosts file* are automatically used when possible.

Platform Configuration

By default, the platform only listens on the local IPC VIP socket. Additional addresses may be bound using the `--vip-address` option, which can be provided multiple times to bind multiple addresses. Each VIP address should follow the standard ZeroMQ convention of prefixing with the socket type (*ipc://* or *tcp://*) and may include any of the following additional URL parameters:

- `domain`: domain name to associate with this endpoint (defaults to “vip”)
- `secretkey`: alternate private/secret key (defaults to generated key for *tcp://*)
- `ipv6`: instructs ZeroMQ to attempt to use IPv6

Example Setup

Suppose agent A needs to connect to a remote platform B. First, agent A must know platform B’s public key (the *server key*) and platform B’s IP address (including port). Also, platform B needs to know agent A’s public key (let’s say it is `HOVXfTspZWcpHQcYT_xGcqypBHzQHTgqEzVb4iXrcDg`).

Given these values, a user on agent A’s platform adds platform B’s information to the *known-hosts file*.

At this point agent A has all the information needed to connect to platform B, but platform B still needs to add an authentication entry for agent A.

If agent A tried to connect to platform B at this point both parties would see an error. Agent A would see an error similar to:

```
No response to hello message after 10 seconds.
A common reason for this is a conflicting VIP IDENTITY.
Shutting down agent.
```

Platform B (if started with `-v` or `-vv`) will show an error:

```
2016-10-19 14:21:20,934 () voltttron.platform.auth INFO: authentication failure:
↪ domain='vip', address='127.0.0.1', mechanism='CURVE', credentials=[
↪ 'HOVXfTspZWcpHQcYT_xGcqypBHzQHTgqEzVb4iXrcDg']
```

Agent A failed to authenticat to platform B because the platform didn’t have agent A’s public in the authentication list.

To add agent A’s public key, a user on platform B runs:

```
(volttron)[user@platform-b]$ volttron-ctl auth add
domain []:
address []:
user_id []: Agent-A
capabilities (delimit multiple entries with comma) []:
roles (delimit multiple entries with comma) []:
groups (delimit multiple entries with comma) []:
mechanism [CURVE]:
credentials []: HOVXfTspZWcpHQcYT_xGcqypBHzQHTgqEzVb4iXrcDg
comments []:
enabled [True]:
```

Now if agent A can successfully connect to platform B, and platform B's log will show:

```
2016-10-19 14:26:16,446 () volttron.platform.auth INFO: authentication success:
↳ domain='vip', address='127.0.0.1', mechanism='CURVE', credentials=[
↳ 'HOVXfTspZWcpHQcYT_xGcqypBHzQHTgqEzVb4iXrcDg'], user_id='Agent-A'
```

For a more details see the [authentication walkthrough](#).

VIP Authorization

VIP *authentication* and authorization go hand in hand. When an agent authenticates to a VOLTRON platform that agent proves its identity to the platform. Once authenticated, an agent is allowed to connect to the message bus. VIP authorization is about giving a platform owner the ability to limit the capabilities of authenticated agents.

There are two parts to authorization:

1. Required capabilities (specified in agent's code)
2. Authorization entries (specified via `volttron-ctl auth` commands)

The following example will walk through how to specify required capabilities and grant those capabilities in authorization entries.

Single Capability

For this example suppose there is a temperature agent that can read and set the temperature of a particular room. The agent author anticipates that building managers will want to limit which agents can set the temperature.

In the temperature agent, a required capability is specified by using the `RPC.allow` decorator:

```
@RPC.export
def get_temperature():
    ...

@RPC.allow('CAP_SET_TEMP')
@RPC.export
def set_temperature(temp):
    ...
```

In the code above, any agent can call the `get_temperature` method, but only agents with the `CAP_SET_TEMP` capability can call `set_temperature`. (Note: capabilities are arbitrary strings. This example follows the general style used for Linux capabilities, but it is up to the agent author.)

Now that a required capability has been specified, suppose a VOLTRON platform owner wants to allow a specific agent, say `AliceAgent`, to set the temperature.

The platform owner runs `voltttron-ctl auth add` to add new authorization entries or `voltttron-ctl auth update` to update an existing entry. If `AliceAgent` is installed on the platform, then it already has an authorization entry. Running `voltttron-ctl auth list` shows the existing entries:

```
...
INDEX: 3
{
  "domain": null,
  "user_id": "AliceAgent",
  "roles": [],
  "enabled": true,
  "mechanism": "CURVE",
  "capabilities": [],
  "groups": [],
  "address": null,
  "credentials": "JydrFRRv-kdSejL6Ldxy978pOf8HkWC9fRHUWKmJfxc",
  "comments": null
}
...
```

Currently `AliceAgent` cannot set the temperature because it does not have the `CAP_SET_TEMP` capability. To grant this capability the platform owner runs `voltttron-ctl auth update 3`:

```
(For any field type "clear" to clear the value.)
domain []:
address []:
user_id [AliceAgent]:
capabilities (delimit multiple entries with comma) []: CAP_SET_TEMP
roles (delimit multiple entries with comma) []:
groups (delimit multiple entries with comma) []:
mechanism [CURVE]:
credentials [JydrFRRv-kdSejL6Ldxy978pOf8HkWC9fRHUWKmJfxc]:
comments []:
enabled [True]:
updated entry at index 3
```

Now `AliceAgent` can call `set_temperature` via RPC. If other agents try to call that method they will get the following exception:

```
error: method "set_temperature" requires capabilities set(['CAP_SET_TEMP']),
but capability list [] was provided
```

Multiple Capabilities

Expanding on the temperature-agent example, the `set_temperature` method can require agents to have multiple capabilities:

```
@RPC.allow(['CAP_SET_TEMP', 'CAP_FOO_BAR'])
@RPC.export
def set_temperature():
    ...
```

This requires an agent to have both the `CAP_SET_TEMP` and the `CAP_FOO_BAR` capabilities. Multiple capabilities can also be specified by using multiple `RPC.allow` decorators:

```
@RPC.allow('CAP_SET_TEMP')
@RPC.allow('CAN_FOO_BAR')
@RPC.export
def temperature():
    ...
```

VIP Enhancements

Outline a vision of how VOLTTRON Message Bus should work

When creating VIP for VOLTTRON 3.0 we wanted to address two security concerns and one user request:

- Security Concern 1: Agents can spoof each other on the VOLTTRON message bus and fake messages.
- Security Concern 2: Agents can subscribe to topics that they are not authorized to subscribe to.
- User Request 1: Several users requested means to transfer large amounts of data between agents without using the message bus.

VOLTTRON Interconnect Protocol (VIP) was created to address these issues but unfortunately, it broke the easy to use pub-sub messaging model of VOLTTRON. Additionally to use the security features of VOLTTRON in 3.0 code has become an ordeal especially when multiple platforms are concerned. Finally, VIP has introduced the requirement for knowledge of specific other platforms to agents written by users in order to be able to communicate. The rest of this memo focuses on defining the way VOLTTRON message bus will work going forward indefinitely and should be used as the guiding principles for any future work on VIP and VOLTTRON.

VOLTTRON Message Bus Guiding Principles:

1. All communications between two or more different VOLTTRON platforms
MUST go through the VIP Router. Said another way, a user agent (application) should have NO capability to reach out to an agent on a different VOLTTRON platform directly. | All communications between two or more VOLTTRON platforms must be in the form of topics on the message bus. Agents MUST not use a distinct platform address or name to communicate via a direct connection between two platforms.
2. VOLTTRON will use two TCP ports. One port is used to extend VIP across platforms. A second port is used for the VOLTTRON discovery protocol (more on this to come on a different document). VIP will establish bi-directional communication via a single TCP port.
3. In order to solve the bootstrapping problem that CurveMQ has punted on, we will modify VIP to operate similar (behaviorally) to SSH.

A. On a single VOLTTRON platform, the platform's public key will be made available via an API so that all agents will be able to communicate with the platform. Additionally, the behavior of the platform will be changed so that agents on the same platform will automatically be added to auth.json file. No more need for user to add the agents manually to the file. The desired behavior is similar to how SSH handles known_hosts. Note that this behavior still addresses the security request 1 & 2.

B. When connecting VOLTTRON platforms, VOLTTRON Discovery Protocol (VDP) will be used to discover the other platforms public key to establish the router to router connection. Note that since we BANNED agent to agent communication between two platforms, we have prevented an $O(N^2)$ communication pattern and key bootstrapping problem.

1. Authorization determines what agents are allowed to access what topics. Authorization MUST be managed by the VOLTTRON Central platform on a per organization basis. It is not recommended to have different authorization profiles on different VOLTTRON instances belonging to the same organization.

2. VOLTRON message bus uses topics such as and will adopt an information model agreed upon by the VOLTRON community going forward. Our initial information model is based on the OpenEIS schema going forward. A different document will describe the information model we have adopted going forward. All agents are free to create their own topics but the VOLTRON team (going forward) will support the common VOLTRON information model and all agents developed by PNNL will be converted to use the new information model.
3. Two connected VOLTRON systems will exchange a list of available topics via the message router. This will allow each VIP router to know what topics are available at what VOLTRON platform.
4. Even though each VOLTRON platform will have knowledge of what topics are available around itself, no actual messages will be forwarded between VOLTRON platforms until an agent on a specific platform subscribes to a topic. When an agent subscribes to a topic that has a publisher on a different VOLTRON platform, the VIP router will send a request to its peer routers so that the messages sent to that topic will be forwarded. There will be cases (such as clean energy transactive project) where the publisher to a topic may be multiple hops away. In this case, the subscribe request will be sent towards the publisher through other VIP routers. In order to find the most efficient path, we may need to keep track of the total number of hops (in terms of number of VIP routers).
5. The model described in steps 5/6/7 applies to data collection. For control applications, VOLTRON team only allows control actions to be originated from the VOLTRON instance that is directly connected to that controlled device. This decision is made to increase the robustness of the control agent and to encourage truly distributed applications to be developed.
6. Direct agent to agent communication will be supported by creation of an ephemeral topic under the topic hierarchy. Our measurements have shown repeatedly that the overhead of using the ZeroMQ message pub/sub is minimal and has zero impact on communications throughput.

In summary, by making small changes to the way VIP operates, I believe that we can significantly increase the usability of the platform and also correct the mixing of two communication platforms into VIP. VOLTRON message bus will return to being a pub/sub messaging system going forward. Direct agent to agent communication will be supported through the message bus.

VIP Known Identities

It is critical for systems to have known locations for receiving resources and services from in a networked environment. The following table details the vip identities that are reserved for voltrtron specific usage.

VIP Identity	Sphere of Influence	Notes
plat-form	Plat-form	
plat-form.agent	Plat-form	The PlatformAgent is responsible for this identity. It is used to allow the VolttronCentralAgent to control and individual platform.
volt-tron.central	Multi-Network	The default identity for a VolttronCentralAgent. The PlatformAgent by default will use this as it's manager, but can be overridden in the configuration file of individual agents.
plat-form.historian	plat-form	An individual platform may have many historians available to it, however the only one that will be available through Volttron Central by default will be called this. Note that this does not require a specific type of historian, just that it's VIP Identity.
control	plat-form	Control is used to control the individual platform. From the command line when issuing any volttron-ctl operations or when using Volttron Central.
pub-sub	plat-form	Pub/Sub subsystem router
plat-form.actuator	actuator	Agent which coordinates sending control commands to devices.
con-fig.store	plat-form	The configuration subsystem service agent on the platform.
plat-form.driver	devices	The default identity for the Master Driver Agent.

VIP - VOLTTRON™ Interconnect Protocol

This document specifies VIP, the VOLTTRON™ Interconnect Protocol. The use case for VIP is to provide communications between *agents*, *controllers*, *services*, and the supervisory *platform* in an abstract fashion so that additional protocols can be built and used above VIP. VIP defines how *peers* connect to the *router* and the messages they exchange.

- Name: github.com/VOLTTRON/volttron/wiki/VOLTTRON-Interconnect-Protocol
- Editor: Brandon Carpenter <brandon (dot) carpenter (at) pnnl (dot) gov>
- State: draft
- See also: [ZeroMQ](#), [ZMTP](#), [CurveZMQ](#), [ZAP](#)

Preamble

Copyright (c) 2015 Battelle Memorial Institute All rights reserved

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT

SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The views and conclusions contained in the software and documentation are those of the authors and should not be interpreted as representing official policies, either expressed or implied, of the FreeBSD Project.

This material was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the United States Department of Energy, nor Battelle, nor any of their employees, nor any jurisdiction or organization that has cooperated in the development of these materials, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness or any information, apparatus, product, software, or process disclosed, or represents that its use would not infringe privately owned rights.

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or Battelle Memorial Institute. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

PACIFIC NORTHWEST NATIONAL LABORATORY operated by BATTELLE for the UNITED STATES DEPARTMENT OF ENERGY under Contract DE-AC05-76RL01830

The key words “MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”, “SHOULD”, “SHOULD NOT”, “RECOMMENDED”, “MAY”, and “OPTIONAL” in this document are to be interpreted as described in [RFC 2119](#).

Overall Design

What Problems does VIP Address?

When VOLTRON agents, controllers, or other entities needed to exchange data, they previously used the first generation pub/sub messaging mechanism and ad-hoc methods to set up direct connections. While the pub/sub messaging is easy to implement and use, it suffers from several limitations:

- It requires opening two listening sockets: one each for publishing and subscribing.
- There is no trivial way to prevent message spoofing.
- There is no trivial way to enable private messaging
- It is not ideal for peer-to-peer communications.

These limitations have severe security implications. For improved security in VOLTRON, the communications protocol must provide a method for secure data exchange that is intuitive and simple to implement and use.

ZeroMQ already provides many of the building blocks to implement encrypted and authenticated communications over a shared socket. It already includes a socket type implementing the router pattern. What remains is a protocol built on ZeroMQ to provide a single connection point, secure message passing, and retain the ability for entities to come and go as they please.

VIP is just that protocol, specifically targeting the limitations above.

Why ZeroMQ?

Rather than reinvent the wheel, VIP makes use of many features already implemented in ZeroMQ, including ZAP and CurveMQ. While VIP doesn't require the use of ZAP or CurveMQ, their use substantially improves security by encrypting traffic over public networks and limiting connections to authenticated peers.

ZeroMQ also provides reliable transports with built-in framing, automatic reconnection, in-process zero-copy message passing, abstractions for underlying protocols, and so much more. While some of these features create other pain points, they are minimal compared with the effort of either reimplementing or cobbling together libraries.

VIP is a routing protocol

VIP uses the ZeroMQ router pattern. Specifically, the router binds a ROUTER socket and peers connect using a DEALER or ROUTER socket. Unless the peer is connecting a single socket to multiple routers, using the DEALER socket is easiest, but there are instances where using a ROUTER is more appropriate. One must just exercise care to include the proper address envelope to ensure proper routing.

Extensible Security

VIP makes no assumptions about the security mechanisms used. It works equally well over encrypted or unencrypted channels. Any connection-level authentication and encryption is handled by ZAP. Message-level authentication can be implemented in the protocols and services using VIP or by utilizing message properties set in ZAP replies.

ZeroMQ Compatibility

For enhanced security, VOLTTRON recommends libzmq version 4.1 or greater, however, most features of VIP are available with older versions. The following is an incomplete list of core features available with recent versions of libzmq.

- Version 3.2:
 - Basic, unauthenticated, unencrypted routing
 - Use ZMQ_ROUTER_BEHAVIOR socket option instead of ZMQ_ROUTER_MANDATORY
- Version 4.0:
 - Adds authentication and encryption via ZAP
- Version 4.1:
 - Adds message properties allowing correlating authentication tokens to messages

Message Format and Version Detection

VIP uses a simple, multi-frame format for its messages. The first one (for peers) or two (for router) frames contain the delivery address(es) and are followed immediately by the VIP signature `VIP1`. The first characters of the signature are used to match the protocol and the last character digit indicates the protocol version, which will be incremented as the protocol is revised. This allows for fail-fast behavior and backward compatibility while being simple to implement in any language supported by ZeroMQ.

Formal Specification

Architecture

VIP defines a message-based dialog between a *router* that transfers data between *peers*. The *router* and *peers* SHALL communicate using the following socket types and transports:

- The router SHALL use a ROUTER socket.
- Peers SHALL use a DEALER or ROUTER socket.
- The router SHALL bind to one or more endpoints using inproc, tcp, or ipc address types.
- Peers SHALL connect to these endpoints.
- There MAY be any number of peers.

Message Format

A routing exchange SHALL consist of a peer sending a message to the router followed by the router receiving the message and sending it to the destination peer.

Messages sent to the router by peers SHALL consist of the following message frames:

- The *recipient*, which SHALL contain the socket identity of the destination peer.
- The protocol signature, which SHALL contain the four octets “VIP1”.
- The *user id*, which SHALL be an implementation-defined value.
- The *request id*, which SHALL contain an opaque binary blob.
- The *subsystem*, which SHALL contain a string.
- The *data*, which SHALL be zero or more subsystem-specific opaque frames.

Messages received from a peer by the router will automatically have a *sender* frame prepended to the message by the ROUTER socket. When the router forwards the message, the sender and recipient fields are swapped so that the *recipient* is in the first frame and the *sender* is in the second frame. The *recipient* frame is automatically stripped by the ROUTER socket during delivery. Peers using ROUTER sockets must prepend the message with an *intermediary* frame, which SHALL contain the identity of a router socket.

Messages received from the router by peers SHALL consist of the following message frames:

- The *sender*, which SHALL contain the socket identity of the source peer.
- The protocol signature, which SHALL contain the four octets “VIP1”.
- The *user id*, which MAY contain a UTF-8 encoded string.
- The *request id*, which SHALL contain an opaque binary blob.
- The *subsystem*, which SHALL contain a non-empty string.
- The *data*, which SHALL be zero or more subsystem-specific opaque frames.

The various fields have these meanings:

- sender: the ZeroMQ DEALER or ROUTER identity of the sending (source) peer.
- recipient: the ZeroMQ DEALER or ROUTER identity of the recipient (destination) peer.
- intermediary: the ZeroMQ ROUTER identity of the intermediary router.

- user id: VIP authentication metadata set in the authenticator. See the discussion below for more information on this value.
- request id: the meaning of this field is defined by the sending peer. Replies SHALL echo the request id without modifying it.
- subsystem: this specifies the peer subsystem the data is intended for. The length of a subsystem name SHALL NOT exceed 255 characters and MUST only contain ASCII characters.
- data: provides the data for the given subsystem. The number of frames required is defined by each subsystem.

User ID

The value in the *user id* frame depends on the implementation and the version of ZeroMQ. If ZAP is used with libzmq 4.1.0 or newer, peers should send an empty string for the user id and the ZAP authenticator will replace it with an authentication token which receiving peers may use to authorize access. If ZAP is not used or a version of libzmq is used which lacks support for retrieving the user id metadata, an authentication subsystem may be used to authenticate peers. The authentication subsystem SHALL provide peers with private tokens that must be sent with each message in the user id frame and which the router will substitute with a public token before forwarding. If the message cannot be authenticated, the user id received by peers SHALL be a zero-length string.

Socket Types

Peers communicating via the router will typically use DEALER sockets and should not require additional handling. However, a DEALER peer may only connect to a single router. Peers may use ROUTER sockets to connect to multiple endpoints, but must prepend the routing ID of the destination.

When using a DEALER socket:

- A peer SHALL not send in intermediary address.
- A peer SHALL connect to a single endpoint.

When using a ROUTER socket:

- A peer SHALL prepend the intermediary routing ID of to the message frames.
- A peer MAY connect to multiple endpoints.

Routing Identities

Routing identities are set on a socket using the `ZMQ_IDENTITY` socket option and MUST be set on both ROUTER and DEALER sockets. The following additional requirements are placed on the use of peer identities:

- Peers SHALL set a valid identity rather than rely on automatic identity generation.
- The router MAY drop messages with automatically generated identities, which begin with the zero byte ('0').

A zero length identity is invalid for peers and is, therefore, unroutable. It is used instead to address the router itself.

- Peers SHALL use a zero length recipient to address the router.
- Messages sent from the router SHALL have a zero length sender address.

Error Handling

The documented default behavior of ZeroMQ ROUTER sockets when entering the mute state (when the send buffer is full) is to silently discard messages without blocking. This behavior, however, is not consistently observed. Quietly discarding messages is not the desired behavior anyway because it prevents peers from taking appropriate action to the error condition.

- Routers SHALL set the ZMQ_SNDBUFSIZE socket option to 0.
- Routers SHALL forward EAGAIN errors to sending peers.

It is also the default behavior of ROUTER sockets to silently drop messages addressed to unknown peers.

- Routers SHALL set the ZMQ_ROUTER_MANDATORY socket option.
- Routers SHALL forward EHOSTUNREACH errors to sending peers, unless the recipient address matches the sender.

Most subsystems are optional and some way of communicating unsupported subsystems to peers is needed.

- The error code 93, EPROTONOSUPPORT, SHALL be returned to peers to indicate unsupported or unimplemented subsystems.

The errors above are reported via the *error* subsystem. Other errors MAY be reported via the *error* subsystem, but subsystems SHOULD provide mechanisms for reporting subsystem-specific errors whenever possible.

An error message must contain the following:

- The recipient frame SHALL contain the socket identity of the original sender of the message.
- The sender frame SHALL contain the socket identity of the reporting entity, usually the router.
- The request ID SHALL be copied from the from the message which triggered the error.
- The subsystem frame SHALL be the 5 octets 'error'.
- The first data frame SHALL be a string representation of the error number.
- The second data frame SHALL contain a UTF-8 string describing the error.
- The third data frame SHALL contain the identity of the original recipient, as it may differ from the reporter.
- The fourth data frame SHALL contain the subsystem copied from the subsystem field of the offending message.

Subsystems

Peers may support any number of communications protocols or subsystems. For instance, there may be a remote procedure call (RPC) subsystem which defines its own protocol. These subsystems are outside the scope of VIP and this document with the exception of the *hello* and *ping* subsystems.

- A router SHALL implement the hello subsystem.
- All peers and routers SHALL implement the ping subsystem.

The hello Subsystem

The hello subsystem provides one simple RPC-style routine for peers to probe the router for version and identity information.

A peer hello request message must contain the following:

- The recipient frame SHALL have a zero length value.

- The request id MAY have an opaque binary value.
- The subsystem SHALL be the 5 characters “hello”.
- The first data frame SHALL be the five octets ‘hello’ indicating the operation.

A peer hello reply message must contain the following:

- The sender frame SHALL have a zero length value.
- The request id SHALL be copied unchanged from the associated request.
- The subsystem SHALL be the 7 characters “hello”.
- The first data frame SHALL be the 7 octets ‘welcome’.
- The second data frame SHALL be a string containing the router version number.
- The third data frame SHALL be the router’s identity blob.
- The fourth data frame SHALL be the peer’s identity blob.

The hello subsystem can help a peer with the following tasks:

- Test that a connection is established.
- Discover the version of the router.
- Discover the identity of the router.
- Discover the identity of the peer.
- Discover authentication metadata.

For instance, if a peer will use a ROUTER socket for its connections, it must first know the identity of the router. The peer might first connect with a DEALER socket, issue a hello, and use the returned identity to then connect the ROUTER socket.

The ping Subsystem

The *ping* subsystem is useful for testing the presence of a peer and the integrity and latency of the connection. All endpoints, including the router, must support the ping subsystem.

A peer ping request message must contain the following:

- The recipient frame SHALL contain the identity of the endpoint to query.
- The request id MAY have an opaque binary value.
- The subsystem SHALL be the 4 characters “ping”.
- The first data frame SHALL be the 4 octets ‘ping’.
- There MAY be zero or more additional data frames containing opaque binary blobs.

A ping response message must contain the following:

- The sender frame SHALL contain the identity of the queried endpoint.
- The request id SHALL be copied unchanged from the associated request.
- The subsystem SHALL be the 4 characters “ping”.
- The first data frame SHALL be the 4 octets ‘pong’.
- The remaining data frames SHALL be copied from the ping request unchanged, starting with the second data frame.

Any data can be included in the ping and should be returned unchanged in the pong, but limited trust should be placed in that data as it is possible a peer might modify it against the direction of this specification.

Discovery

VIP does not define how to discover peers or routers. Typical options might be to hard code the router address in peers or to pass it in via the peer configuration. A well known (i.e. statically named) directory service might be used to register connected peers and allow for discovery by other peers.

Example Exchanges

These examples show the messages *as sent on the wire* as sent or received by peers using DEALER sockets. The messages received or sent by peers or routers using ROUTER sockets will have an additional address at the start. We do not show the frame sizes or flags, only frame contents.

Example of hello Request

This shows a hello request sent by a peer, with identity “alice”, to a connected router, with identity “router”.

```

+--+
| |           Empty recipient frame
+-----+
| VIP1 |     Signature frame
+-----+
| |           Empty user ID frame
+-----+
| 0001 |     Request ID, for example "0001"
+-----+
| hello |     Subsystem, "hello" in this case
+-----+
| hello |     Operation, "hello" in this case
+-----+

```

This example assumes a DEALER socket. If a peer uses a ROUTER socket, it SHALL prepend an additional frame containing the router identity, similar to the following example.

This shows the example request received by the router:

```

+-----+
| alice |     Sender frame, "alice" in this case
+-----+
| |           Empty recipient frame
+-----+
| VIP1 |     Signature frame
+-----+
| |           Empty user ID frame
+-----+
| 0001 |     Request ID, for example "0001"
+-----+
| hello |     Subsystem, "hello" in this case
+-----+
| hello |     Operation, "hello" in this case
+-----+

```

This shows an example reply sent by the router:

```

+-----+
| alice |      Recipient frame, "alice" in this case
+-----+
| |          Empty sender frame
+-----+
| VIP1 |      Signature frame
+-----+
| |          Empty authentication metadata in user ID frame
+-----+
| 0001 |      Request ID, for example "0001"
+-----+
| hello |      Subsystem, "hello" in this case
+-----+
| welcome |    Operation, "welcome" in this case
+-----+
| 1.0 |       Version of the router
+-----+
| router |     Router ID, "router" in this case
+-----+
| alice |     Peer ID, "alice" in this case
+-----+

```

This shows an example reply received by the peer:

```

++
| |          Empty sender frame
+-----+
| VIP1 |      Signature frame
+-----+
| |          Empty authentication metadata in user ID frame
+-----+
| 0001 |      Request ID, for example "0001"
+-----+
| hello |      Subsystem, "hello" in this case
+-----+
| welcome |    Operation, "welcome" in this case
+-----+
| 1.0 |       Version of the router
+-----+
| router |     Router ID, "router" in this case
+-----+
| alice |     Peer ID, "alice" in this case
+-----+

```

Example of ping Subsystem

This shows a ping request sent by the peer “alice” to the peer “bob” through the router “router”.

```

+-----+
| bob |       Recipient frame, "bob" in this case
+-----+
| VIP1 |      Signature frame
+-----+
| |          Empty user ID frame

```

(continues on next page)

(continued from previous page)

```

+-----+
| 0002 |           Request ID, for example "0002"
+-----+
| ping |           Subsystem, "ping" in this case
+-----+
| ping |           Operation, "ping" in this case
+-----+
| 1422573492 |    Data, a single frame in this case (Unix timestamp)
+-----+

```

This shows the example request received by the router:

```

+-----+
| alice |         Sender frame, "alice" in this case
+-----+
| bob |          Recipient frame, "bob" in this case
+-----+
| VIP1 |         Signature frame
+-----+
| |             Empty user ID frame
+-----+
| 0002 |         Request ID, for example "0002"
+-----+
| ping |         Subsystem, "ping" in this case
+-----+
| ping |         Operation, "ping" in this case
+-----+
| 1422573492 |   Data, a single frame in this case (Unix timestamp)
+-----+

```

This shows the example request forwarded by the router:

```

+-----+
| bob |          Recipient frame, "bob" in this case
+-----+
| alice |        Sender frame, "alice" in this case
+-----+
| VIP1 |         Signature frame
+-----+
| |             Empty authentication metadata in user ID frame
+-----+
| 0002 |         Request ID, for example "0002"
+-----+
| ping |         Subsystem, "ping" in this case
+-----+
| ping |         Operation, "ping" in this case
+-----+
| 1422573492 |   Data, a single frame in this case (Unix timestamp)
+-----+

```

This shows the example request received by "bob":

```

+-----+
| alice |        Sender frame, "alice" in this case
+-----+
| VIP1 |         Signature frame
+-----+

```

(continues on next page)

(continued from previous page)

```

| |                               Empty authentication metadata in user ID frame
+-----+
| 0002 |                           Request ID, for example "0002"
+-----+
| ping |                           Subsystem, "ping" in this case
+-----+
| ping |                           Operation, "ping" in this case
+-----+
| 1422573492 |                       Data, a single frame in this case (Unix timestamp)
+-----+

```

If “bob” were using a ROUTER socket, there would be an additional frame prepended to the message containing the router identity, “router” in this case.

This shows an example reply from “bob” to “alice”

```

+-----+
| alice |                           Recipient frame, "alice" in this case
+-----+
| VIP1 |                           Signature frame
+-----+
| |                               Empty user ID frame
+-----+
| 0002 |                           Request ID, for example "0002"
+-----+
| ping |                           Subsystem, "ping" in this case
+-----+
| pong |                           Operation, "pong" in this case
+-----+
| 1422573492 |                       Data, a single frame in this case (Unix timestamp)
+-----+

```

The message would make its way back through the router in a similar fashion to the request.

Reference Implementation

Reference VIP router: <https://github.com/VOLTTRON/volttron/blob/master/volttron/platform/vip/router.py>

Reference VIP peer: <https://github.com/VOLTTRON/volttron/blob/master/volttron/platform/vip/socket.py>

Remote Procedure Calls

Remote procedure calls (RPC) is a new feature added with VOLTRON 3.0. The new VOLTRON Interconnect Protocol **VIP** introduced the ability to create new point-to-point protocols, called subsystems, enabling the implementation of **JSON-RPC 2.0**. This provides a simple method for agent authors to write methods and expose or export them to other agents, making request-reply or notify communications patterns as simple as writing and calling methods.

Exporting Methods

The *export()* method, defined on the RPC subsystem class, is used to mark a method as remotely accessible. This *export()* method has a dual use. The class method can be used as a decorator to statically mark methods when the agent class is defined. The instance method dynamically exports methods, and can be used with methods not defined

on the agent class. Each take an optional export name argument, which defaults to the method name. Here are the two export method signatures:

Instance method:

```
RPC.export(method, name=None)
```

Class method:

```
RPC.export(name=None)
```

And here is an example agent definition using both methods:

```
from volttron.platform.vip import Agent, Core, RPC

def add(a, b):
    '''Add two numbers and return the result'''
    return a + b

class ExampleAgent(Agent):
    @RPC.export
    def say_hello(self, name):
        '''Build and return a hello string'''
        return 'Hello, %s!' % (name,)

    @RPC.export('say_bye')
    def bye(self, name):
        '''Build and return a goodbye string'''
        return 'Goodbye, %s.' % (name,)

    @Core.receiver('setup')
    def onsetup(self, sender, **kwargs):
        self.vip.rpc.export('add')
```

Calling exported methods

The RPC subsystem provides three methods for calling exported RPC methods.

```
RPC.call(peer, method, *args, **kwargs)
```

Call the remote *method* exported by *peer* with the given arguments. Returns a gevent *AsyncResult* object.

```
RPC.batch(peer, requests)
```

Batch call remote methods exported by *peer*. *requests* must be an iterable of 4-tuples (*notify*, *method*, *args*, *kwargs*), where *notify* is a boolean indicating whether this is a notification or standard call, *method* is the method name, *args* is a list and *kwargs* is a dictionary. Returns a list of *AsyncResult* objects for any standard calls. Returns *None* if all requests were notifications.

```
RPC.notify(peer, method, *args, **kwargs)
```

Send a one-way notification message to *peer* by calling *method* without without returning a result.

Here are some examples:

```
self.vip.rpc.call(peer, 'say_hello', 'Bob').get()
results = self.vip.rpc.batch(peer, [(False, 'say_bye', 'Alice', {}), (True, 'later',
→ [], {})])
self.vip.rpc.notify(peer, 'ready')
```

Inspection

A list of methods is available by calling the *inspect* method. Additional information can be returned for any method by appending `‘.inspect’` to the method name. Here are a couple examples:

```
self.vip.rpc.call(peer, 'inspect')    # Returns a list of exported methods
self.vip.rpc.call(peer, 'say_hello.inspect') # Return metadata on say_hello method
```

Implementation

See the `rpc` module for implementation details.

Also see [RPC by example](#) for additional examples.

Messaging and Topics

Introduction

Agents in VOLTTRON™ communicate with each other using a publish/subscribe mechanism built on the Zero MQ Python library. This allows for great flexibility as topics can be created dynamically and the messages sent can be any format as long as the sender and receiver understand it. An agent with data to share publishes to a topic, then any agents interested in that data subscribe to that topic.

While this flexibility is powerful, it also could also lead to confusion if some standard is not followed. The current conventions for communicating in the VOLTTRON are:

- Topics and subtopics follow the format: `topic/subtopic/subtopic`
- Subscribers can subscribe to any and all levels. Subscriptions to “topic” will include messages for the base topic and all subtopics. Subscriptions to “topic/subtopic1” will only receive messages for that subtopic and any children subtopics. Subscriptions to empty string (“”) will receive ALL messages. This is not recommended.
 - All agents should subscribe to the “platform” topic. This is the topic the VOLTTRON will use to send messages to agents, such as “shutdown”.

Agents should set the “From” header. This will allow agents to filter on the “To” message sent back. This is especially useful for requests to the ArchiverAgent so agents do not receive replies not meant for their request.

Topics

In VOLTTRON

- `platform` - Base topic used by the platform to inform agents of platform events
- `platform/shutdown` - General shutdown command. All agents should exit upon receiving this. Message content will be a reason for the shutdown

- `platform/shutdown_agent` - This topic will provide a specific agent id. Agents should subscribe to this topic and exit if the id in the message matches their id.
- `devices` - Base topic for data being published by drivers
- `datalogger` - Base topic for agents wishing to record time series data
- `record` - Base topic for agents to record data in an arbitrary format.

Controller Agent Topics

See the documentation for the `[[ActuatorAgent]]`.

3.5.6 VOLTRON Security

There are various security-related topics throughout VOLTRON's documentation. This is a quick roadmap for finding security documentaion.

A core component of VOLTRON is its *message bus*. The security of this message bus is crucial to the entire system. The *VOLTRON Interconnect Protocol* provides communication over the message bus. VIP was built with security in mind from the ground up. VIP uses encrypted channels and enforces agent *authentication* by default for all network communication. VIP's *authorization* mechanism allows system administrators to limit agent capabilities with fine granularity.

Even with these security mechanisms built into VOLTRON, it is important for system administrators to *harden VOLTRON's underlying OS*.

Additional documentation related to VIP authentication and authorization is avaiable here:

Key Stores

Note: most VOLTRON users should not need to directly interact with agent key stores. These are notes for VOLTRON platform developers. This is not a stable interface and the implementation details are subject to change.

Each agent has its own encryption key-pair that is used to *authenticate* itself to the VOLTRON platform. A key-pair comprises a public key and a private (secret) key. These keys are saved in a key store, which is implemented by the `KeyStore` class. Each agent has its own key store.

Key Store Locations

There are two main locations key stores will be saved. Installed agents' key stores are in the the agent's data directory:

```
$VOLTRON_HOME/agents/<AGENT_UUID>/<AGENT_NAME>/keystore.json
```

Agents that are not installed, such as platform services and stand-alone agents, store their key stores here:

```
$VOLTRON_HOME/keystores/<VIP_IDENTITY>/keystore.json
```

Generating a Key Store

Agents automatically retrieve keys from their key store unless both the `publickey` and `secretkey` parameters are specified when the agent is initialized. If an agent's key store does not exist it will automatically be generated upon access.

Users can generate a key pair by running the `volttron-ctl auth keypair` command.

Known Hosts File

Before an agent can connect to a VOLTTRON platform that agent must know the platform's VIP address and public key (known as the *server key*). It can be tedious to manually keep track of server keys and match them with their corresponding addresses.

The purpose of the known-hosts file is to save a mapping of platform addresses to server keys. This way the user only has to specify a server key one time.

Saving a Server Key

Suppose a user wants to connect to a platform at `192.168.0.42:22916`, and the platform's public key is `uhjbcUm3kT5QWj5Py9w0XZ7c1p6EP8pdo4Hq4dNEIiQ`. To save this address-to-server-key association, the user can run:

```
volttron-ctl auth add-known-host --host 192.168.0.42:22916 --serverkey_
↪uhjbcUm3kT5QWj5Py9w0XZ7c1p6EP8pdo4Hq4dNEIiQ
```

Now agents on this system will automatically use the correct server key when connecting to the platform at `192.168.0.42:22916`.

Server Key for Local Platforms

When a platform starts it automatically adds its public key to the known-hosts file. Thus agents connecting to the local VOLTTRON platform (on the same system and using the same `$VOLTTRON_HOME`) will automatically be able to retrieve the platform's public key.

Know-Host-File Details

*Note: the following details regarding the known-hosts file are subject to change. These notes are primarily for developers, but they may be helpful if troubleshooting an issue. **The known-hosts file should not be edited directly.***

File Location

The known-hosts-file is stored at `$VOLTTRON_HOME/known_hosts`.

File Contents

Here are the contents of an example known-hosts file:

```
{
  "@": "FSG7LHhy3v8tdNz3gK35G6-oxUcyln54pYRKu5fBJzU",
  "127.0.0.1:22916": "FSG7LHhy3v8tdNz3gK35G6-oxUcyln54pYRKu5fBJzU",
  "127.0.0.2:22916": "FSG7LHhy3v8tdNz3gK35G6-oxUcyln54pYRKu5fBJzU",
  "127.0.0.1:12345": "FSG7LHhy3v8tdNz3gK35G6-oxUcyln54pYRKu5fBJzU",
  "192.168.0.42:22916": "uhjbcUm3kT5QWj5Py9w0XZ7c1p6EP8pdo4Hq4dNEIiQ"
}
```

The first four entries are for the local platform. (They were automatically added when the platform started.) The first entry with the @ key is for IPC connections, and the entries with the 127.0.0.* keys are for local TCP connections. Note that a single VOLTRON platform can bind to multiple TCP addresses, and each address will be automatically added to the known-hosts file. The last entry is for a remote VOLTRON platform. (It was added in the *Saving a Server Key* section.)

Protecting Pub/Sub Topics

VIP *authorization* enables VOLTRON platform owners to protect pub/sub topics. More specifically, a platform owner can limit who can publish to a given topic. This protects subscribers on that platform from receiving messages (on the protected topic) from unauthorized agents.

Example

To protect a topic, add the topic name to `$VOLTRON_HOME/protected_topics.json`. For example, the following protected-topics file declares that the topic `foo` is protected:

```
{
  "write-protect": [
    {"topic": "foo", "capabilities": ["can_publish_to_foo"]}
  ]
}
```

Note: The capability name `can_publish_to_foo` is not special. It can be any string, but it is easier to manage capabilities with meaningful names.

Now only agents with the capability `can_publish_to_foo` can publish to the topic `foo`. To add this capability to authenticated agents, run `volttron-ctl auth update` (or `volttron-ctl auth add` for new authentication entries), and enter `can_publish_to_foo` in the capabilities field:

```
capabilities (delimit multiple entries with comma) []: can_publish_to_foo
```

Agents that have the `can_publish_to_foo` capabilities can publish to topic `foo`. That is, such agents can call:

```
self.vip.pubsub.publish('pubsub', 'foo', message='Here is a message')
```

If unauthorized agents try to publish to topic `foo` they will get an exception:

```
to publish to topic "foo" requires capabilities ['can_publish_to_foo'], but
capability list [] was provided
```

Regular Expressions

Topic names in `$VOLTRON_HOME/protected_topics.json` can be specified as regular expressions. In order to use a regular expression, the topic name must begin and end with a `/`. For example:

```
{
  "write-protect": [
    {"topic": "/foo/*.*/", "capabilities": ["can_publish_to_foo"]}
  ]
}
```

This protects topics such as `foo/bar` and `foo/anything`.

Security Update Notes

This is a list of updates to security-related functionality in VOLTTRON that either break backward compatibility or may have noticeable impact to the user.

Version 3.5rc1

- `$VOLTRON_HOME/auth.json` should not be edited with a text editor. Use `volttron-ctl` commands `auth-list`, `auth-add`, `auth-remove`, and `auth-update` to view and manipulate that file.
- `#-style` comments are no longer supported in `$VOLTRON_HOME/auth.json`. Use the comments and enabled fields. (See the *agent authentication walkthrough*.)

Version 4.0

- The `$VOLTRON_HOME/curve.key` file has been replaced with a *key store*. Use the `scripts/update_curve_key.py` script to update an existing key pair.
- A `mechanism` field has been added to the auth file. Therefore, the `credentials` field no longer is prepended with a mechanism such as “CURVE:”. VOLTTRON automatically updates the auth entries to use the new field.
 - Entries with a regular expression in the `credentials` field cannot be upgraded.
- Security-related commands for `volttron-ctl` have been moved to a `auth` subcommand. (See the *auth command documentation*.)

3.5.7 VOLTTRON PNNL Licensed Code

Agent Mobility

The mobility module can enable the deployment of agents within a site or to other sites (possibly in different building, cities, etc.) remotely. The mobility feature allows authorized VOLTTRON platforms to send and deploy agents, allowing for greater management and deployment ease and flexibility.

This feature requires that you have installed the [VOLTTRON™ Restricted](#).

To create the required keys (minimum requirement to run VOLTTRON with Restricted module installed) enter the following commands in a command terminal:

1. Create ssh directory in `VOLTRON_HOME` (see [PlatformConfiguration](#) for details on configuring the platform):

```
mkdir -p ~/.volttron/ssh
```

2. Generate ssh key and add to `id_rsa` file:

```
ssh-keygen -t rsa -N '' -f ~/.volttron/ssh/id_rsa
```

3. Create empty file for authorized keys and know hosts:

```
touch ~/.volttron/ssh/{authorized_keys,known_hosts}
```

Then, for each host you wish to authorize, its public key must be added to the `authorized_keys` file on the host to which it needs to connect. The public key has a `.pub` extension. The added hosts must have VOLTTRON instances installed, with the Restricted code installed and enabled (the authorized host has created the keys detailed in the steps above):

- Copy host information securely:

```
scp otherhost.example.com:~/.volttron/ssh/id_rsa.pub ./otherhost.pub
```

- Append host key(s) to authorized_keys file in \$VOLTRON_HOME/ssh:

```
“cat otherhost.pub >> ~/.volttron/ssh/authorized_keys“““
```

See the [PingPongAgent example](#) for information on how to add this feature to your custom agents.

Agent Package Signing

The creation of a signed agent package requires four certificates. The developer (creator) certificate is used to sign the agent code and allows the platform to verify that the agent code has not been modified since being distributed. The admin certificate is used for allowing the agent into a scope of influence. The initiator certificate is used when the agent is ready to be deployed into a specific platform. The platform certificate is used to sign the possibly modified data that an agent would like to carry with it during moving from platform to platform. All of these certificates must be signed by a “known” Certificate Authority (CA).

In order to facilitate the development of agents Volttron Restricted includes packaging commands for creating the platform CA as well as the CA signed certificates for use in the agent signing process.

When the Volttron Restricted package is installed on a platform the volttron-pkg command will be expanded to

```
usage: volttron-pkg [-h] [-I FILE] [-L FILE] [-q] [-v]
                [-verbooseness LEVEL] | {package,repacage,configure,create_ca,create_cert,sign,verify}
```

The additional (sub)commands:

- create_ca - Creates a platform specific root CA. When this command is executed the user will be required to respond to prompts in order to fill out the certificate’s data.
- create_cert - Allows the creation of a ca signed certificate. A type of certificate must be specified as (–creator | –admin | –initiator | –platform) and the name(–name) of the certificate may be specified. The name will be used as the filename for the certificate on the platform.
- sign - Signs the agent package at the specified level.
- (ALWAYS REQUIRED) Agent package to be signed.
- (ALWAYS REQUIRED) Signing level must be specified as one of (–creator | –admin | –initiator | –platform) and must be presented in the correct order. In other words an admin cannot sign the package until the creator has signed it.
- –contract (resource contract) a file containing the definition of the necessary agent resources needed to execute properly. This option is only available to the creator.
- –config-file a file used to define custom configuration for the starting of agent on the platform. This option is available to the initiator.
- –certs_dir Allows the specification of where the certificate store is located. If this is not specified the default certificate store will be used.
- verify Allows the user to verify a package is valid.
- package - The agent package to validate against.

Agent Signing Example Using the Listener Agent

If VOLTRON Restricted is installed and the security features are enabled, all agents must be signed prior to launching them. The following steps will describe how to sign an agent and will use the Listener agent as an example. From a terminal, in the `voltrtron` directory (`~/voltrtron`), enter the following commands:

1. Package the agent:

```
voltrtron-pkg package examples/ListenerAgent
```

2. Sign the agent as creator:

```
voltrtron-pkg sign --creator --contract resource_contract ~/.voltrtron/  
packaged/listeneragent-3.0-py2-none-any.whl
```

3. Sign the agent as admin:

```
voltrtron-pkg sign --admin ~/.voltrtron/packaged/listeneragent-3.  
0-py2-none-any.whl
```

4. Sign the agent as initiator:

```
voltrtron-pkg sign --initiator --config-file examples/ListenerAgent/config  
~/.voltrtron/packaged/listeneragent-3.0-py2-none-any.whl
```

5. Set the configuration file:

```
voltrtron-pkg configure ~/.voltrtron/packaged/listeneragent-3.  
0-py2-none-any.whl examples/ListenerAgent /config
```

6. Install agent into platform (with the platform running):

```
voltrtron-ctl install ~/.voltrtron/packaged/listeneragent-3.0-py2-none-any.  
whl
```

Upon successful completion of this command the terminal output will inform one of the install directory, the agent UUID (unique identifier for an agent, each instance of an agent will have a different UUID) and the agent name:

```
Installed ~/.voltrtron/packaged/weatheragent-3.0-py2-none-any.whl as  
a9d67c55-7f58-4591-80af-3c1ff8a81740 listeneragent-3.0
```

Now the agent can be started, enter the following command to start the agent:

```
voltrtron-ctl start --name listeneragent-3.0
```

or

```
voltrtron-ctl start --uuid <UUID>
```

The previous steps document the signing of an agent using the Listener agent as an example. When deploying other agents the following parameters, used in the example, will need to be modified according to which agent you are signing and starting. The following is a brief description of the components of the commands in the signing example above:

- `examples/ListenerAgent` - Path to agents top-level directory.
- `~/.voltrtron/packaged/listeneragent-3.0-py2-none-any.whl` - Path to the agent wheel (check the `~/.voltrtron/packaged` directory after step 1 in this example to view the name of the wheel).
- `examples/ListenerAgent/config` - Path to the agents json style configuration file (note this configuration file is named `config`).

Resource Monitor

The *VOLTRON™ Restricted* additions provide additional protection against an agent consuming too many resources to the point of the host system becoming unresponsive or unstable. The resource monitor uses Linux control groups (or cgroups) to limit the CPU cycles and memory an individual agent may consume, preventing its possible overconsumption from adversely affecting other agents and services on the system. The execution requirements of an agent are set when *provisioning an agent* for service.

When a request is made to move an agent to a new platform, part of the validation of the agent includes checking its execution requirements against resources currently available on the system. If the resources are available and the agent has passed all other validation, the agent will be executed and retain those resource guarantees throughout its lifetime on that platform. If the agent, however, requests memory or CPU cycles that are not available, its move request is denied and it will not execute on the requested platform.

Once an agent has been assigned resources, it is the responsibility of that agent to manage use of its resources. While an agent may exceed its resource guarantees when system utilization is low, when resources given to other agents are required, an agent exceeding the use in its contract may be terminated.

Execution Requirements

The execution requirements are specified as a JSON formatted document embedded in the agent during initial provisioning and takes the following form:

```
{
  "requirements": {
    "cpu.bogomips": 100,
    "memory.soft_limit_in_bytes": 2000000
  }
}
```

The contract *must* contain the `requirements` object, specifying the soft requirements, and might optionally specify a `hard_requirements` object.

Each agent, including newly developed agents, must maintain their own requirements. The execution requirements for an agent are located in a file in the individual agent directory, called *exereqs.json*.

For example, the execution requirements for the Listener Agent are located at *volttron/examples/ListenerAgent/exereqs.json*

Soft requirements

Soft requirements are considered *soft* on the platform because they change depending on the number of agents and other services are running on the system. They may also be negotiated on the fly in a future release. A list of the current resources which may be reserved follows:

- **cpu.bogomips** - The CPU requirements of an agent indicated as either an exact integer ($N \geq 1$) in MIPS (millions of instructions per second) or a floating-point percentage ($0.0 < N < 1.0$) of the total available bogomips on a system. Bogomips is a rough calculation performed at system boot indicating the likely number of calculations a system may perform each second.
- **memory.soft_limit_in_bytes** - The maximum amount of random access memory (RAM) an agent requires to perform its tasks measured in bytes and given as an integer. Additional resources may be added in a future release.

Hard requirements

Hard requirements are based on system attributes that are very unlikely to change except after a system reboot. It is rare that an agent would need to set hard requirements and is usually only necessary for architecture-specific code. Each hard requirement is tested for a match.

- **kernel.name** - Kernel name as given by `uname`.
- **kernel.release** - Kernel release as given by `uname`.
- **kernel.version** - Kernel version as given by `uname`.
- **architecture** - Kernel architecture as given by `uname`.
- **os** - Always 'GNU/Linux'
- **platform.version** - Version of VOLTTRON in use.
- **memory.total** - Total amount of memory on the system in bytes.
- **bogomips.total** - Total of all bogomips reported for all processors on the system.

Example using requirements

In this example we modify the execution requirements for the Listener Agent.

1. Open a terminal and type the following command:

```
cat /proc/meminfo | grep MemTotal
```

The output will be the total memory available on the system. Save this number.

2. In a text editor, open `voltron/examples/ListenerAgent/exereqs.json`
3. Replace the requirements with the following text:

```
{
  "requirements": {
    "cpu.bogomips": 100,
    "memory.soft_limit_in_bytes": 2000000
  },
  "hard_requirements": {
    "os": "GNU/Linux",
    "memory.total": 2064328
  }
}
```

4. Replace the number for "memory.total" with the number from step 1, so that the requirement matches the memory for your system.
5. Save and close the file. Now, if the total memory on the system is changed, such as with a hardware update, the requirement will fail. Note that the hard requirements are separate, and follow the same format as the soft requirements.

3.5.8 VOLTTRON Configuration Store

The configuration store provides storage for agent configurations and an agent interface to facilitate dynamic agent configuration.

Configuration Store Command Line Tools

Command line management of the Configuration Store is done with the *voltron-ctl config* sub-commands.

Store Configuration

To store a configuration in the Configuration Store use the *store* sub-command:

```
voltron-ctl config store <agent vip identity> <configuration name> <infile>
```

- **agent vip identity** - The agent store to add the configuration to.
- **configuration name** - The name to give the configuration in the store.
- **infile** - The file to ingest into the store.

Optionally you may specify the file type of the file. Defaults to `--json`.

- `--json` - Interpret the file as JSON.
- `--csv` - Interpret the file as CSV.
- `--raw` - Interpret the file as raw data.

Delete Configuration

To delete a configuration in the Configuration Store use the *delete* sub-command:

```
voltron-ctl config delete <agent vip identity> <configuration name>
```

- **agent vip identity** - The agent store to delete the configuration from.
- **configuration name** - The name of the configuration to delete.

To delete all configurations for an agent in the Configuration Store use `--all` switch in place of the configuration name:

```
voltron-ctl config delete <agent vip identity> --all
```

Get Configuration

To get the current contents of a configuration in the Configuration Store use the *get* sub-command:

```
voltron-ctl config get <agent vip identity> <configuration name>
```

- **agent vip identity** - The agent store to retrieve the configuration from.
- **configuration name** - The name of the configuration to get.

By default this command will return the json representation of what is stored.

- `--raw` - Return the raw version of the file.

List Configurations

To get the current list of agents with configurations in the Configuration Store use the *list* sub-command:

```
voltttron-ctl config list
```

To get the current list of configurations for an agent include the Agent's VIP IDENTITY:

```
voltttron-ctl config list <agent vip identity>
```

- **agent vip identity** - The agent store to retrieve the configuration from.

Edit Configuration

To edit a configuration in the Configuration Store use the *edit* sub-command:

```
voltttron-ctl config edit <agent vip identity> <configuration name>
```

- **agent vip identity** - The agent store containing the configuration.
- **configuration name** - The name of the configuration to edit.

The configuration must exist in the store to be edited.

By default *edit* will try to open the file with the *nano* editor. The *edit* command will respect the *EDITOR* environment variable. You may override this with the *-editor* option.

Platform Configuration Store

The Platform Configuration Store is a mechanism provided by the platform to facilitate the dynamic configuration of agents. The Platform Configuration Store works by informing agents of changes to their configuration store and the agent responding to those changes by updating any settings, subscriptions, or processes that are affected by the configuration of the Agent.

Support for the Configuration Store is not automatically available in existing agents. An agent must be updated in order to support this feature. Currently only the Master Driver Agent, the Aggregation Agent, and the Actuator Agent support the Configuration Store.

Configurations and Agents

Each agent has its own configuration store (or just store). Agents are not given access to any other agent's store.

The existence of a store is not dependent on the existence of an agent installed on the platform.

Each store has a unique identity. Stores are matched to agents at agent runtime via the agent's VIP IDENTITY. Therefore the store for an agent is the store with the same identity as the agent's VIP IDENTITY.

When a user updates a configuration in the store the platform immediately informs the agent of the change. The platform will not send another update until the Agent finishes processing the first. The platform will send updates to the agent, one file at a time, in the order the changes were received.

Configuration Names

Every configuration in an agent's store has a unique name. When a configuration is added to an agent's store with the same name as an existing configuration it will replace the existing configuration. The store will remove any leading or trailing whitespace, "/", and "" from the name.

Configuration File Types

The configuration store will automatically parse configuration before presenting them to an agent. While the configuration store does support storing raw data and giving to the agent unparsed. Most Agents will require the configuration to be parsed. Any Agent that requires raw data will specifically mention it in its documentation.

This system removes the requirement that configuration files for an agent be in a specific format. For instance a registry configuration for a driver may be JSON instead of CSV if that is more convenient for the user. This will work as long as the JSON parses into an equivalent set of objects as an appropriate CSV file.

Currently the store supports parsing JSON and CSV files with support for more files types to come.

JSON

The store uses the same JSON parser that agents use to parse their configuration files. Therefore it supports Python style comments and must create an object or list when parsed.

```
{
  "result": "PREEMPTED", #This is a comment.
  "info": null,
  "data": {
    "agentID": "my_agent", #This is another comment.
    "taskID": "my_task"
  }
}
```

CSV

A CSV file is represented as a list of objects. Each object represents a row in the CSV file.

For instance this simple CSV file:

Table 20: Example CSV

Volttron Point Name	Modbus Register	Writable	Point Address
ReturnAirCO2	>f	FALSE	1001
ReturnAirCO2Stpt	>f	TRUE	1011

Is the equivalent to this JSON file:

```
[
  {
    "Volttron Point Name": "ReturnAirCO2",
    "Modbus Register": ">f",
    "Writable": "FALSE",
    "Point Address": "1001"
  },
]
```

(continues on next page)

(continued from previous page)

```

{
  "Volttron Point Name": "ReturnAirCO2Stpt",
  "Modbus Register": ">f",
  "Writable": "TRUE",
  "Point Address": "1011"
}
]

```

File references

The Platform Configuration Store supports referencing one configuration file from another. If a referenced file exists the contents of that file will replace the file reference when the file is processed by the agent. Otherwise the reference will be replaced with null (or in Python, None).

Only configurations that are parsed by the platform (currently JSON or CSV) will be examined for references. If the file referenced is another parsed file type (JSON or CSV, currently) then the replacement will be the parsed contents of the file, otherwise it will be the raw contents of the file.

In a JSON object the name of a value will never be considered a reference.

A file reference is any value string that starts with “config://”. The rest of the string is the name of another configuration. The configuration name is converted to lower case for comparison purposes.

Consider the following configuration files named “devices/vav1.config” and “registries/vav.csv”, respectively:

```

{
  "driver_config": {"device_address": "10.1.1.5",
                    "device_id": 500},

  "driver_type": "bacnet",
  "registry_config": "config://registries/vav.csv",
  "campus": "pnnl",
  "building": "isb1",
  "unit": "vav1"
}

```

Table 21: vav.csv

Volttron Point Name	Modbus Register	Writable	Point Address
ReturnAirCO2	>f	FALSE	1001
ReturnAirCO2Stpt	>f	TRUE	1011

The resulting configuration returns when an agent asks for “devices/vav1.config”.

```

{
  "driver_config": {"device_address": "10.1.1.5",
                    "device_id": 500},

  "driver_type": "bacnet",
  "registry_config": [
    {
      "Volttron Point Name": "ReturnAirCO2",
      "Modbus Register": ">f",
      "Writable": "FALSE",
      "Point Address": "1001"
    }
  ]
}

```

(continues on next page)

(continued from previous page)

```

        },
        {
            "Volttron Point Name": "ReturnAirCO2Stpt",
            "Modbus Register": ">f",
            "Writable": "TRUE",
            "Point Address": "1011"
        }
    ],
    "campus": "pnnl",
    "building": "isb1",
    "unit": "vav1"
}

```

Circular references are not allowed. Adding a file that creates a circular reference will cause that file to be rejected by the platform.

If a configuration is changed in anyway and that configuration is referred to by another configuration then the agent considers the referring configuration as changed. Thus a set of configurations with references can be considered one large configuration broken into pieces for the users convenience.

Multiple configurations may all reference a single configuration. For instance, when configuring drivers in the Master Driver you may have multiple drivers reference the same registry if appropriate.

Modifying the Configuration Store

Currently the configuration store must be modified through the command line. See *Configuration Store Command Line Tools*.

3.5.9 MultiPlatform Message Bus Communication

The multi platform message bus communication allows the user to connect to remote VOLTRON platforms seamlessly. This bypasses the need for an agent wanting to send/receive messages to/from remote platforms from having to setup the connection to remote platform directly. Instead, the router module in each platform will maintain connections to the remote platforms internally, that means it will connect, disconnect and monitor the status of each connection.

Multi-Platform Communication

To connect to remote VOLTRON platforms, we would need platform discovery information of the remote platforms. This information contains the platform name, VIP address and serverkey of the remote platforms and we need to provide this as part of Multiplatform configuration.

Configuration

The configuration and authentication for multi-platform connection can be setup either manually or by running the platforms in set up mode. Both the setups are described below.

Setup Mode For Automatic Authentication

Note: It is necessary for each platform to have a web server if running in setup mode.

For ease of use and to support multi-scale deployment, the process of obtaining the platform discovery information and authenticating the new platform connection is automated. We can now bypass the manual process of adding auth keys (i.e., either by using the `voltron-ctl` utility or directly updating the `auth.json` config file).

A config file containing list of web addresses (one for each platform) need to be made available in `VOLTRON_HOME` directory.

Name of the file: `external_address.json`

Directory path: Each platform's `VOLTRON_HOME` directory.

For example: `/home/voltron/.voltron1`

Contents of the file:

```
[
"http://<ip1>:<port1>",
"http://<ip2>:<port2>",
"http://<ip3>:<port3>",
.....
]
```

We then start each VOLTRON platform with setup mode option in this way.

```
voltron -vv -l voltron.log --setup-mode&
```

Each platform will obtain the platform discovery information of the remote platform that it is trying to connect through a HTTP discovery request and store the information in a configuration file (`$VOLTRON_HOME/external_platform_discovery.json`). It will then use the VIP address and serverkey to connect to the remote platform. The remote platform shall authenticate the new connection and store the auth keys (public key) of the connecting platform for future use.

The platform discovery information will be stored in `VOLTRON_HOME` directory and looks like below:

Name of config file: `external_platform_discovery.json`

Contents of the file:

```
{"<platform1 name>": {"vip-address":"tcp://<ip1>:<vip port1>",
"instance-name":"<platform1 name>",
"serverkey":"<serverkey1>"
},
"<platform2 name>": {"vip-address":"tcp://<ip2>:<vip port2>",
"instance-name":"<platform2 name>",
"serverkey":"<serverkey2>"
},
"<platform3 name>": {"vip-address":"tcp://<ip3>:<vip port3>",
"instance-name":"<platform3 name>",
"serverkey":"<serverkey3>"
},
.....
}
```

Each platform will use this information for future connections.

Once the keys have been exchanged and stored in the auth module, we can restart all the VOLTRON platforms in normal mode.

```

volttron-ctl shutdown --platform

volttron -vv -l volttron.log&

```

Manual Configuration of External Platform Information

Platform discovery configuration file can also be built manually and it needs to be added inside VOLTRON_HOME directory of each platform.

Name of config file: external_platform_discovery.json

Contents of the file:

```

{ "<platform1 name>": { "vip-address": "tcp://<ip1>:<vip port1>",
                      "instance-name": "<platform1 name>",
                      "serverkey": "<serverkey1>"
                    },
  "<platform2 name>": { "vip-address": "tcp://<ip2>:<vip port2>",
                      "instance-name": "<platform2 name>",
                      "serverkey": "<serverkey2>"
                    },
  "<platform3 name>": { "vip-address": "tcp://<ip3>:<vip port3>",
                      "instance-name": "<platform3 name>",
                      "serverkey": "<serverkey3>"
                    },
  .....
}

```

With this configuration, platforms can be started in normal mode.

```

volttron -vv -l volttron.log&

```

For external platform connections to be authenticated, we would need to add the credentials of the connecting platforms in each platform using the volttron-ctl auth utility. For more details [Agent authentication walkthrough](#).

See also:

[Multi-Platform Walkthrough](#)

Multi-Platform PubSub Communication

Multi-Platform pubsub communication allows an agent on one platform to subscribe to receive messages from another platform without having to setup connection to the remote platform directly. The connection will be internally managed by the VOLTRON platform router module. Please refer here [Multi-Platform Communication Setup](#) for more details regarding setting up of Multi-Platform connections.

External Platform Message Subscription

To subscribe for topics from remote platform, the subscriber agent has to add an additional input parameter -all_platforms to the pubsub subscribe method.

Here is an example,

```
self.vip.pubsub.subscribe('pubsub', 'foo', self.on_match, all_platforms=True)
```

There is no change in the publish method of PubSub subsystem. If all the configurations are correct and the publisher agent on the remote platform is publishing message to topic="foo", then the subscriber agent will start receiving those messages.

Multi-Platform RPC Communication

Multi-Platform RPC communication allows an agent on one platform to make RPC call on an agent in another platform without having to setup connection to the remote platform directly. The connection will be internally managed by the VOLTTRON platform router module. Please refer here *Multi-Platform Communication Setup* for more details regarding setting up of Multi-Platform connections.

Calling External Platform RPC Method

If an agent in one platform wants to use an exported RPC method of an agent in another platform, it has to provide the platform name of the remote platform when using RPC subsystem call/notify method.

Here is an example:

```
self.vip.rpc.call(peer, 'say_hello', 'Bob', external_platform='platform2').get()
self.vip.rpc.notify(peer, 'ready', external_platform='platform2')
```

Here, 'platform2' is the platform name of the remote platform.

3.5.10 Message Bus Refactor

Refactoring of the existing message bus became necessary as we needed to reduce long term costs of maintenance, enhancement and support of the message bus. It made sense to move to a more widely used, industry accepted messaging library such as RabbitMQ that has many of the features that we need already built in.

1. It has many different messaging patterns and routing topologies.
2. It offers flexibility in deployment and supports large scale deployment
3. It has well-developed SSL based authentication plugin.

The goal of the message bus refactor task is to

1. Maintain essential features of current message bus and minimize transition cost
2. Leverage an existing and growing community dedicated to the further development of RabbitMQ
3. Move services provided currently by VOLTTRON agents to services natively provided by RabbitMQ
4. Decrease VOLTTRON development time spent on supporting message bus which is now a commodity technology
5. Address concerns from community about ZeroMQ.

Message Bus Plugin Framework

The message bus plugin framework aims to decouple the VOLTRON specific code from the message bus implementation without compromising the existing features of the platform. The concept of the plugin framework is similar to that used in historian or driver framework i.e, we should be easily able to support multiple message buses and be able to use any of them by following few installation and setup steps.

Message Bus Refactor

It consists of five components

1. New connection class per message bus
2. Extensions to platform router functionality
3. Extensions to core agent functionality
4. A proxy agent for each message bus to support backward compatibility
5. Authentication related changes

Connection class

A connection class that has methods to handle

1. Connection to new message bus.
2. Set properties such as message transmission rate, send/receive buffer sizes, open socket limits etc.
3. Send/receive messages from the underlying layer.
4. Error handling functionality.
5. Disconnect from the message bus

Platform Level Changes

A new message bus flag is introduced to indicate the type of message bus used by the platform. If no message bus flag is added in the platform config file, the platform uses default ZeroMQ based message bus.

Path of the config: \$VOLTRON_HOME/config

```
[volttron]
vip-address = tcp://127.0.0.1:22916
instance-name = volttron1
message-bus = rmq
```

Please note, the valid message types are ‘zmq’ and ‘rmq’.

On startup, platform checks for the type of message bus and creates appropriate router module. Please note, ZeroMQ router functionality remains unchanged. However, a new router module with limited functionality is added for RabbitMQ message bus. The actual routing of messages is handed over to the RabbitMQ broker and router module will only handle some of the necessary subsystem messages such as “hello”, “peerlist”, “query” etc. If a new message bus needs to be added then the complexity of the router module depends on whether the messaging library uses a broker based or broker less (as in case of ZeroMQ) protocol.

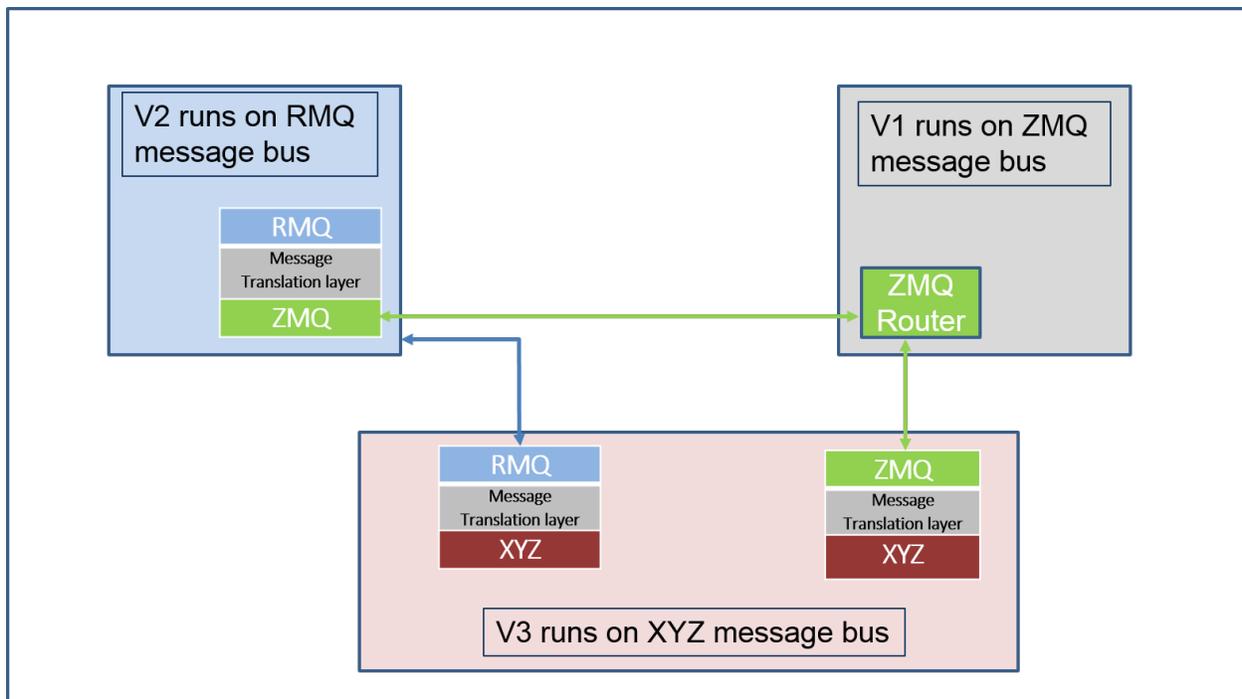
Agent Core Changes

The application specific code of the agent remains unchanged. The agent core functionality is modified to check the type of message bus and connect to and use the appropriate message bus. On startup, the agent Core checks the type of message bus, connects to appropriate message bus and routes messages to appropriate subsystem. All subsystem messages are encapsulated inside a message bus agnostic VIP message object. If a new message bus needs to be added, then we would have to extend the Agent Core to connect to new message bus.

Compatibility Between VOLTRON Instances Running On Different Message Buses

All the agents connected to local platform uses the same message bus that the platform is connected to. But if we need agents running on different platforms with different message buses to communicate with each other then we need some kind of proxy entity or bridge that establishes the connection, handles the message routing and performs the message translation between the different message formats. To achieve that, we have a proxy agent that acts as a bridge between the local message bus and remote message bus. The role of the proxy agent is to

- Maintain connections to internal and external message bus.
- Route messages from internal to external platform.
- Route messages from external to internal platform.



The above figure shows three VOLTRON instances with V1 connected to ZMQ message bus, V2 connected to RMQ message bus and V3 connected to XYZ (some message bus of the future) and all three want to connect to each other. Then V2 and V3 will have proxy agents that get connected to the local bus and to the remote bus and forward messages from one to another.

Authentication Related Changes

RabbitMQ Overview

NOTE: Some of the RabbitMQ summary/overview documentation and supporting images added here are taken from [RabbitMQ official documentation](#).

RabbitMQ is the most popular messaging library with over 35,000 production deployments. It is highly scalable, easy to deploy, runs on many operating systems and cloud environments. It supports many kinds of distributed deployment methodologies such as clusters, federation and shovels.

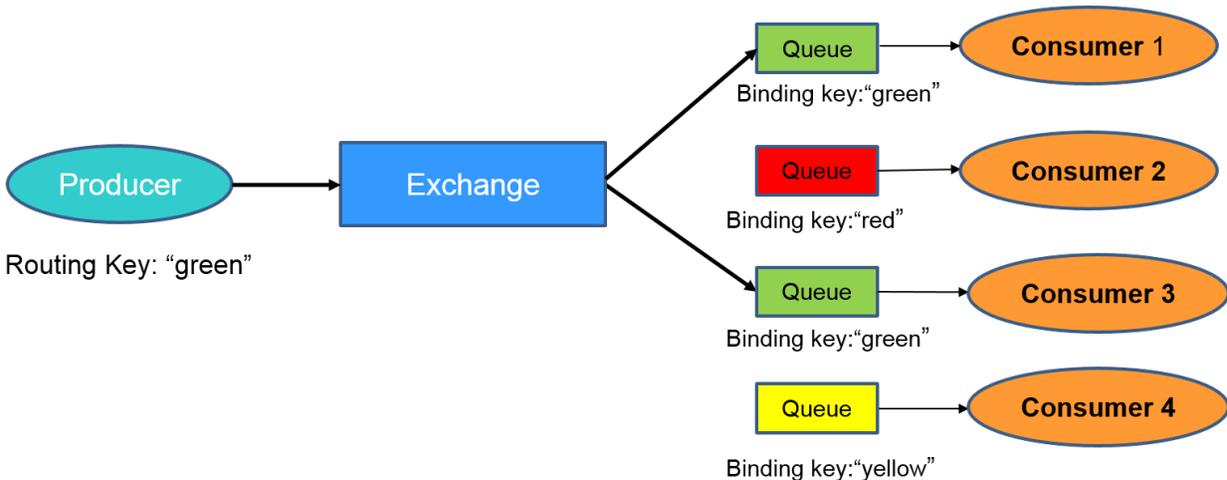
RabbitMQ uses Advanced Message Queueing Protocol (AMQP) and works on the basic producer consumer model. A consumer is a program that consumes/receives messages and producer is a program that sends the messages. Following are some important definitions that we need to know before we proceed.

- Queue - Queues can be considered like a post box that stores messages until consumed by the consumer. Each consumer must create a queue to receives messages that it is interested in receiving. We can set properties to the queue during it's declaration. The queue properties are
 - Name - Name of the queue
 - Durable - Flag to indicate if the queue should survive broker restart.
 - Exclusive - Used only for one connection and it will be removed when connection is closed.
 - Auto-queue - Flag to indicate if auto-delete is needed. The queue is deleted when last consumer un-subscribes from it.
 - Arguments - Optional, can be used to set message TTL (Time To Live), queue limit etc.
- Bindings - Consumers bind the queue to an exchange with binding keys or routing patterns. Producers send messages and associate them with a routing key. Messages are routed to one or many queues based on a pattern matching between a message routing key and binding key.
- Exchanges - Exchanges are entities that are responsible for routing messages to the queues based on the routing pattern/binding key used. They look at the routing key in the message when deciding how to route messages to queues. There are different types of exchanges and one must choose the type of exchange depending on the application design requirements
 1. Fanout - It blindly broadcasts the message it receives to all the queues it knows.
 2. Direct - Here, the message is routed to a queue if the routing key of the message exactly matches the binding key of the queue.
 3. Topic - Here, the message is routed to a queue based on pattern matching of the routing key with the binding key. The binding key and the routing key pattern must be a list of words delimited by dots, for example, "car.subaru.outback" or "car.subaru.*", "car.#". A message sent with a particular routing key will be delivered to all the queues that are bound with a matching binding key with some special rules as
 - ‘*’ (star) - can match exactly one word in that position. ‘#’ (hash) - can match zero or more words
 4. Headers - If we need more complex matching then we can add a header to the message with all the attributes set to the values that need to be matched. The message is considered matching if the values of the attributes in the header is equal to that of the binding. Header exchange ignore the routing key.

We can set some properties to the exchange during it's declaration.

- Name - Name of the exchange
- Durable - Flag to indicate if the exchange should survive broker restart.
- Auto-delete - Flag indicates if auto-delete is needed. If set to true, the exchange is deleted when the last queue is unbound from it.
- Arguments - Optional, used by plugins and broker-specific features

Lets use an example to understand how they all fit together. Consider an example where there are four consumers (Consumer 1 - 4) interested in receiving messages matching the pattern “green”, “red” or “yellow”. In this example, we are using a direct exchange that will route the messages to the queues only when there is an exact match of the routing key of the message with the binding key of the queues. Each of the consumers declare a queue and bind the queue to the exchange with a binding key of interest. Lastly, we have a producer that is continuously sending messages to exchange with routing key “green”. The exchange will check for an exact match and route the messages to only Consumer 1 and Consumer 3.



For more information about queues, bindings, exchanges, please refer to [RabbitMQ tutorial](#).

Distributed RabbitMQ Brokers

RabbitMQ allows multiple distributed RabbitMQ brokers to be connected in three different ways - with clustering, with federation and using shovel. We take advantage of these built-in plugins for multi-platform VOLTRON communication. For more information about the differences between clustering, federation, and shovel, please refer to RabbitMQ documentation [Distributed RabbitMQ brokers](#).

Clustering

Clustering connects multiple brokers residing in multiple machines to form a single logical broker. It is used in applications where tight coupling is necessary i.e, where each node shares the data and knows the state of all other nodes in the cluster. A new node can connect to the cluster through a peer discovery mechanism if configured to do so in the RabbitMQ config file. For all the nodes to be connected together in a cluster, it is necessary for them to share the same Erlang cookie and be reachable through it's DNS hostname. A client can connect to any one of the nodes in the cluster and perform any operation (to send/receive messages from other nodes etc.), the nodes will route the operation internally. In case of a node failure, clients should be able to reconnect to a different node, recover their topology and continue operation.

Please note, this feature is not integrated into VOLTRON. But we hope to support it in the future. For more detailed information about clustering, please refer to RabbitMQ documentation [Clustering plugin](#).

Federation

Federation plugin is used in applications that does not require as much of tight coupling as clustering. Federation has several useful features.

- Loose coupling - The federation plugin can transmit messages between brokers (or clusters) in different administrative domains:
 - they may have different users and virtual hosts;
 - they may run on different versions of RabbitMQ and Erlang.
- WAN friendliness - They can tolerate network intermittent connectivity.
- Specificity - Not everything needs to be federated (made available to other brokers). There can be local-only components.
- Scalability - Federation does not require $O(n^2)$ connections for n brokers, so it scales better.

The federation plugin allows you to make exchanges and queues federated. A federated exchange or queue can receive messages from one or more upstreams (remote exchanges and queues on other brokers). A federated exchange can route messages published upstream to a local queue. A federated queue lets a local consumer receive messages from an upstream queue.

Before we move forward, let's define upstream and downstream servers.

- Upstream server - The node that is publishing some message of interest
- Downstream server - The node connected to a different broker that wants to receive messages from the upstream server

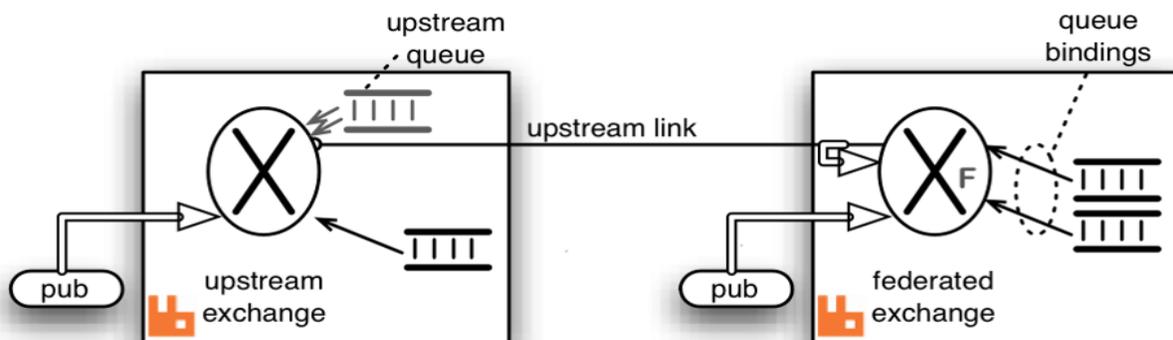
A federation link needs to be established from downstream server to the upstream server. The data flows in single direction from upstream server to downstream server. For bi-directional data flow, we would need to create federation links on both the nodes.

We can receive messages from upstream server to downstream server by either making an exchange or a queue *federated*.

For more detailed information about federation, please refer to RabbitMQ documentation [Federation plugin](#).

Federated Exchange

When we make an exchange on the downstream server *federated*, the messages published to the upstream exchanges are copied to the federated exchange, as though they were published directly to it.



Above figure explains message transfer using federated exchange. The box on the right acts as the downstream server and the box on the left acts as the upstream server. A federation/upstream link is established between the downstream server and the upstream server by using federation management plugin. An exchange on the downstream server is made *federated* using federation policy configuration. The federated exchange only receives the messages for which it has subscribed for. An upstream queue is created on the upstream server with a binding key same as subscription made on the federated exchange. For example, if an upstream server is publishing messages with binding key “foo”

and a client on the downstream server is interested in receiving messages of the binding key “foo”, then it creates a queue and binds the queue to the federated with the same binding key. This binding is sent to the upstream and the upstream queue binds to the upstream exchange with that key.

Publications to either exchange may be received by queues bound to the federated exchange, but publications directly to the federated exchange cannot be received by queues bound to the upstream exchange.

For more information about federated exchanges and different federation topologies, please read [Federated Exchanges](#).

Federated Queue

Federated queue provides a way of balancing load of a single queue across nodes or clusters. A federated queue lets a local consumer receive messages from an upstream queue. A typical use would be to have the same “logical” queue distributed over many brokers. Such a logical distributed queue is capable of having higher capacity than a single queue. A federated queue links to other upstream queues.

A federation or upstream link needs to be created like before and a federated queue needs to be setup on the downstream server using federation policy configuration. The federated queue will only retrieve messages when it has run out of messages locally, it has consumers that need messages, and the upstream queue has “spare” messages that are not being consumed.

For more information about federated queues, please read [Federated Queues](#).

Shovel

Shovel plugin allows you to reliably and continually move messages from a source in one broker to destination in another broker. A shovel behaves like a well-written client application, that

- connects to its source and destination broker
- consumes messages from the source queue
- re-publishes messages to the destination if the messages match the routing key.

Shovel plugin uses Erlang client under the hood. In case of shovel, apart from configuring the hostname, port and virtual host of the remote node, we will also have to provide list of routing keys that we want to forward to remote node. The primary advantages of shovels are

- Loose coupling - A shovel can move messages between brokers (or clusters) in different administrative domains:
 - they may have different users and virtual hosts;
 - they may run on different versions of RabbitMQ and Erlang.
- WAN friendliness - They can tolerate network intermittent connectivity.

Shovels are also useful in case if one of the nodes is behind NAT. We can setup shovel on the node behind NAT to forward messages to the node outside NAT. Shovels do not allow you to adapt to subscriptions like a federation link and we need to create a new shovel per subscription.

For more detailed information about shovel, please refer to RabbitMQ documentation [Shovel plugin](#).

Authentication in RabbitMQ

By default RabbitMQ supports SASL PLAIN authentication with user name and password. RabbitMQ supports other SASL authentication mechanism using plugins. In VOLTRON we use one such external plugin based on x509 certificates (<https://github.com/rabbitmq/rabbitmq-auth-mechanism-ssl>). This authentication is based on a technique called public key cryptography which consists of a key pair - a public key and a private key. Data that has been encrypted

with a public key can only be decrypted with the corresponding private key and vice versa. The owner of key pair makes the public key available and keeps the private confidential. To send a secure data to a receiver, a sender encrypts the data with the receiver's public key. Since only the receiver has access to his own private key only the receiver can decrypted. This ensures that others, even if they can get access to the encrypted data, cannot decrypt it. This is how public key cryptography achieves confidentiality.

Digital certificate is a digital file that is used to prove ownership of a public key. Certificates act like identification cards for it owner/entity. Certificates are hence crucial to determine that a sender is using the right public key to encrypt the data in the first place. Digital Certificates are issued by Certification Authorities(CA). Certification Authorities fulfil the role of the Trusted Third Party by accepting Certificate applications from entities, authenticating applications, issuing Certificates and maintaining status information about the Certificates issued. Each CA has its own public private key pair and its public key certificate is called a root CA certificate. The CA attests to the identity of a Certificate applicant when it signs the Digital Certificate using its private key. In x509 based authentication, a signed certificate is presented instead of username/password for authentication and if the server recognizes the the signer of the certificate as a trusted CA, accepts and allows the connection. Each server/system can maintain its own list of trusted CAs (i.e. list of public certificates of CAs). Certificates signed by any of the trusted CA would be considered trusted. Certificates can also be signed by intermediate CAs that are in turn signed by a trusted.

This section only provides a brief overview about the SSL based authentication. Please refer to the vast material available online for detailed description. Some useful links to start:

- https://en.wikipedia.org/wiki/Public-key_cryptography
- <https://robertheaton.com/2014/03/27/how-does-https-actually-work/>

Management Plugin

The rabbitmq-management plugin provides an HTTP-based API for management and monitoring of RabbitMQ nodes and clusters, along with a browser-based UI and a command line tool, *rabbitmqadmin*. The management interface allows you to

- Create, Monitor the status and delete resources such as virtual hosts, users, exchanges, queues etc.
- Monitor queue length, message rates and connection information and more
- Manage users and add permissions (read, write and configure) to use the resources
- Manage policies and runtime parameters
- Send and receive messages (for trouble shooting)

For more detailed information about the management plugin, please refer to RabbitMQ documentation [Management Plugin](#).

RabbitMQ Based VOLTRON

RabbitMQ VOLTRON uses Pika library for RabbitMQ message bus implementation. To setup VOLTRON instance to use RabbitMQ message bus, we need to first configure VOLTRON to use RabbitMQ message library. The contents of the RabbitMQ configuration file looks like below.

Path: \$VOLTRON_HOME/rabbitmq_config.yml

```
#host parameter is mandatory parameter. fully qualified domain name
host: mymachine.pnl.gov

# mandatory. certificate data used to create root ca certificate. Each volttron
# instance must have unique common-name for root ca certificate
certificate-data:
```

(continues on next page)

(continued from previous page)

```

country: 'US'
state: 'Washington'
location: 'Richland'
organization: 'PNNL'
organization-unit: 'VOLTTTRON Team'
# voltttron1 has to be replaced with actual instance name of the VOLTTTRON
common-name: 'voltttron1_root_ca'
#
# optional parameters for single instance setup
#
virtual-host: 'voltttron' # defaults to voltttron

# use the below four port variables if using custom rabbitmq ports
# defaults to 5672
amqp-port: '5672'

# defaults to 5671
amqp-port-ssl: '5671'

# defaults to 15672
mgmt-port: '15672'

# defaults to 15671
mgmt-port-ssl: '15671'

# defaults to true
ssl: 'true'

# defaults to ~/rabbitmq_server/rabbitmq_server-3.7.7
rmq-home: "~/rabbitmq_server/rabbitmq_server-3.7.7"

```

Each VOLTTTRON instance resides within a RabbitMQ virtual host. The name of the virtual host needs to be unique per VOLTTTRON instance if there are multiple virtual instances within a single host/machine. The hostname needs to be able to resolve to a valid IP. The default port of AMQP port without authentication is 5672 and with authentication is 5671. The default management HTTP port without authentication is 15672 and with authentication is 15671. These needs to be set appropriately if default ports are not used. The ‘ssl’ flag indicates if SSL based authentication is required or not. If set to True, information regarding SSL certificates needs to be also provided. SSL based authentication is described in detail in [Authentication And Authorization With RabbitMQ Message Bus](#).

To configure the VOLTTTRON instance to use RabbitMQ message bus, run the following command.

```
vcfg --rabbitmq single [optional path to rabbitmq_config.yml]
```

At the end of the setup process, RabbitMQ broker is setup to use the configuration provided. A new topic exchange for the VOLTTTRON instance is created within the configured virtual host.

On platform startup, VOLTTTRON checks for the type of message bus to be used. If using RabbitMQ message bus, the RabbitMQ platform router is instantiated. The RabbitMQ platform router,

- Connects to RabbitMQ broker (with or without authentication)
- Creates a VIP queue and binds itself to the “VOLTTTRON” exchange with binding key “<instance-name>.router”. This binding key makes it unique across multiple VOLTTTRON instances in a single machine as long as each instance has a unique instance name.
- Handles messages intended for router module such as “hello”, “peerlist”, “query” etc.
- Handles unrouteable messages - Messages which cannot be routed to any destination agent are captured and an error message indicating “Host Unreachable” error is sent back to the caller.

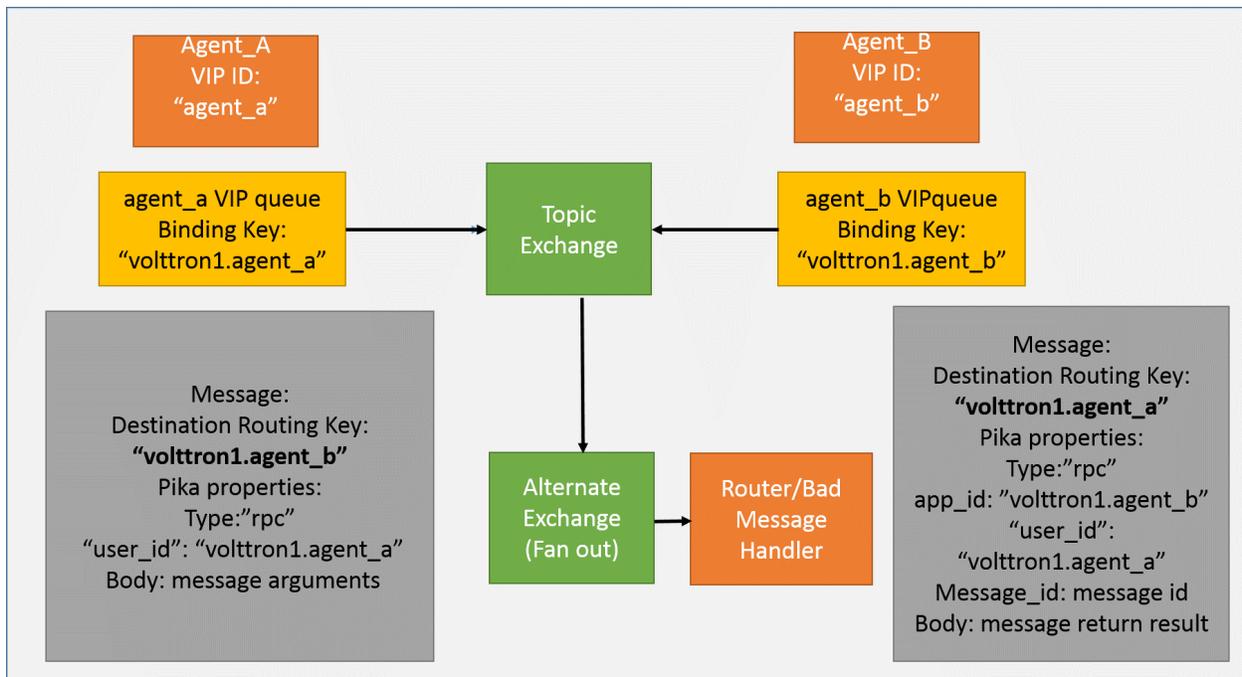
- Disconnects from the broker when the platform shuts down.

When any agent is installed and started, the Agent Core checks for the type of message bus used. If it is RabbitMQ message bus then

- It creates a RabbitMQ user for the agent.
- If SSL based authentication is enabled, client certificates for the agent is created.
- Connect to the RabbitMQ broker with appropriate connection parameters
- Creates a VIP queue and binds itself to the “VOLTRON” exchange with binding key “<instance-name>.<agent identity>”.
- Sends and receives messages using Pika library methods.
- Checks for the type of subsystem in the message packet that it receives and calls the appropriate subsystem message handler.
- Disconnects from the broker when the agent stops or platform shuts down.

RPC In RabbitMQ VOLTRON

The agent functionality remain unchanged irrespective of the underlying message bus used. That means they can continue to use the same RPC interfaces without any change.



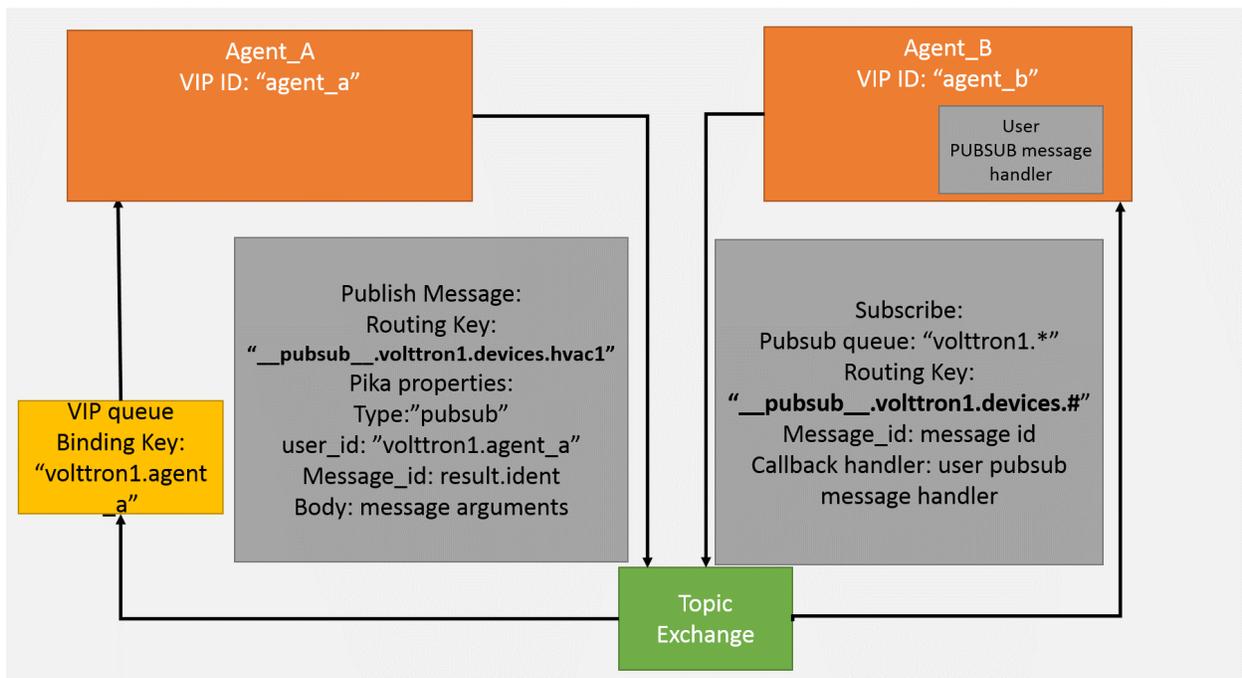
Consider two agents with VIP identities “agent_a” and “agent_b” connected to VOLTRON platform with instance name “volttron1”. Agent A and B each have a VIP queue with binding key “volttron1.agent_a” and “volttron1.agent_b”. Following is the sequence of operation when Agent A wants to make RPC call to Agent B.

1. Agent A make RPC call to Agent B. `agent_a.vip.rpc.call("agent_b", set_point, "point_name", 2.5)`
2. RPC subsystem wraps this call into a VIP message object and sends it to Agent B.
3. The VOLTRON exchange routes the message to Agent B as the destination routing in the VIP message object matches with the binding key of Agent B.

4. Agent Core on Agent B receives the message, unwraps the message to find the subsystem type and calls the RPC subsystem handler.
5. RPC subsystem makes the actual RPC call “set_point()” and gets the result. It then wraps into VIP message object and sends it back to the caller.
6. The VOLTRON exchange routes it back to Agent A.
7. Agent Core on Agent A calls the RPC subsystem handler which in turn hands over the RPC result to Agent A application.

PUBSUB In RabbitMQ VOLTRON

The agent functionality remains unchanged irrespective of the platform using ZeroMQ based pubsub or RabbitMQ based pubsub i.e, agents continue to use the same PubSub interfaces and use the same topic format delimited by “/”. Since RabbitMQ expects binding key to be delimited by ‘.’, RabbitMQ PUBSUB internally replaces ‘/’ with “.”. Additionally, all agent topics converted to “__pubsub__.<instance_name>.<remainder of topic>” to differentiate from main Agent VIP queue binding.



Consider two agents with VIP identities “agent_a” and “agent_b” connected to VOLTRON platform with instance name “volttron1”. Agent A and B each have a VIP queue with binding key “volttron1.agent_a” and “volttron1.agent_b”. Following is the sequence of operation when Agent A subscribes to a topic and Agent B publishes to same the topic.

1. **Agent B makes subscribe call for topic “devices”**. `agent_b.vip.pubsub.subscribe("pubsub", prefix="devices", callback=self.onmessage)`
2. Pubsub subsystem creates binding key from the topic “__pubsub__.volttron1.devices.#”
3. It creates a queue internally and binds the queue to the VOLTRON exchange with the above binding key.
4. Agent B is publishing messages with topic: “devices/hvac1”. `agent_b.vip.pubsub.publish("pubsub", topic="devices/hvac1", headers={}, message="foo")`.
5. PubSub subsystem internally creates a VIP message object and publishes on the VOLTRON exchange.

6. RabbitMQ broker routes the message to Agent B as routing key in the message matches with the binding key of the topic subscription.
7. The pubsub subsystem unwraps the message and calls the appropriate callback method of Agent A.

If agent wants to subscribe to topic from remote instances, it uses `agent.vip.subscribe("pubsub", "devices.hvac1", all_platforms=True)`

It is internally set to “`__pubsub__.*.<remainder of topic>`”

Pubsub subsystem for ZeroMQ message bus performs $O(N)$ comparisons where N is the number of unique subscriptions. RabbitMQ Topic Exchange was enhanced in version 2.6.0 to reduce the overhead of additional unique subscriptions to almost nothing in most cases. We speculate they are using a tree structure to store the binding keys which would reduce the search time to $O(1)$ in most cases and $O(\ln)$ in the worst case. VOLTRON PUBSUB with ZeroMQ could be updated to match this performance scalability with some effort.

Multi-Platform Communication In RabbitMQ VOLTRON

With ZeroMQ based VOLTRON, multi-platform communication was accomplished in three different ways.

1. Direct connection to remote instance - Write an agent that would connect to remote instance directly.
2. Special agents - Use special agents such as forward historian/data puller agents that would forward/receive messages to/from remote instances. In RabbitMQ-VOLTRON, we make use of shovel plugin to achieve this behavior. Please refer to *Shovel Plugin* to get an overview of shovels.
3. Multi-Platform RPC and PubSub - Configure VIP address of all remote instances that an instance has to connect to in it's `$VOLTRON_HOME/external_discovery.json` and let the router module in each instance manage the connection and take care of the message routing for us. In RabbitMQ-VOLTRON, we make use of federation plugin to achieve this behavior. Please refer to *Federation Plugin* get an overview of federation.

Using Federation Plugin

We can connect multiple VOLTRON instances using the federation plugin. Before setting up federation links, we need to first identify upstream server and downstream server. Upstream Server is the node that is publishing some message of interest and downstream server is the node that wants to receive messages from the upstream server. A federation link needs to be established from a downstream VOLTRON instance to upstream VOLTRON instance. To setup a federation link, we will need to add upstream server information in a RabbitMQ federation configuration file

Path: `$VOLTRON_HOME/rabbitmq_federation_config.yml`

```
# Mandatory parameters for federation setup
federation-upstream:
  rabbit-4:
    port: '5671'
    virtual-host: voltron4
  rabbit-5:
    port: '5671'
    virtual-host: voltron5
```

To configure the VOLTRON instance to setup federation, run the following command.

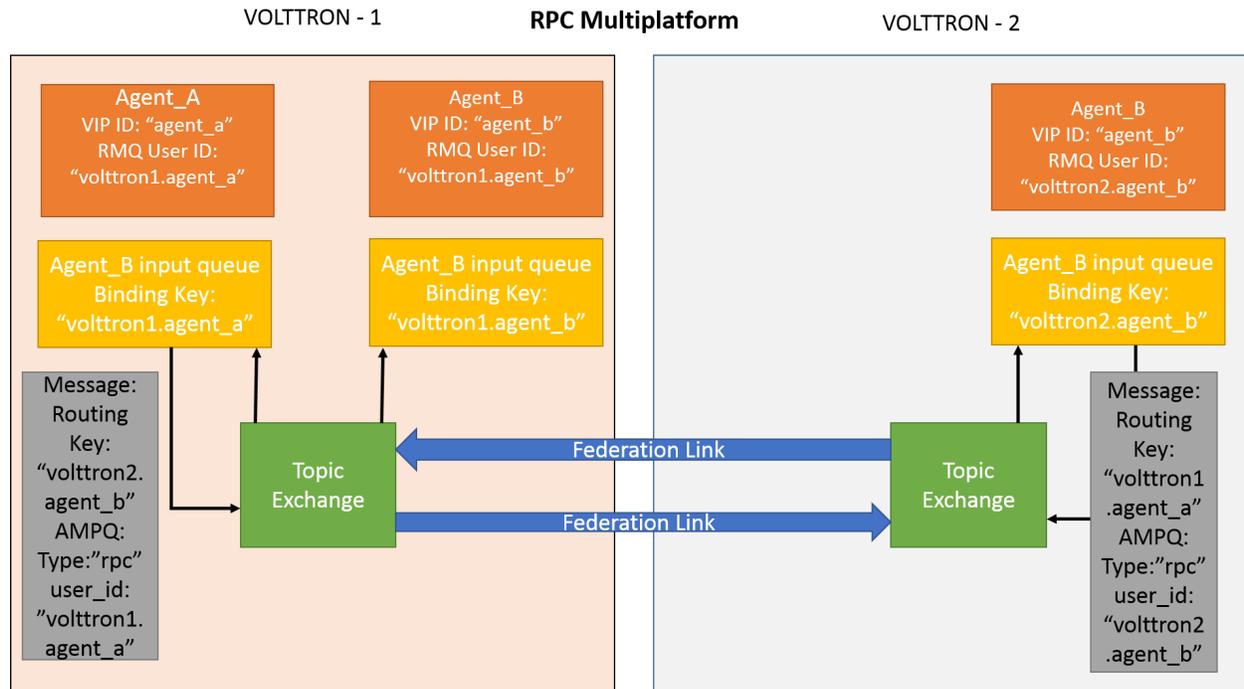
```
vfg -rabbitmq federation [optional path to rabbitmq_federation_config.yml]
```

This will setup federation links to upstream servers and sets policy to make the VOLTRON exchange *federated*. Once a federation link is established to remote instance, the messages published on the remote instance become available to local instance as if it were published on the local instance.

For detailed instructions to setup federation, please refer to README section <>.

Multi-Platform RPC With Federation

For multi-platform RPC communication, federation links need to be established on both the VOLTTRON nodes. Once the federation links are established, RPC communication becomes fairly simple.



Consider Agent A on volttron instance "volttron1" on host "host_A" wants to make RPC call on Agent B on VOLTTRON instance "volttron2" on host "host_B".

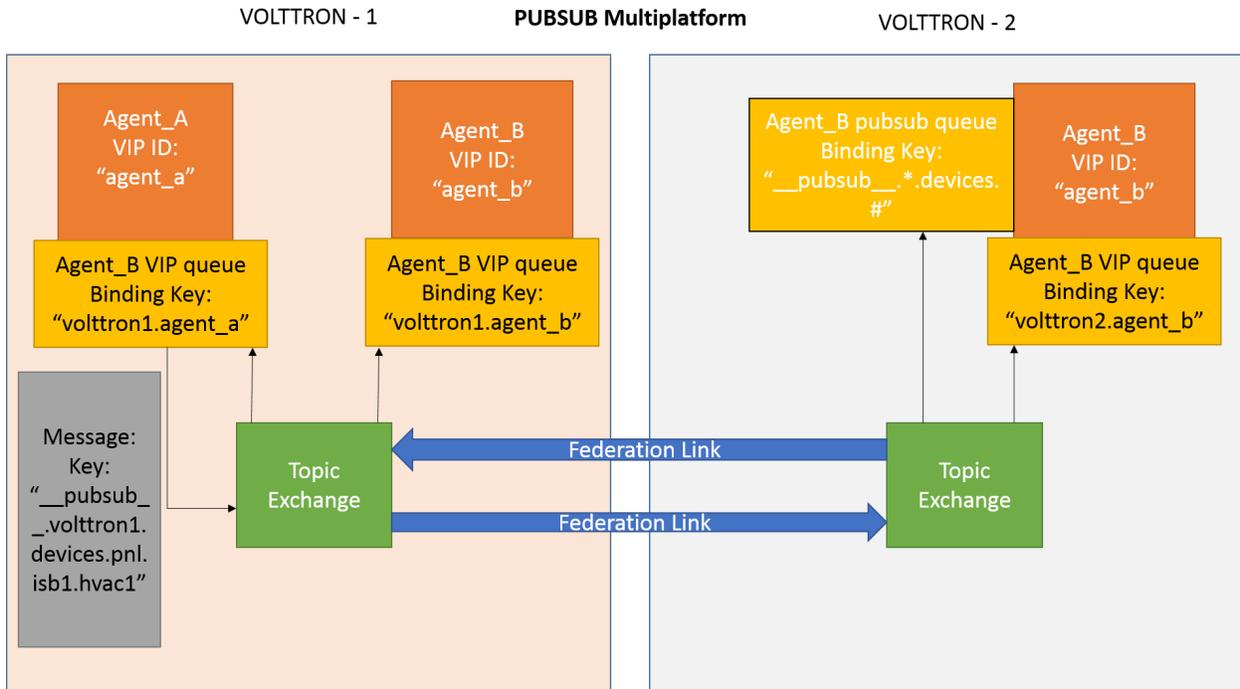
1. Agent A makes RPC call. .. code-block:: Python

```
kwargs = {"external_platform": self.destination_instance_name} agent_a.vip.rpc.call("agent_b",
set_point, "point_name", 2.5, **kwargs)
```

2. The message is transferred over federation link to VOLTTRON instance "volttron2" as both the exchanges are made *federated*.
3. RPC subsystem of Agent B calls the actual RPC method and gets the result. It encapsulates the message result into VIP message object and sends it back to Agent A on VOLTTRON instance "volttron1".
4. The RPC subsystem on Agent A receives the message result and gives it to Agent A application.

Multi-Platform PubSub With Federation

For multi-platform PubSub communication, it is sufficient to have federation link from downstream server to upstream server. In case of bi-directional data flow, links have to established in both the directions.



Consider Agent B on voltrtron instance “volttron2” on host “host_B” wants to subscribe to messages from VOLTRON instance “volttron1” on host “host_A”. Firstly, federation link needs to be established from “volttron2” to “volttron1”.

1. Agent B makes a subscribe call.

```
agent_b.vip.subscribe.call("pubsub", prefix="devices", all_platforms=True)
```

2. The PubSub subsystem converts the prefix to “_pubsub_*.devices.#”. Here, “*” indicates that agent is subscribing to “devices” topic from all the VOLTRON platforms.
3. A new queue is created and bound to VOLTRON exchange with above binding key. Since the VOLTRON exchange is a *federated exchange*, any subscribed message on the upstream server becomes available on the federated exchange and Agent B will be able to receive it.
4. Agent A publishes message to topic “devices/pnl/isb1/hvac1”
5. PubSub subsystem publishes this message on its VOLTRON exchange.
6. Due to the federation link, message is received by the Pubsub subsystem of Agent B.

Using Shovel Plugin

Shovels act as well written client application which moves messages from source to destination broker. Below configuration shows how to setup a shovel to forward PubSub messages or perform multi-platform RPC communication from local to a remote instance. It expects hostname, port and virtual host of remote instance.

Path: \$VOLTRON_HOME/rabbitmq_shovel_config.yml

```
# Mandatory parameters for shovel setup
shovel:
  rabbit-2:
    port: '5671'
    virtual-host: volttron
    # Configuration to forward pubsub topics
```

(continues on next page)

(continued from previous page)

```
pubsub:
  # Identity of agent that is publishing the topic
  platform.driver:
    - devices
  # Configuration to make remote RPC calls
  rpc:
    # Remote instance name
    voltron2:
      # List of pair of agent identities (local caller, remote callee)
      - [scheduler, platform.actuator]
```

To forward PubSub messages, the topic and agent identity of the publisher agent is needed. To perform RPC, instance name of the remote instance and agent identities of the local agent and remote agent are needed.

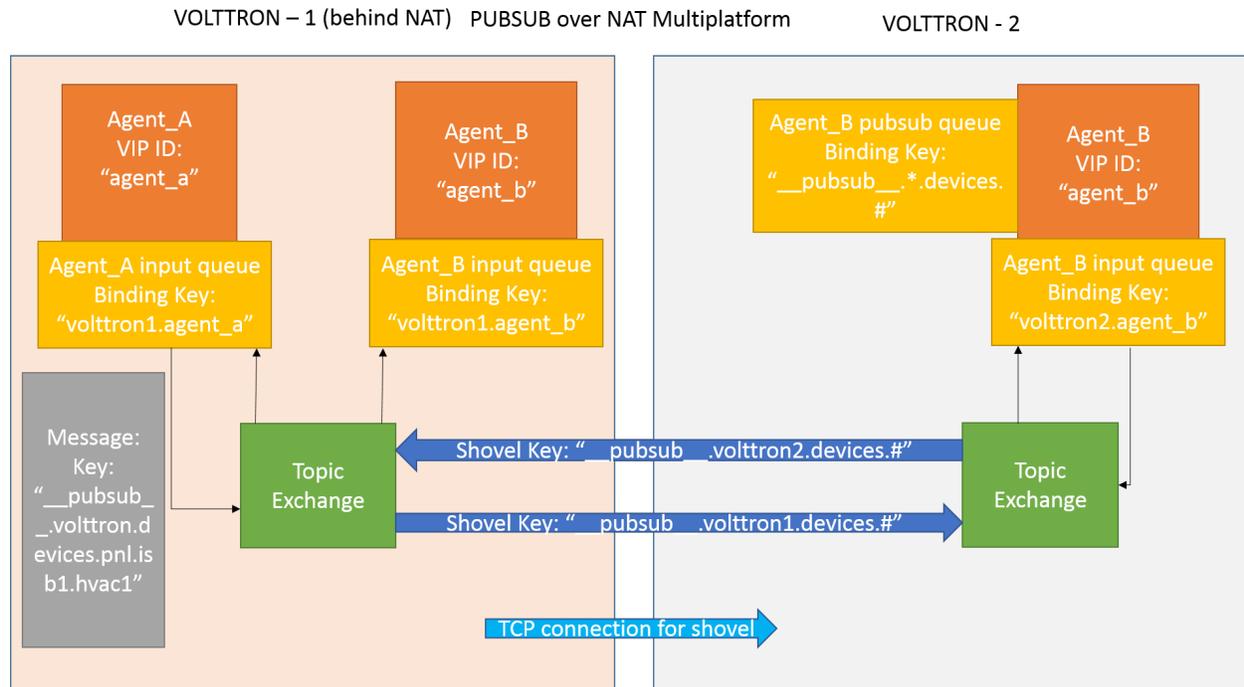
To configure the VOLTRON instance to setup shovel, run the following command.

```
vcfg --rabbitmq shovel [optional path to rabbitmq_shovel_config.yml]
```

This setups up a shovel that forwards messages (either PubSub or RPC) from local exchange to remote exchange.

Multi-Platform PubSub With Shovel

After the shovel link is established for Pubsub, the below figure shows how the communication happens. Please note, for bi-directional pubsub communication, shovel links need to be created on both the nodes. The “blue” arrows show the shovel binding key. The pubsub topic configuration in `$VOLTRON_HOME/rabbitmq_shovel_config.yml` get internally converted to shovel binding key, “`__pubsub__.<local instance name>.<actual topic>`”.



Now consider a case where shovels are setup in both the directions for forwarding “devices” topic.

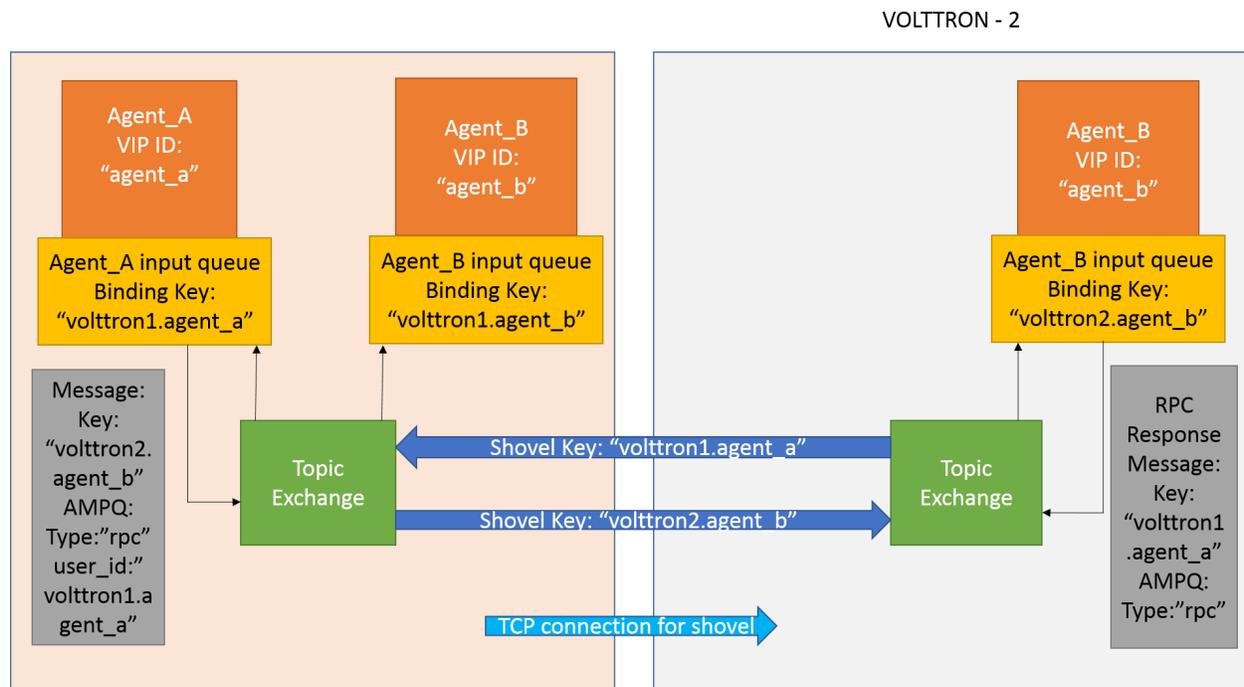
1. Agent B makes a subscribe call to receive messages with topic “devices” from all connected platforms.

```
agent_b.vip.subscribe.call("pubsub", prefix="devices", all_platforms=True)
```

- The PubSub subsystem converts the prefix to “__pubsub__.*.devices.#” “*” indicates that agent is subscribing to “devices” topic from all the VOLTRON platforms.
- A new queue is created and bound to VOLTRON exchange with above binding key.
- Agent A publishes message to topic “devices/pnml/isb1/hvac1”
- PubSub subsystem publishes this message on it’s VOLTRON exchange.
- Due to a shovel link from VOLTRON instance “voltrtron1” to “voltrtron2”, the message is forwarded from voltrtron exchange “voltrtron1” to “voltrtron2” and picked up by Agent A on “voltrtron2”.

Multi-Platform RPC With Shovel

After the shovel link is established for multi-platform RPC, the below figure shows how the RPC communication happens. Please note it is mandatory to have shovel links on both directions as it is request-response type of communication. We will need to set the agent identities for caller and callee in the `$VOLTRON_HOME/rabbitmq_shovel_config.yml`. The “blue” arrows show the resulting the shovel binding key.



Consider Agent A on voltrtron instance “voltrtron1” on host “host_A” wants to make RPC call on Agent B on VOLTRON instance “voltrtron2” on host “host_B”.

- Agent A makes RPC call. .. code-block:: Python

```
kwargs = {"external_platform": self.destination_instance_name} agent_a.vip.rpc.call("agent_b",
set_point, "point_name", 2.5, **kwargs)
```

- The message is transferred over shovel link to VOLTRON instance “voltrtron2”.
- RPC subsystem of Agent B calls the actual RPC method and gets the result. It encapsulates the message result into VIP message object and sends it back to Agent A on VOLTRON instance “voltrtron1”.
- The RPC subsystem on Agent A receives the message result and gives it to Agent A application.

RabbitMQ Management Tool Integrated Into VOLTRON

Some of the important native RabbitMQ control and management commands are now integrated with “volttron-ctl” utility. Using volttron-ctl RabbitMQ management utility, we can control and monitor the status of RabbitMQ message bus.

```

volttron-ctl rabbitmq --help
usage: volttron-ctl command [OPTIONS] ... rabbitmq [-h] [-c FILE] [--debug]
                                                [-t SECS]
                                                [--msgdebug MSGDEBUG]
                                                [--vip-address ZMQADDR]
                                                ...

subcommands:

  add-vhost          add a new virtual host
  add-user           Add a new user. User will have admin privileges
                    i.e,configure, read and write
  add-exchange       add a new exchange
  add-queue          add a new queue
  list-vhosts        List virtual hosts
  list-users         List users
  list-user-properties
                    List users
  list-exchanges     add a new user
  list-exchange-properties
                    list exchanges with properties
  list-queues        list all queues
  list-queue-properties
                    list queues with properties
  list-bindings      list all bindings with exchange
  list-federation-parameters
                    list all federation parameters
  list-shovel-parameters
                    list all shovel parameters
  list-policies      list all policies
  remove-vhosts      Remove virtual host/s
  remove-users       Remove virtual user/s
  remove-exchanges   Remove exchange/s
  remove-queues      Remove queue/s
  remove-federation-parameters
                    Remove federation parameter
  remove-shovel-parameters
                    Remove shovel parameter
  remove-policies    Remove policy

```

Authentication And Authorization With RabbitMQ Message Bus

Authentication In RabbitMQ VOLTRON

RabbitMQ VOLTRON uses SSL based authentication, rather than the default username and password authentication. VOLTRON adds SSL based configuration entries into the ‘rabbitmq.conf’ file during the setup process. The necessary SSL configurations can be seen by running the following command:

```
cat ~/rabbitmq_server/rabbitmq_server-3.7.7/etc/rabbitmq/rabbitmq.conf
```

The configurations required to enable SSL:

```
listeners.ssl.default = 5671
ssl_options.cacertfile = VOLTRON_HOME/certificates/certs/volttron1-trusted-cas.crt
ssl_options.certfile = VOLTRON_HOME/certificates/certs/volttron1-server.crt
ssl_options.keyfile = VOLTRON_HOME/certificates/private/volttron1-server.pem
ssl_options.verify = verify_peer
ssl_options.fail_if_no_peer_cert = true
```

Parameter explanations

- `listeners.ssl.default`: port for listening for SSL connections
- `ssl_options.cacertfile`: path to trusted Certificate Authorities (CA)
- `ssl_options.certfile`: path to server public certificate
- `ssl_options.keyfile`: path to server's private key
- `ssl_options.verify`: whether verification is enabled
- `ssl_options.fail_if_no_peer_cert`: upon client's failure to provide certificate, SSL connection either rejected (true) or accepted (false)
- `auth_mechanisms.1`: type of authentication mechanism. EXTERNAL means SSL authentication is used

SSL in RabbitMQ VOLTRON

To configure RabbitMQ-VOLTRON to use SSL based authentication, we need to add SSL configuration in `rabbitmq_config.yml`.

```
#host parameter is mandatory parameter. fully qualified domain name
host: mymachine.pnl.gov

# mandatory. certificate data used to create root ca certificate. Each volttron
# instance must have unique common-name for root ca certificate
certificate-data:
  country: 'US'
  state: 'Washington'
  location: 'Richland'
  organization: 'PNNL'
  organization-unit: 'VOLTRON Team'
  # volttron1 has to be replaced with actual instance name of the VOLTRON
  common-name: 'volttron1_root_ca'

virtual-host: 'volttron' # defaults to volttron

# use the below four port variables if using custom rabbitmq ports
# defaults to 5672
amqp-port: '5672'

# defaults to 5671
amqp-port-ssl: '5671'

# defaults to 15672
mgmt-port: '15672'

# defaults to 15671
mgmt-port-ssl: '15671'
```

(continues on next page)

(continued from previous page)

```
# defaults to true
ssl: 'true'

# defaults to ~/rabbitmq_server/rabbitmq_server-3.7.7
rmq-home: "~/rabbitmq_server/rabbitmq_server-3.7.7"
```

The parameters of interest for SSL based configuration are

- certificate-data: subject information needed to create certificates
- ssl: Flag set to 'true' for SSL based authentication
- amqp-port-ssl: Port number for SSL connection (defaults to 5671)
- mgmt-port-ssl: Port number for HTTPS management connection (defaults to 15671)

We can then configure the VOLTRON instance to use SSL based authentication with the below command.

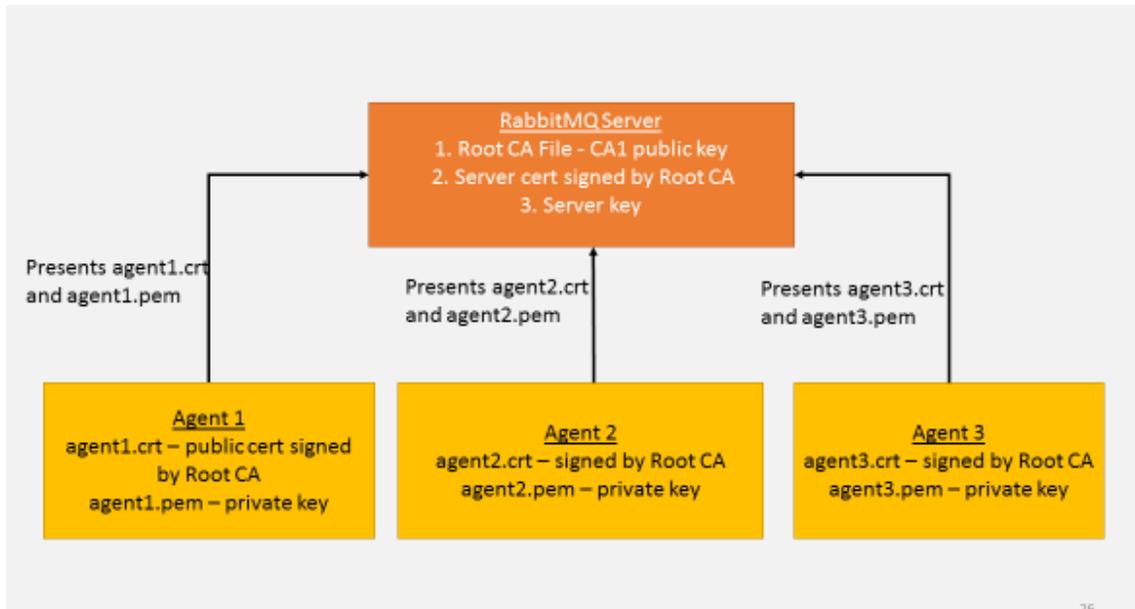
```
vcfg -rabbitmq single <optional path to rabbitmq_config.yml>
```

When one creates a single instance of RabbitMQ, the following is created / re-created in the VOLTRON_HOME/certificates directory:

- Public and private certificates of root Certificate Authority (CA)
- Public and private (automatically signed by the CA) server certificates needed by RabbitMQ broker
- Admin certificate for the RabbitMQ instance
- Public and private (automatically signed by the CA) certificates for VOLTRON platform service agents.
- Trusted CA certificate

The public files can be found at VOLTRON_HOME/certificates/certs and the private files can be found at VOLTRON_HOME/certificates/private. The trusted-cas.crt file is used to store the root CAs of all VOLTRON instances that the RabbitMQ server has to connected to. The trusted ca is only created once, but can be updated. Initially, the trusted ca is a copy of the the root CA file, but when an external VOLTRON instance needs to be connected to an instance, then external VOLTRON instance's root CA have to be appended to this file in order for RabbitMQ broker to trust the new connection.

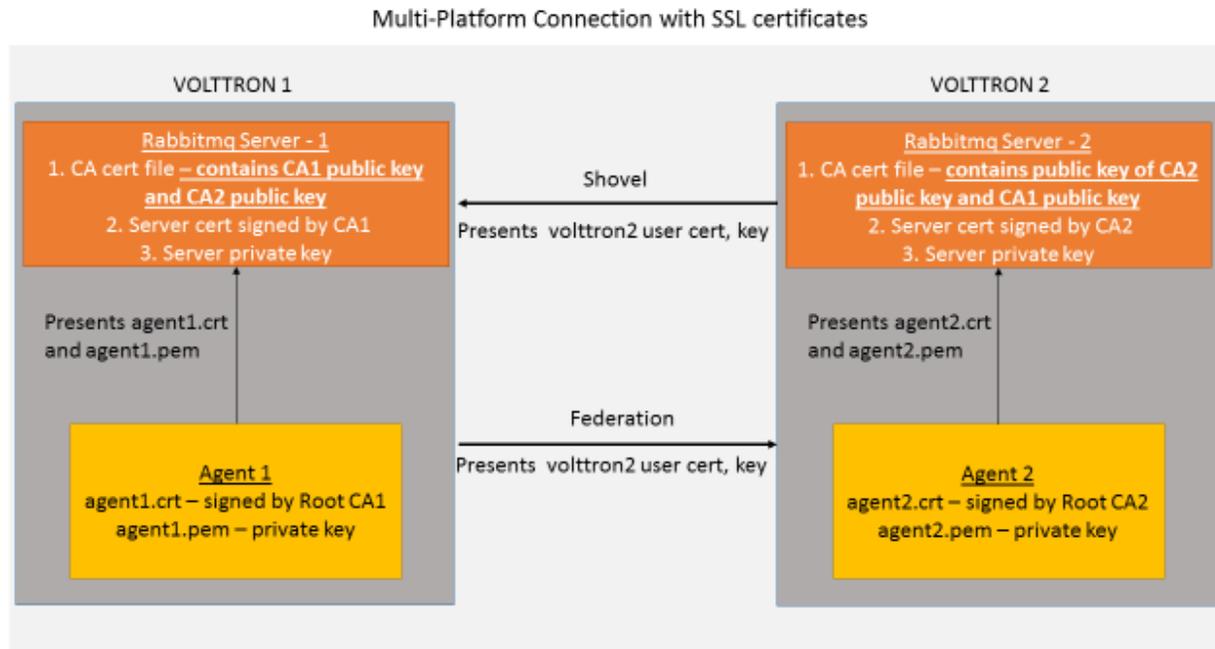
RabbitMQ Server SSL certificates



Every RabbitMQ has a single self signed root ca and server certificate signed by the root CA. This is created during VOLTRON setup and the RabbitMQ server is configured and started with these two certificates. Every time an agent is started, the platform automatically creates a pair of public-private certificates for that agent that is signed by the same root CA. When an agent communicates with the RabbitMQ message bus it presents its public certificate and private key to the server and the server validates if it is signed by a root CA it trusts – ie., the root certificate it was started with. Since there is only a single root CA for one VOLTRON instance, all the agents in this instance can communicate with the message bus over SSL.

Multi-Platform Communication With RabbitMQ SSL

For multi-platform communication over federation and shovel, we need connecting instances to trust each other.



Suppose there are two VMs (VOLTRON1 and VOLTRON2) running single instances of RabbitMQ, and VOLTRON1 and VOLTRON2 want to talk to each other via either the federation or shovel plugins. In order for VOLTRON1 to talk to VOLTRON2, VOLTRON1's root certificate must be appended to VOLTRON2's trusted CA certificate, so that when VOLTRON1 presents its root certificate during connection, VOLTRON2's RabbitMQ server can trust the connection. VOLTRON2's root CA must be appended to VOLTRON1's root CA and it must in turn present its root certificate during connection, so that VOLTRON1 will know it is safe to talk to VOLTRON2.

Agents trying to connect to remote instance directly, need to have a public certificate signed by the remote instance for authenticated SSL based connection. To facilitate this process, the VOLTRON platform exposes a web based server api for requesting, listing, approving and denying certificate requests. For more detailed description, refer to Agent communication to Remote RabbitMQ instance

Authorization in RabbitMQ VOLTRON

To be implemented in VOLTRON

For more detailed information about access control, please refer to RabbitMQ documentation [Access Control](#).

3.5.11 Open ADR

OpenADR (Automated Demand Response) is a standard for alerting and responding to the need to adjust electric power consumption in response to fluctuations in grid demand. OpenADR communications are conducted between Virtual Top Nodes (VTNs) and Virtual End Nodes (VENs).

In this implementation, a VOLTRON agent, OpenADRVenAgent, is made available as a VOLTRON service. It acts as a VEN, communicating with its VTN via EiEvent and EiReport services in conformance with a subset of the OpenADR 2.0b specification.

A VTN server has also been implemented, with source code in the `kisensum/openadr` folder of the `voltron-applications` git repository. As has been described below, it communicates with the VEN and provides a web user interface for defining and reporting on Open ADR events.

The OpenADR 2.0b specification (<http://www.openadr.org/specification>) is available from the OpenADR Alliance. This implementation also generally follows the DR program characteristics of the Capacity Program described in Section 9.2 of the OpenADR Program Guide (http://www.openadr.org/assets/openadr_drprogramguide_v1.0.pdf).

The OpenADR Capacity Bidding program relies on a pre-committed agreement about the VEN's load shed capacity. This agreement is reached in a bidding process transacted outside of the OpenADR interaction, typically with a long-term scope, perhaps a month or longer. The VTN can "call an event," indicating that a load-shed event should occur in conformance with this agreement. The VTN indicates the level of load shedding desired, when the event should occur, and for how long. The VEN responds with an "optIn" acknowledgment. (It can also "optOut," but since it has been pre-committed, an "optOut" may incur penalties.)

OpenADR VEN Agent: Installation and Configuration

The VEN agent can be configured, built and launched using the VOLTRON agent installation process described in http://voltron.readthedocs.io/en/develop/devguides/agent_development/Agent-Development.html#agent-development.

The VEN agent depends on some third-party libraries that are not in the standard VOLTRON installation. They should be installed in the VOLTRON virtual environment prior to building the agent:

```
(volttron) $ cd $VOLTRON_ROOT/services/core/OpenADRVenAgent
(volttron) $ pip install -r requirements.txt
```

where `$VOLTRON_ROOT` is the base directory of the cloned VOLTRON code repository.

The VEN agent is designed to work in tandem with a "control agent," another VOLTRON agent that uses VOLTRON RPC calls to manage events and supply report data. A sample control agent has been provided in the `test/ControlAgentSim` subdirectory under `OpenADRVenAgent`.

The VEN agent maintains a persistent store of event and report data in `$VOLTRON_HOME/data/openadr.sqlite`. Some care should be taken in managing the disk consumption of this data store. If no events or reports are active, it is safe to take down the VEN agent and delete the file; the persistent store will be reinitialized automatically on agent startup.

Configuration Parameters

The VEN agent's configuration file contains JSON that includes several parameters for configuring VTN server communications and other behavior. A sample configuration file, `openadrven.config`, has been provided in the agent directory.

The VEN agent supports the following configuration parameters:

Parameter	Example	Description
db_path	“\$VOLTTRON_HOME/ope-nadr.sqlite”	Pathname of the agent’s sqlite database. Shell variables will be expanded if they are present in the pathname.
ven_id	“0”	The OpenADR ID of this virtual end node. Identifies this VEN to the VTN. If automated VEN registration is used, the ID is assigned by the VTN at that time. If the VEN is registered manually with the VTN (i.e., via configuration file settings), then a common VEN ID should be entered in this config file and in the VTN’s site definition.
ven_name	“ven01”	Name of this virtual end node. This name is used during automated registration only, identifying the VEN before its VEN ID is known.
vtn_id	“vtn01”	OpenADR ID of the VTN with which this VEN communicates.
vtn_address	“http://openadr-vtn.ki-evi.com:8000”	URL and port number of the VTN.
send_registration	“False”	(“True” or “False”) If “True”, the VEN sends a one-time automated registration request to the VTN to obtain the VEN ID. If automated registration will be used, the VEN should be run in this mode initially, then shut down and run with this parameter set to “False” thereafter.
security_level	“standard”	If ‘high’, the VTN and VEN use a third-party signing authority to sign and authenticate each request. The default setting is “standard”: the XML payloads do not contain Signature elements.
poll_interval	30 secs	(integer) How often the VEN should send an OadrPoll request to the VTN. The poll interval cannot be more frequent than the VEN’s 5-second process loop frequency.
log_xml	“False”	(“True” or “False”) Whether to write each inbound/outbound request’s XML data to the agent’s log.
opt_in_timeout	1800 secs	(integer) How long to wait before making a default optIn/optOut decision.
opt_in_default	optOut	(“True” or “False”) Which optIn/optOut choice to make by default.
request_events_on_startup	“False”	(“True” or “False”) Whether to ask the VTN for a list of current events during VEN startup.
report_parameters	(see below)	A dictionary of definitions of reporting/telemetry parameters.

Reporting Configuration

The VEN’s reporting configuration, specified as a dictionary in the agent configuration, defines each telemetry element (metric) that the VEN can report to the VTN, if requested. By default, it defines reports named “telemetry” and “telemetry_status”, with a report configuration dictionary containing the following parameters:

“telemetry” report: parameters	Example	Description
report_name	“TELEMETRY_USAGE”	Friendly name of the report.
report_name_metadata	“META-DATA_TELEMETRY_USAGE”	Friendly name of the report’s metadata, when sent to the VEN’s oadrRegisterReport request.
report_specifier_id	“telemetry”	Uniquely identifies the report’s data set.
report_interval_secs_default	“300”	How often to send a reporting update to the VTN.
telemetry_parameters (baseline_power_kw): r_id	“baseline_power”	(baseline_power) Unique ID of the metric.
telemetry_parameters (baseline_power_kw): report_type	“baseline”	(baseline_power) The type of metric being reported.
telemetry_parameters (baseline_power_kw): reading_type	“Direct Read”	(baseline_power) How the metric was calculated.
telemetry_parameters (baseline_power_kw): units	“powerReal”	(baseline_power) The reading’s data type.
telemetry_parameters (baseline_power_kw): method_name	“get_baseline_power”	(baseline_power) The VEN method to use when extracting the data for reporting.
telemetry_parameters (baseline_power_kw): min_frequency	30	(baseline_power) The metric’s minimum sampling frequency.
telemetry_parameters (baseline_power_kw): max_frequency	60	(baseline_power) The metric’s maximum sampling frequency.
telemetry_parameters (current_power_kw): r_id	“actual_power”	(current_power) Unique ID of the metric.
telemetry_parameters (current_power_kw): report_type	“reading”	(current_power) The type of metric being reported.
telemetry_parameters (current_power_kw): reading_type	“Direct Read”	(current_power) How the metric was calculated.
telemetry_parameters (current_power_kw): units	“powerReal”	(baseline_power) The reading’s data type.
telemetry_parameters (current_power_kw): method_name	“get_current_power”	(current_power) The VEN method to use when extracting the data for reporting.
telemetry_parameters (current_power_kw): min_frequency	30	(current_power) The metric’s minimum sampling frequency.
telemetry_parameters (current_power_kw): max_frequency	60	(current_power) The metric’s maximum sampling frequency.

“telemetry_status” report: parameters	Example	Description
report_name	“TELEMETRY_STATUS”	Friendly name of the report.
report_name_metadata	“META-DATA_TELEMETRY_STATUS”	Friendly name of the report’s metadata, when sent by the VEN’s oadrRegisterReport request.
report_specifier_id	“telemetry_status”	Uniquely identifies the report’s data set.
report_interval_secs_default	“300”	How often to send a reporting update to the VTN.
telemetry_parameters (Status): r_id	“Status”	Unique ID of the metric.
telemetry_parameters (Status): report_type	“x-resourceStatus”	The type of metric being reported.
telemetry_parameters (Status): reading_type	“x-notApplicable”	How the metric was calculated.
telemetry_parameters (Status): units	“”	The reading’s data type.
telemetry_parameters (Status): method_name	“”	The VEN method to use when extracting the data for reporting.
telemetry_parameters (Status): min_frequency	60	The metric’s minimum sampling frequency.
telemetry_parameters (Status): max_frequency	120	The metric’s maximum sampling frequency.

OpenADR VEN Agent: Operation

Events:

- The VEN maintains a persistent record of DR events.
- Event updates (including creation) trigger publication of event JSON on the VOLTTRON message bus.
- Another VOLTTRON agent (a “control agent”) can get notified immediately of event updates by subscribing to event publication. It can also call `get_events()` to retrieve the current status of each active DR event.

Reporting:

- The VEN reports device status and usage telemetry to the VTN, relying on information received periodically from other VOLTTRON agents.
- The VEN config defines telemetry values (data points) that can be reported to the VTN.
- The VEN maintains a persistent record of telemetry values over time.
- Other VOLTTRON agents are expected to call `report_telemetry()` to supply the VEN with a regular stream of telemetry values for reporting.
- The VTN can identify which of the VEN’s supported data points needs to be actively reported at a given time, including their reporting frequency.
- Another VOLTTRON agent (a “control agent”) can get notified immediately of changes in telemetry reporting requirements by subscribing to publication of “telemetry parameters.” It can also call `get_telemetry_parameters()` to retrieve the current set of reporting requirements.
- The VEN persists these reporting requirements so that they survive VOLTTRON restarts.

VOLTRON Agent Interface

The VEN implements the following VOLTRON PubSub and RPC calls.

PubSub: Event Update

When an event is created/updated, the event is published with a topic that includes 'openadr/event/{ven_id}'.

Event JSON structure:

```
{
  "event_id"      : String,
  "creation_time" : DateTime - UTC,
  "start_time"   : DateTime - UTC,
  "end_time"     : DateTime - UTC,
  "priority"     : Integer,      # Values: 0, 1, 2, 3. Usually expected to be 1.
  "signals"      : String,      # Values: json string describing one or more
↳ signals.
  "status"       : String,      # Values: unresponded, far, near, active, completed,
↳ canceled.
  "opt_type"     : String       # Values: optIn, optOut, none.
}
```

If an event status is 'unresponded', the VEN is awaiting a decision on whether to optIn or optOut. The downstream agent that subscribes to this PubSub message should communicate that choice to the VEN by calling `respond_to_event()` (see below). The VEN then relays the choice to the VTN.

PubSub: Telemetry Parameters Update

When the VEN telemetry reporting parameters have been updated (by the VTN), they are published with a topic that includes 'openadr/status/{ven_id}'.

These parameters include state information about the current report.

Telemetry parameters structure:

```
{
  'telemetry': '{
    "baseline_power_kw": {
      "r_id"      : "baseline_power",      # ID of the reporting metric
      "report_type" : "baseline",         # Type of reporting metric, e.
↳ g. baseline or reading
      "reading_type" : "Direct Read",     # (per OpenADR telemetry_
↳ usage report requirements)
      "units"       : "powerReal",        # (per OpenADR telemetry_
↳ usage report requirements)
      "method_name" : "get_baseline_power", # Name of the VEN agent
↳ method that gets the metric
      "min_frequency" : (Integer),        # Data capture frequency in
↳ seconds (minimum)
      "max_frequency" : (Integer)        # Data capture frequency in
↳ seconds (maximum)
    },
    "current_power_kw": {
      "r_id"      : "actual_power",      # ID of the reporting metric
```

(continues on next page)

(continued from previous page)

```

        "report_type"      : "reading",          # Type of reporting metric, e.
↪g. baseline or reading
        "reading_type"    : "Direct Read",      # (per OpenADR telemetry_
↪usage report requirements)
        "units"           : "powerReal",        # (per OpenADR telemetry_
↪usage report requirements)
        "method_name"     : "get_current_power", # Name of the VEN agent,
↪method that gets the metric
        "min_frequency"   : (Integer),          # Data capture frequency in
↪seconds (minimum)
        "max_frequency"   : (Integer)          # Data capture frequency in
↪seconds (maximum)
    }
}'
'report parameters': '{
    "status"              : (String),          # active, inactive, completed,
↪ or cancelled
    "report_specifier_id" : "telemetry",       # ID of the report definition
    "report_request_id"   : (String),          # ID of the report request;
↪supplied by the VTN
    "request_id"          : (String),          # Request ID of the most
↪recent VTN report modification
    "interval_secs"       : (Integer),         # How often a report update
↪is sent to the VTN
    "granularity_secs"    : (Integer),         # How often a report update
↪is sent to the VTN
    "start_time"          : (DateTime - UTC),  # When the report started
    "end_time"            : (DateTime - UTC),  # When the report is
↪scheduled to end
    "last_report"         : (DateTime - UTC),  # When a report update was
↪last sent
    "created_on"          : (DateTime - UTC)   # When this set of
↪information was recorded in the VEN db
}',
'manual_override'       : (Boolean)          # VEN manual override status,
↪as supplied by Control Agent
'online'                 : (Boolean)         # VEN online status, as
↪supplied by Control Agent
}

```

Telemetry value definitions such as `baseline_power_kw` and `current_power_kw` come from the VEN agent config.

RPC Calls

`respond_to_event()`

```

@RPC.export
def respond_to_event(self, event_id, opt_in=True):
    """
        Respond to an event, opting in or opting out.

        If an event's status=unresponded, it is awaiting this call.
        When this RPC is received, the VEN sends an eventResponse to
        the VTN, indicating whether optIn or optOut has been chosen.
        If an event remains unresponded for a set period of time,
    """

```

(continues on next page)

(continued from previous page)

```
it times out and automatically opts in to the event.
```

```
Since this call causes a change in the event's status, it triggers
a PubSub call for the event update, as described above.
```

```
@param event_id: (String) ID of an event.
@param opt_type: (Boolean) Whether to opt in to the event (default True).
"""
```

get_events()

```
@RPC.export
def get_events(self, active_only=True, started_after=None, end_time_before=None):
    """
        Return a list of events.

        By default, return only event requests with status=active or
        ↪ status=unresponded.

        If an event's status=active, a DR event is currently in progress.

        @param active_only: (Boolean) Default True.
        @param started_after: (DateTime) Default None.
        @param end_time_before: (DateTime) Default None.
        @return: (JSON) A list of events -- see 'PubSub: event update'.
    """
```

get_telemetry_parameters()

```
@RPC.export
def get_telemetry_parameters(self):
    """
        Return the VEN's current set of telemetry parameters.

        @return: (JSON) Current telemetry parameters -- see 'PubSub: telemetry parameters
        ↪ update'.
    """
```

set_telemetry_status()

```
@RPC.export
def set_telemetry_status(self, online, manual_override):
    """
        Update the VEN's reporting status.

        @param online: (Boolean) Whether the VEN's resource is online.
        @param manual_override: (Boolean) Whether resource control has been overridden.
    """
```

report_telemetry()

```
@RPC.export
def report_telemetry(self, telemetry_values):
    """
        Update the VEN's report metrics.
```

(continues on next page)

(continued from previous page)

```
Examples of telemetry_values are:
```

```
{
    'baseline_power_kw': '6.2',
    'current_power_kw': '6.145',
    'start_time': '2017-12-05 16:11:42.977298+00:00',
    'end_time': '2017-12-05 16:12:12.977298+00:00'
}
```

```
@param telemetry_values: (JSON) Current value of each report metric.
"""
```

OpenADR VTN Server: Installation and Configuration

The Kisensum VTN server is a Django application written in Python 3 and utilizing a Postgres database.

Get Source Code

To install the VTN server, first get the code by cloning `voltttron-applications` from github and checking out the `openadr` software.

```
$ cd ~/repos
$ git clone https://github.com/voltttron/voltttron-applications
$ cd voltttron-applications
$ git checkout master
```

Install Python 3

After installing Python3 on the server, configure an `openadr` virtual environment:

```
$ sudo pip install virtualenvwrapper
$ mkdir ~/.virtualenvs (if it doesn't exist already)
```

Edit `~/.bashrc` and add these lines:

```
export WORKON_HOME=$HOME/.virtualenvs
export PROJECT_HOME=$HOME/repos/voltttron-applications/kisensum/openadr
source virtualenvwrapper.sh
```

Create the `openadr` project's virtual environment:

```
$ source ~/.bashrc
$ mkvirtualenv -p /usr/bin/python3 openadr
$ setvirtualenvproject openadr ~/repos/voltttron-applications/kisensum/openadr
$ workon openadr
```

From this point on, use **workon openadr** to operate within the `openadr` virtual environment.

Create a local site override for Django's base settings file as follows. First, create `~/.virtualenvs/openadr/settings` in a text editor, adding the following line to it:

```
openadr.settings.site
```

Then, edit `~/.virtualenvs/openadr/postactivate`, adding the following lines:

```
PROJECT_PATH=`cat "$VIRTUAL_ENV/$VIRTUALENVWRAPPER_PROJECT_FILENAME"`
PROJECT_ROOT=`dirname $PROJECT_PATH`
PROJECT_NAME=`basename $PROJECT_PATH`
SETTINGS_FILENAME=".settings"
ENV_FILENAME=".env_postactivate.sh"

# Load the default DJANGO_SETTINGS_MODULE from a .settings
# file in the django project root directory.
export OLD_DJANGO_SETTINGS_MODULE=$DJANGO_SETTINGS_MODULE
if [ -f $VIRTUAL_ENV/$SETTINGS_FILENAME ]; then
    export DJANGO_SETTINGS_MODULE=`cat "$VIRTUAL_ENV/$SETTINGS_FILENAME"`
fi
```

Finally, create `$PROJECT_HOME/openadr/openadr/openadr/settings/site.py`, which holds overrides to `base.py`, the Django base settings file. At a minimum, this file should contain the following:

```
from .base import *
ALLOWED_HOSTS = ['*']
```

A more restrictive `ALLOWED_HOSTS` setting (e.g. `'ki-evi.com'`) should be used in place of `*` if it is known.

Use Pip to Install Third-Party Software

```
$ workon openadr
$ pip install -r requirements.txt
```

Set up a Postgres Database

Install postgres.

Create a postgres user.

Create a postgres database named `openadr`.

(The user name, user password, and database name must match what is in `$PROJECT_HOME/openadr/openadr/settings/base.py` or the override settings in `$PROJECT_HOME/openadr/openadr/settings/local.py`.)

You may have to edit `/etc/postgresql/9.5/main/pg_hba.conf` to be `'md5'` authorization for `'local'`.

Migrate the Database and Create an Initial Superuser

```
$ workon openadr
$ cd openadr
$ python manage.py migrate
$ python manage.py createsuperuser
```

This is the user that will be used to login to the VTN application for the first time, and will be able to create other users and groups.

Configure Rabbitmq

rabbitmq is used by celery, which manages the openadr server's periodic tasks.

Install and run rabbitmq as follows (for further information, see <http://www.rabbitmq.com/download.html>):

```
$ sudo apt-get install rabbitmq-server
```

Start the rabbitmq server if it isn't already running:

```
$ sudo rabbitmq-server -detached (note the single dash)
```

Start the VTN Server

```
$ workon openadr
$ cd openadr
$ python manage.py runserver 0.0.0.0:8000
```

Start Celery

```
$ workon openadr
$ cd openadr
$ celery -A openadr worker -B
```

Configuration Parameters

The VTN supports the following configuration parameters, which can be found in **base.py** and overridden in **site.py**:

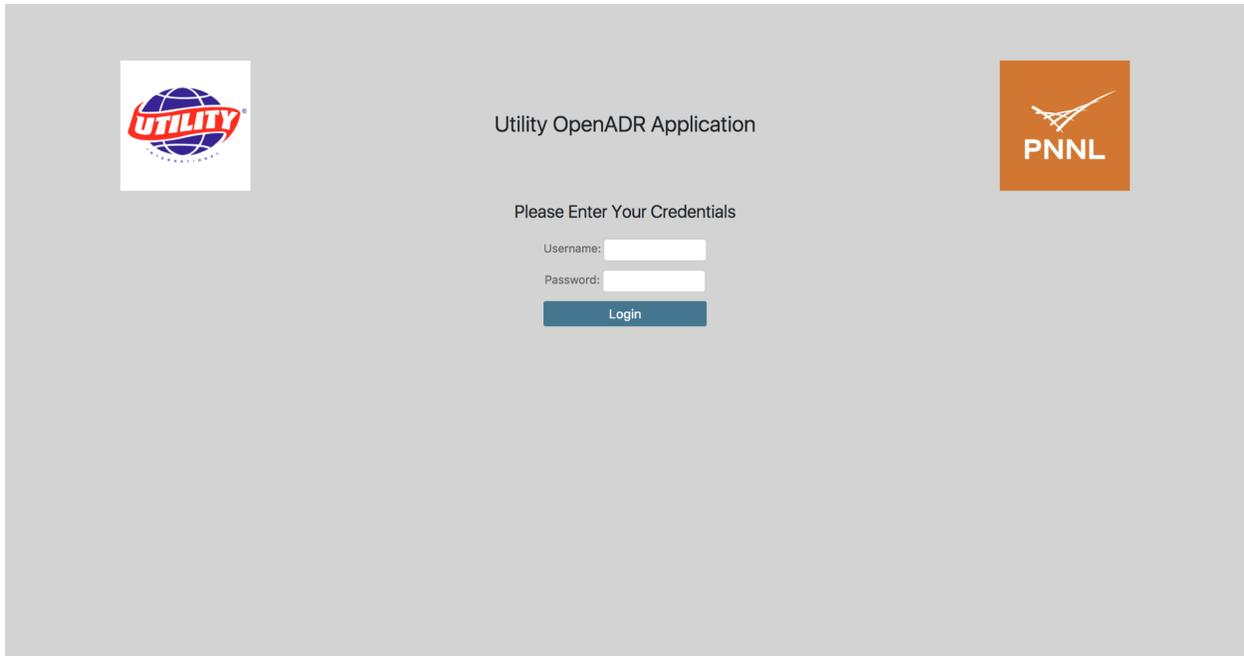
Parameter	Ex-ample	Description
VTN_ID	"vtn01"	OpenADR ID of this virtual top node. Virtual end nodes must know this VTN_ID to be able to communicate with the VTN.
ON-LINE_INTERVAL_MINUTES	15	The amount of time, in minutes, that determines how long the VTN will wait until displaying a given VEN as offline. In other words, if the VTN does not receive any communication from a given VEN within ONLINE_INTERVAL_MINUTES minutes, the VTN will display said VEN as offline.
GRAPH_TIMECHUNK_SECONDS	3600	Displays DR Event graph data by averaging individual VENs' telemetry by GRAPH_TIMECHUNK_SECONDS seconds. This value should be adjusted according to how often VENs are sending the VTN telemetry.

OpenADR VTN Server: User Guide

This guide assumes that you have a valid user account to access and log in to the VTN application website.

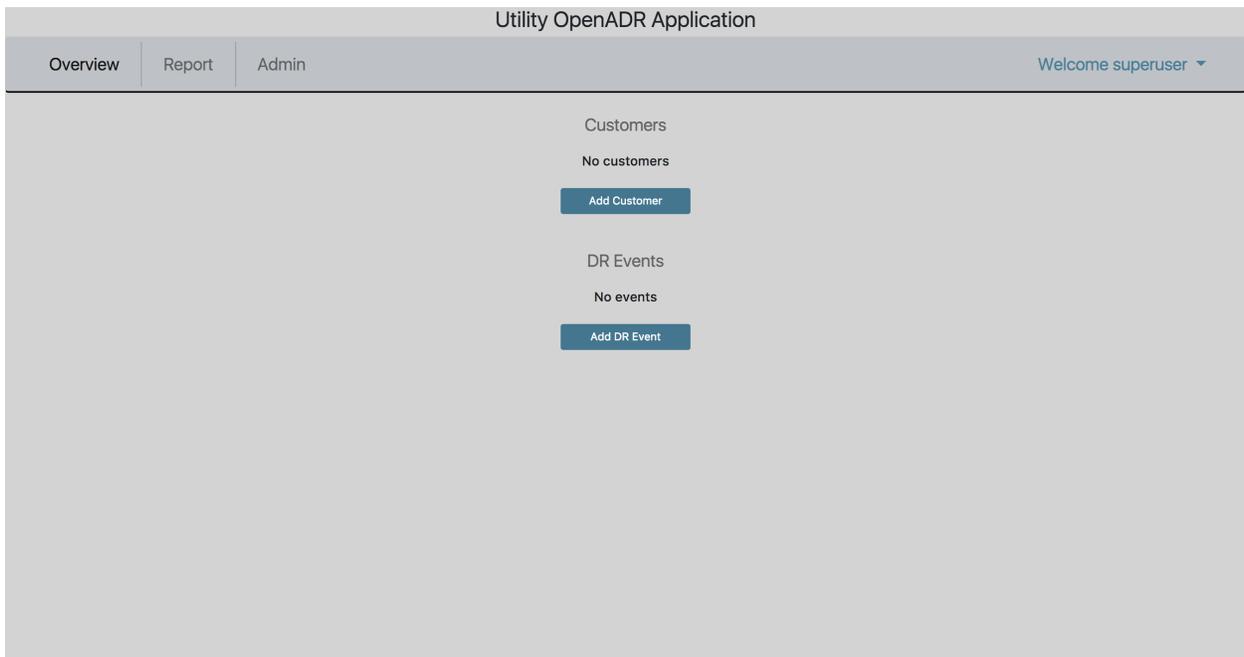
Login Screen

In order to begin using the VTN application, navigate to http://yourhostname*:<or>*ip:8000/vtn.



Overview Screen

Once logged in for the first time, this is the 'Overview' screen.



In order to begin scheduling DR events, one must first create at least one customer, with at least one associated site/VEN, and at least one sort of demand response (DR) program. A VTN will not be able to tell a VEN about DR Events if the VTN doesn't know about the VEN. A VTN knows about a VEN after a Site for the VEN has been created in the VTN application, and the VEN has contacted the VTN.

The rest of this document describes how to set up Customers, Sites, DR Programs, and DR Events, as well as how to export event data.

Create a Customer

Creating a Customer can be done by clicking on ‘Add Customer’ on the **Overview** screen.

The standard interface for adding a Customer:

Utility OpenADR Application

Overview | Report | Admin Welcome superuser ▾

Add Customer

Name

Utility ID

Contact Name

Phone Number

Customers will appear on the **Overview** screen after they have been added.

Utility OpenADR Application

Overview | Report | Admin Welcome superuser ▾

Customers

CUSTOMER	UTILITY ID	CONTACT NAME	PHONE NUMBER	SITES	ONLINE	OFFLINE
Acme Brands	UtilityID02	Joan Graham	2222222222	0	0	0
Best Electronics	UtilityID01	Carl Customer	5555555555	1	0	1
Globex Corp.	UtilityID03	Stanley Martinez	3333333333	0	0	0

DR Events

No events

Create a Site

At first, Customers will not have any Sites. To add a Site for a Customer, click on the Customer's name from the **Overview** screen, and then click 'Create New Site'.

Utility OpenADR Application

Overview Report Admin Welcome superuser ▾

Customer: Globex Corp.

Name
Globex Corp.

Utility ID
UtilityID03

Contact Name
Stanley Martinez

Phone Number
3333333333

Save

No sites

Create New Site

On the **Create Site** screen, DR Programs will appear in the 'DR Programs' multiple-select box if they have been added. This will be discussed soon. Selecting one or more DR Programs here means, when creating a DR Event with a specified DR Program, the site will be an available option for the given DR Event.

A site's 'VEN Name' is permanent. In order to change a Site's VEN Name, the Site must be deleted and re-added.

Utility OpenADR Application

Overview Report Admin Welcome superuser ▾

Add New Site

Customer
Globex Corp.

Name
Headquarters

Site ID
SID04

Location code
LOC04

IP Address
178.121.5.166

Address Line 1
2232 Allen Meadows

Address Line 2

City
East Andrea

State
WA

Zip
83431

DR Programs
Peak Day Pricing
Capacity Bidding
Direct Load Control

VEN Name
ven03

Contact Name
Irma Gural

Phone Number
9999999999

Hold down "Control", or "Command" on a Mac, to select more than one.

Save New Site

After creating a Site for a given customer, the Site will appear offline until communication has been established with

the Site's VEN within a configurable interval (default is 15 minutes).

Utility OpenADR Application

Overview | Report | Admin Welcome superuser ▾

Customer: Globex Corp.

Name

Utility ID

Contact Name

Phone Number

[Save](#)

Sites

SITE	STATUS	CONTACT NAME	PHONE NUMBER
Headquarters	Offline	Irma Gural	9999999999

[Create New Site](#)

Note: When editing a Site, you will notice an extra field on the Screen labeled 'VEN ID'. This field is assigned automatically upon creation of a Site and is used by the VTN to communicate with and identify the VEN.

Utility OpenADR Application

Overview | Report | Admin Welcome superuser ▾

Headquarters

Customer

Name

ID

Location code

IP Address

DR Programs

Capacity Bidding

Direct Load Control

Peak Day Pricing

Hold down "Control", or "Command" on a Mac, to select more than one.

Address Line 1

Address Line 2

City

State

Zip

VEN Name

VEN ID

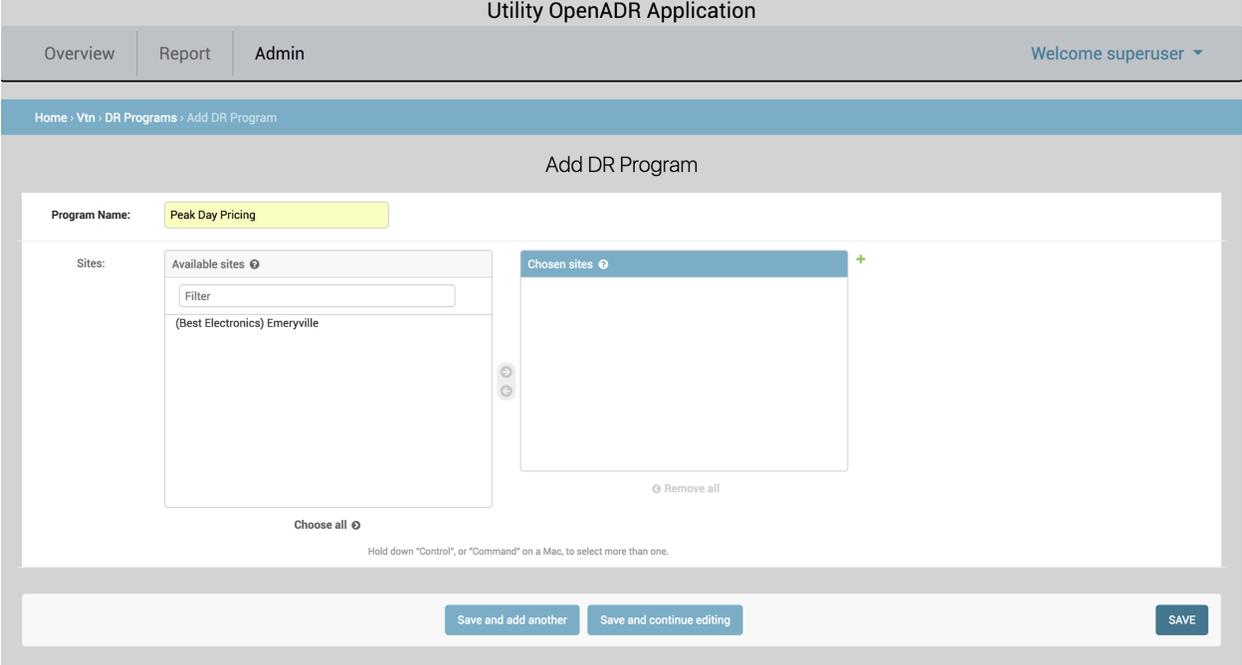
Contact Name

Phone Number

[Save](#)

Create a DR Program

DR Programs must be added via the Admin interface. DR Programs can be added with or without associated sites. In other words, a DR Program can be created with no sites, and sites can be added later, either by Creating/Editing a Site and selecting the DR Program, or by Creating/Editing the DR Program and adding the Site.

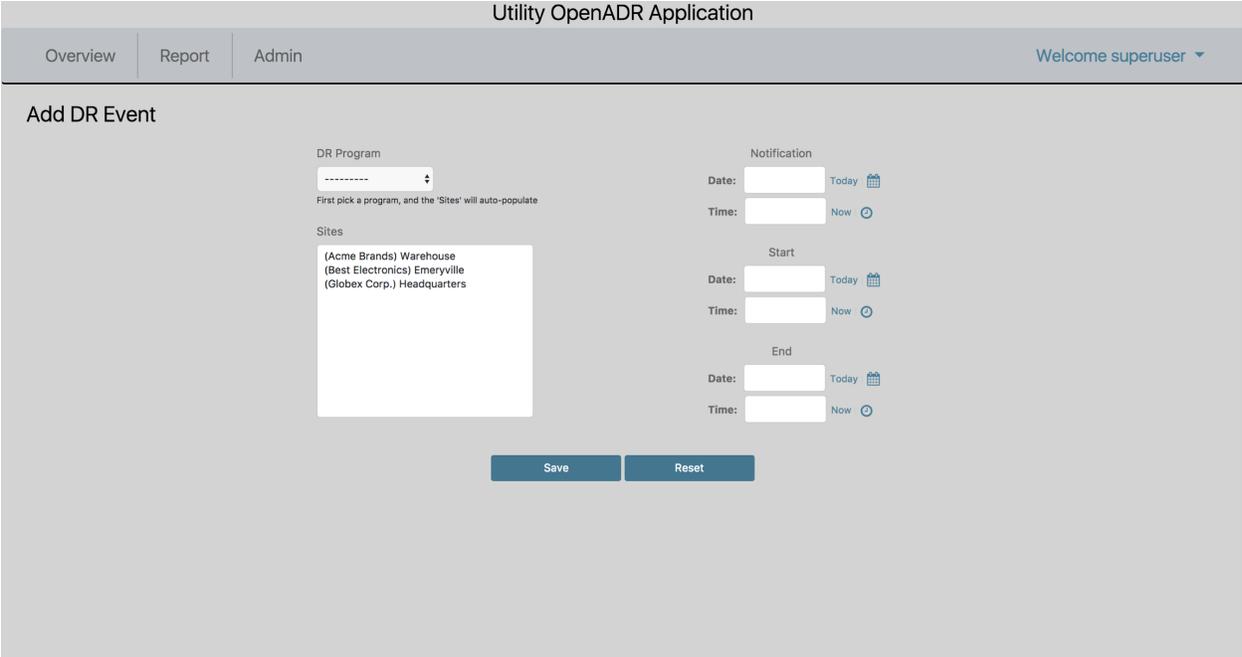


Create a DR Event

Once a Customer (with at least one site) and a DR Program have been added, a DR Event can be created. This is done by navigating to the **Overview** screen and clicking ‘Add DR Event’.

On the **Add DR Event** screen, the first step is to select a DR Program from the drop-down menu. Once a DR Program is selected, the ‘Sites’ multi-select box will auto-populate with the Sites that are associated with that DR Program.

Note that the Notification Time is the absolute soonest time that a VEN will be notified of a DR Event. VENS will not ‘know’ about DR Events that apply to them until they have ‘polled’ the VTN after the Notification Time.



Active DR events are displayed on the **Overview** screen. DR Events are considered active if they have not been cancelled and if they have not been completed.

Utility OpenADR Application

Overview | Report | Admin Welcome superuser ▾

Customers

CUSTOMER	UTILITY ID	CONTACT NAME	PHONE NUMBER	SITES	ONLINE	OFFLINE
Acme Brands	UtilityID02	Joan Graham	2222222222	1	0	1
Best Electronics	UtilityID01	Carl Customer	5555555555	1	0	1
Globex Corp.	UtilityID03	Stanley Martinez	3333333333	0	0	0

[Add Customer](#)

DR Events

PROGRAM	NOTIFICATION	START	END	SITES	STATUS
Peak Day Pricing	Dec. 5, 2017, 7 a.m.	Dec. 6, 2017, 5 p.m.	Dec. 6, 2017, 8 p.m.	2	Scheduled

[Add DR Event](#)

Exporting event telemetry to a .csv is available on the **Report** tab. In the case of this VTN and its associated VENS, the telemetry that will be reported include **baseline power (kw)** and **measured power (kw)**.

Utility OpenADR Application

Overview | **Report** | Admin Welcome superuser ▾

Export Report Data

Date Range: DR Program: All ▾ [Filter](#) [Clear Filter\(s\)](#)

DR PROGRAM	NOTIFICATION	START	END	SITES	ACTION
Peak Day Pricing	Dec. 5, 2017, 7 a.m.	Dec. 6, 2017, 5 p.m.	Dec. 6, 2017, 8 p.m.	2	Export

3.6 Platform Specifications

3.6.1 Weather service specification

Description

The weather service agent will provide API to access current weather data, historical data and weather forecast data. There are several weather data providers, some paid and some free. Weather data providers differs from one and other

1. In the kind of features provided - current data, historical data, forecast data
2. The data points returned
3. The naming schema used to represent the data returned
4. Units of data returned
5. Frequency of data updates

The weather service agent would have a design similar to historians. There would be a single base weather service that defines the api signatures and the ontology of the weather data points. There would be one concrete weather service agents for each weather provider. Users can install one or more provider specific agent to access weather data.

The initial implementations would be for [NOAA](#) and would support current and forecast. NOAA does not support accessing historical weather data through their api.

The next implementation would be for [darksky.net](#) and when implementing this, caching of history data would be added to base weather agent.

Features

Base weather agent features:

1. Caching

The weather service will provide basic caching capability so that repeated request for same data can be returned from cache instead of network round trip to the weather data provider. This is also useful to limit the number of request made to the provider as most weather data provider have restrictions on number of requests for developer/free api keys. The size of the cache can be restricted by setting an optional configuration parameter 'max_size_gb'

2. Name mapping

Data points returned by concrete weather agents would be mapped to standard names based on [CF standard names table](#) Name mapping would be done using a CSV file. See [Configuration](#) section for an example configuration

3. Unit conversion

If data returned from the provider is of the format {"data_point_name":value}, base weather agent can do unit conversions on the value. Both name mapping and unit conversions can be specified as a csv file and packaged with the concrete implementing agent. This file is not mandatory. See [Configuration](#) section for an example configuration

Core weather data retrieval features :

1. Retrieve current weather data.
2. Retrieve hourly weather forecast data.
3. Retrieve historical weather data.
4. Periodic polling of current weather data for one or more locations. Users can configure one or more locations in a config file and weather agent will periodically poll for current weather data for the configured locations and publish the results to message bus.

The set of points returned from the above queries would depend on the specific weather data provider, however the point names returned would be from the standard schema.

Note:

1. Since individual weather data provider can support slightly different sets of features, users should be able to query for the list of available features. For example a provider could provide daily weather forecast in addition to the hourly forecast data.

API

1. Get available features

rpc call to weather service method `'get_api_features'`

Parameters - None

Returns - dictionary of api features that can be called for this weather agent.

2. Get current weather data

rpc call to weather service method `'get_current_weather'`

Parameters:

1. **locations** - dictionary containing location details. The format of location accepted would differ between different weather providers and even different APIs supported by the same provider. For example the location input could be either `{"zipcode":value}` or `{"region":value, "country": value}`.

Returns: List of dictionary objects containing current weather data. The actual data points returned depends on the weather service provider.

3. Get hourly forecast data

rpc call to weather service method `'get_hourly_forecast'`

Parameters:

1. **locations** - dictionary containing location details. The format of location accepted would differ between different weather providers and even different APIs supported by the same provider. For example the location input could be either `{"zipcode":value}` or `{"region":value, "country": value}`.

optional parameters:

2. **hours** - The number of hours for which forecast data should be returned. By default, it is 24 hours.

Returns: List of dictionary objects containing forecast data. If weather data provider returns less than requested number of hours result returned would contain a warning message in addition to the result returned by the provider.

4. Get historical weather data

rpc call to weather service method `'get_hourly_historical'`

Parameters:

1. **locations** - dictionary containing location details. For example the location input could be either {"zip-code":value} or {"region":value, "country": value}.
2. **start_date** - start date of requested data
3. **end_date** - end date of requested data

Returns: List of dictionary objects containing historical data.

Note: Based on the weather data provider this api could do multiple calls to the data provider to get the requested data. For example, darksky.net allows history data query by a single date and not a date range.

5. Periodic polling of current weather data

This can be achieved by configuring the locations for which data is requested in the agent's configuration file along with polling interval. Results for each location configured, is published to its corresponding result topic. If no result topic prefix is configured, then results for all locations are posted to the topic weather/poll/current/all. poll_topic_suffixes when provided should be a list of string with the same length as the number of poll_locations. When topic prefix is specified, each location's result is published to weather/poll/current/<poll_topic_suffix for that location> topic_prefix.

Configuration

Example configuration:

```
{
  poll_locations: [
    {"zip": "22212"},
    {"zip": "99353"}
  ],
  poll_topic_suffixes: ["result_22212", "result_99353"],
  poll_interval: 20 #seconds,

  #optional cache arguments
  max_cache_size: ...
}
```

Example configuration for mapping point names returned by weather provider to a standard name and units:

```
Service_Point_Name, Standard_Point_Name, Service_Units, Standard_Units
temperature, air_temperature, fahrenheit, celsius
```

Caching

Weather agent will cache data until the configured size limit is reached (if provided).

1. Current and forecast data:

If current/forecast weather data exists in cache and if the request time is within the update time period of the api (specified by a concrete implementation) then by default cached data would be returned otherwise a new request is made for it. If hours is provided and the amount of cached data records is less than hours, this will also result in a new request.

2. Historical data cache:

Weather api will query the cache for available data for the given time period and fill and missing time period with data from the remote provider.

3. Clearing of cache:

Users can configure the maximum size limit for cache. For each api call, before data is inserted in cache, weather agent will check for this size limit and purge records in this order. - Current data older than update time period - Forecast data older than update time period - History data starting with the oldest cached data

Assumptions

1. User has api key for accessing weather api for a specific weather data provider, if a key is required.
2. Different weather agent might have different requirement for how input locations are specified. For example NOAA expects a station id for querying current weather and requires either a lat/long or gridpoints to query for forecast. weatherbit.io accepts zip code.
3. Not all features might be implemented by a specific weather agent. For example NOAA doesn't make history data available using their weather api.
4. Concrete agents could expose additional api features
5. Optionally, data returned will be based on standard names provided by the CF standard names table (see Ontology). Any points with a name not mapped to a standard name would be returned as is.

3.6.2 Agent VIP IDENTITY Assignment Specification

This document explains how an agent obtains its VIP IDENTITY, how the platform sets an agent's VIP IDENTITY at startup, and what mechanisms are available to the user to set the VIP IDENTITY for any agent.

What is a VIP IDENTITY

A VIP IDENTITY is a platform instance unique identifier for agents. The IDENTITY is used to route messages from one Agent through the VOLTTRON router to the recipient Agent. The VIP IDENTITY provides a consistent, user defined, and human readable character set to build a VIP IDENTITY. VIP IDENTITIES should be composed of both upper and lowercase letters, numbers and the following special characters `_.-`.

Runtime

The primary interface for obtaining a VIP IDENTITY *at runtime* is via the runtime environment of the agent. At startup the utility function `vip_main` shall check for the environment variable `AGENT_VIP_IDENTITY`. If the `AGENT_VIP_IDENTITY` environment variable is not set then the `vip_main` function will fall back to a supplied identity argument. `vip_main` will pass the appropriate identity argument to the agent constructor. If no identity is set the Agent class will create a random VIP IDENTITY using python's `uuid4` function.

An agent that inherits from the platform's base Agent class can get its current VIP IDENTITY by retrieving the value of `self.core.identity`.

The primary use of the 'identity' argument to `vip_main` is for agent development. For development it allows agents to specify a default VIP IDENTITY when run outside the platform. As platform Agents are not started via `vip_main` they will simply receive their VIP IDENTITY via the identity argument when they are instantiated. Using the identity argument of the Agent constructor to set the VIP IDENTITY via agent configuration is no longer supported.

At runtime the platform will set the environment variable `AGENT_VIP_IDENTITY` to the value set at installation time.

Agents not based on the platform's base Agent should set their VIP IDENTITY by setting the identity of the ZMQ socket before the socket connects to the platform. If the agent fails to set its VIP IDENTITY via the ZMQ socket it will be selected automatically by the platform. This platform chosen ID is currently not discoverable to the agent.

Agent Implementation

If an Agent has a preferred VIP IDENTITY (for example the MasterDriverAgent prefers to use "platform.driver") it may specify this as a default packed value. This is done by including a file named IDENTITY containing only the desired VIP IDENTITY in ASCII plain text in the same directory at the setup.py file for the Agent. This will cause the packaged agent wheel to include an instruction to set the VIP IDENTITY at installation time.

This value may be overridden at packaging or installation time.

Packaging

An Agent may have its VIP IDENTITY configured when it is packaged. The packaged value may be used by the platform to set the `AGENT_VIP_IDENTITY` environment variable for the agent process.

The packaged VIP IDENTITY may be overridden at installation time. This overrides any preferred VIP IDENTITY of the agent. This will cause the packaged agent wheel to include an instruction to set the VIP IDENTITY at installation time.

To specify the VIP IDENTITY when packaging use the `-vip-identity` option when running "voltron-pkg package".

Installation

An agent may have its VIP IDENTITY configured when it is installed. This overrides any VIP IDENTITY specified when the agent was packaged.

To specify the VIP IDENTITY when packaging use the `-vip-identity` option when running "voltron-ctl install".

Installation Default VIP IDENTITY

If no VIP IDENTITY has been specified by installation time the platform will assign one automatically.

The platform uses the following template to generate a VIP IDENTITY:

```
"{agent_name}_{n}"
```

{agent_name} is substituted with the name of the actual agent such as "listeneragent-0.1"

{n} is a number to make VIP IDENTITY unique. {n} is set to the first unused number (starting from 1) for all installed instances of an agent. e.g. If there are 2 listener agents installed and the first (VIP IDENTITY listeneragent-0.1_1) is uninstalled leaving the second (VIP IDENTITY "listeneragent-0.1_2") a new listener agent will receive the VIP IDENTITY "listeneragent-0.1_1" when installed. The next installed listener will receive a VIP IDENTITY of "listeneragent-0.1_3".

The # sign is used to prevent confusing the agent version number with the installed instance number.

If an agent is repackaged with a new version number it is treated as a new agent and the number will start again from 1.

VIP IDENTITY Conflicts During Installation

If an agent is assigned a VIP IDENTITY besides the default value given to it by the platform it is possible for VIP IDENTITY conflicts to exist between installed agents. In this case the platform rejects the installation of an agent with a conflicting VIP IDENTITY and reports an error to the user.

VIP IDENTITY Conflicts During Runtime

In the case where agents are not started through the platform (usually during development or when running standalone agents) it is possible to encounter a VIP IDENTITY conflict during runtime. In this case the first agent to use a VIP IDENTITY will function as normal. Subsequent agents will still connect to the ZMQ socket but will be silently rejected by the platform router. The router will not route any message to that Agent. Agents using the platform's base Agent class will detect this automatically during the initial handshake with the platform. This condition will shutdown the Agent with an error indicating a VIP IDENTITY conflict as the most likely cause of the problem.

Auto Numbering With Non-Default VIP IDENTITIES

It is possible to use the auto numbering mechanism that the default VIP IDENTITY scheme uses. Simply include the string "{n}" somewhere in the requested VIP IDENTITY and it will be replaced with a number in the same manner as the default VIP IDENTITY is. Python string.format() escaping rules apply. [See this question on StackOverflow.](#)

Script Features

Both the make-agent.sh and pack_install.sh scripts support reading the desired VIP IDENTITY from the AGENT_VIP_IDENTITY environment variable.

Security/Privacy

Currently, much like the TAG file in an installed agent, there is nothing to stop someone from modifying the IDENTITY file in the installed agent.

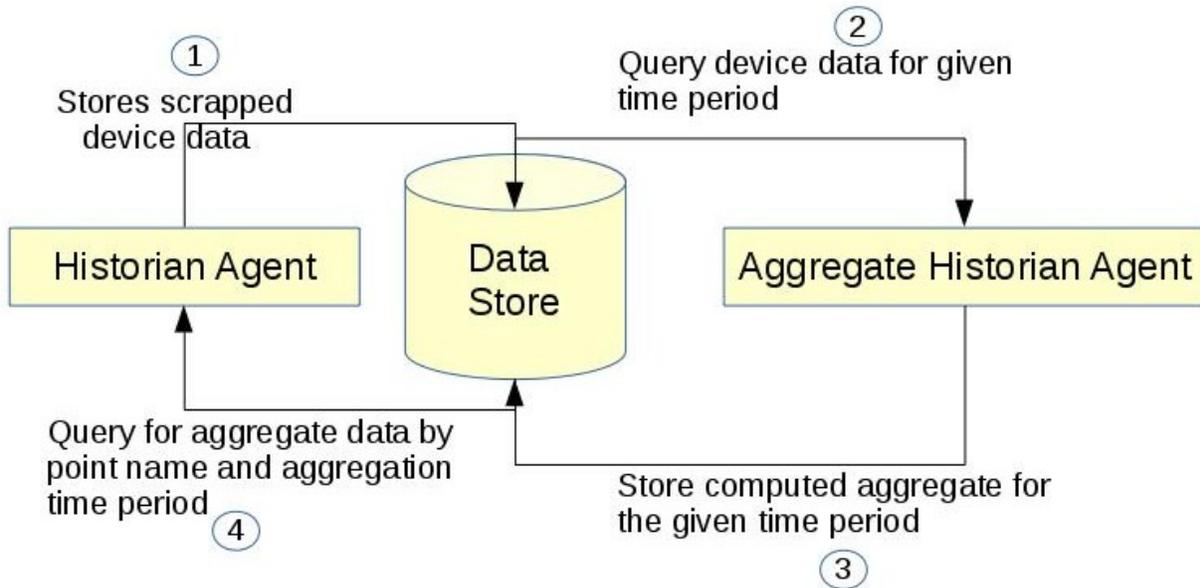
Constraints and Limitations

Currently there is no way for an agent based on the platform base Agent class to recover from a VIP IDENTITY conflict. As that case only affects developers and a very tiny minority of users and is reported via an error message, there are no plans to fix it.

3.6.3 Aggregate Historian Agent Specification

Description

An aggregate historian computes aggregates of data stored in a given volttron historian's data store. It runs periodically to compute aggregate data and store it in new tables/collections in the historian's data store. Each regular historian ([BaseHistorian](#)) needs a corresponding aggregate historian to compute and store aggregates of the data collected by the regular historian.



Software Interfaces

Data Collection - Data store that the aggregate historian uses as input source needs to be up. Access to it should be provided using an account that has create, read, and write privileges. For example, a MongoAggregateHistorian needs to be able to connect to the mongodb used by MongoHistorian using an account that has read and write access to the db used by the MongoHistorian.

Data retrieval Aggregate Historian Agent does not provide api for retrieving the aggregate data collected. Use Historian agent's query interface. Historian's query api will be modified as below

1. topic_name can now be a list of topic names or a single topic
2. Two near optional parameters have been added to the query api - agg_type (aggregation type), agg_period (aggregation time period). Both these parameters are mandatory for query aggregate data.
3. New api to get the list of aggregate topics available for querying

User Interfaces

Aggregation agent requires user to configure the following details as part of the agent configuration file

1. Connection details for historian's data store (same as historian agent configuration)
2. **List of aggregation groups where each group contains:**
 1. Aggregation period - integer followed by m/h/d/w/M (minutes, hours, days, weeks or months)
 2. Boolean parameter to indicate if aggregation periods should align to calendar times
 3. Optional collection start time in utc. If not provided, aggregation collection will start from current time
 4. List of aggregation points with topic name, type of aggregation (sum, avg, etc.), and minimum number of records that should be available for the aggregate to be computed
 5. Topic name can be specified either as a list of specific topic names (topic_names=[topic1, topic2]) or a regular expression pattern (topic_name_pattern="Building1/device_*/Zone*temperature")

6. When aggregation is done for a single topic then name of topic will be used for the computed aggregation as well. You could optionally provide a unique aggregation_topic_name
7. When topic_name_pattern or multiple topics are specified a unique aggregate topic name should be specified for the collected aggregate. Users can query for the collected aggregate data using this aggregate topic name.
8. User should be able to configure multiple aggregations done with the same time period/time interval and these should be time synchronized.

Functional Capabilities

1. Should run periodically to compute aggregate data.
2. Same instance of the agent should be able to collect data at more than one time interval
3. For each configured time period/interval agent should be able to collect different type of aggregation for different topics/points
4. Support aggregation over multiple topics/points
5. Agent should be able to handle and normalize different time units such as minutes, hours, days, weeks and months
6. Agent should be able to compute aggregate both based on wall clock based time intervals and calendar based time interval. For example, agent should be able to calculate daily average based on 12.00AM to 11.59PM of a calendar day or between current time and the same time the previous day.
7. Data should be stored in such a way that users can easily retrieve multiple aggregate topics data within a given time interval

Data Structure

Collected aggregate data should be stored in the historian data store into new collection or tables and should be accessible by historian agent's query interface. Users should easily be able to query aggregate data of multiple points for which data is time synchronized.

Use Cases

Collect monthly average of multiple topic using data from MongoDBHistorian

1. Create a configuration file with connection details from Mongo Historian configuration file and add additional aggregation specific configuration

```
{
  # configuration from mongo historian - START
  "connection": {
    "type": "mongodb",
    "params": {
      "host": "localhost",
      "port": 27017,
      "database": "mongo_test",
      "user": "test",
      "passwd": "test"
    }
  },
}
```

(continues on next page)

(continued from previous page)

```

# configuration from mongo historian - START
"aggregations":[
  # list of aggregation groups each with unique aggregation_period and
  # list of points that needs to be collected
  {
    "aggregation_period": "1M",
    "use_calendar_time_periods": true,
    "utc_collection_start_time":"2016-03-01T01:15:01.000000",
    "points": [
      {
        "topic_names": ["Building/device/point1", "Building/device/point2"],
        "aggregation_topic_name":"building/device/point1_2/month_sum",
        "aggregation_type": "avg",
        "min_count": 2
      }
    ]
  }
]
}

```

In the above example configuration, here is what each field under “aggregations” represent

- **aggregation_period:** can be minutes(m), hours(h), weeks(w), or months(M)
- **use_calendar_time_periods: true or false - Should aggregation period align to calendar time periods. Default False. Exam**
 - if “aggregation_period”:"1h" and “use_calendar_time_periods”: false, example periods: 10.15-11.15, 11.15-12.15, 12.15-13.15 etc.
 - if “aggregation_period”:"1h" and “use_calendar_time_periods”: true, example periods: 10.00-11.00, 11.00-12.00, 12.00-13.00 etc.
 - if “aggregation_period”:"1M" and “use_calendar_time_periods”: true, aggregation would be computed from the first day of the month to last day of the month
 - if “aggregation_period”:"1M" and “use_calendar_time_periods”: false, aggregation would be computed with a 30 day interval based on aggregation collection start time
- **utc_collection_start_time:** The time from which aggregation computation should start. If not provided this would default to current time.
- **points: List of points, its aggregation type and min_count topic_names:** List of topic_names across which aggregation should be computed. **aggregation_topic_name:** Unique name given for this aggregate. Optional if aggregation is for a single topic. **aggregation_type:** Type of aggregation to be done. Please see [Constraints and Limitations](#)
 - min_count:** Optional. Minimum number of records that should exist within the configured time period for a aggregation to be computed.

2. install and starts the aggregate historian using the above configuration

3. Query aggregate data: Query using historian’s query api by passing two additional parameters - agg_type and agg_period

```

result1 = query_agent.vip.rpc.call('platform.historian',
                                  'query',
                                  topic='building/device/point1_2/month_sum',
                                  agg_type='avg',
                                  agg_period='1M',

```

(continues on next page)

(continued from previous page)

```
count=20,
order="FIRST_TO_LAST").get(10)
```

Collect weekly average(sunday to saturday) of single topic using data from MongoDBHistorian

1. Create a configuration file with connection details from Mongo Historian configuration file and add additional aggregation

- **aggregation_period**: "1w",
- **topic_names**: ["Building/device/point1"], #topic for which you want to compute aggregation
- **aggregation_topic_name** need not be provided

2. install and starts the aggregate historian using the above configuration

3. Query aggregate data: Query using historian's query api by passing two additional parameters - **agg_type** and **agg_period**. **topic_name** will be the same as the point name for which aggregation is collected

```
result1 = query_agent.vip.rpc.call('platform.historian',
                                  'query',
                                  topic='Building/device/point1',
                                  agg_type='avg',
                                  agg_period='1w',
                                  count=20,
                                  order="FIRST_TO_LAST").get(10)
```

Collect hourly average for multiple topics based on topic_name pattern

1. Create a configuration file with connection details from Mongo Historian configuration file and add additional aggregation

- **aggregation_period**: "1h",
- Insetead of **topic_names** provide **topic_name_pattern**. For example, **"topic_name_pattern": "Building1/device_a*/point1"**
- **aggregation_topic_name** provide a unique aggregation topic name

2. install and starts the aggregate historian using the above configuration

3. Query aggregate data: Query using historian's query api by passing two additional parameters - **agg_type** and **agg_period**. **topic_name** will be the same as the point name for which aggregation is collected

```
result1 = query_agent.vip.rpc.call('platform.historian',
                                  'query',
                                  topic="unique aggregation_topic_name provided in_
↳configuration",
                                  agg_type='avg',
                                  agg_period='1h',
                                  count=20,
                                  order="FIRST_TO_LAST").get(10)
```

Collect 7 day average of two topics and time synchronize them for easy comparison

1. Create a configuration file with connection details from Mongo Historian configuration file and add additional aggregation specific configuration. The configuration file should be similar to the below example

```
{
  # configuration from mongo historian - START
  "connection": {
    "type": "mongodb",
    "params": {
      "host": "localhost",
      "port": 27017,
      "database": "mongo_test",
      "user": "test",
      "passwd": "test"
    }
  },
  # configuration from mongo historian - START
  "aggregations": [
    # list of aggregation groups each with unique aggregation_period and
    # list of points that needs to be collected
    {
      "aggregation_period": "1w",
      "use_calendar_time_periods": false, #compute for last 7 days, then the next_
      ↪and so on..
      "points": [
        {
          "topic_names": ["Building/device/point1"],
          "aggregation_type": "avg",
          "min_count": 2
        },
        {
          "topic_names": ["Building/device/point2"],
          "aggregation_type": "avg",
          "min_count": 2
        }
      ]
    }
  ]
}
```

2. install and starts the aggregate historian using the above configuration

3. Query aggregate data: Query using historian's query api by passing two additional parameters - `agg_type` and `agg_period`. provide the list of topic names for which aggregate was configured above. Since both the points were configured within a single "aggregations" array element, their aggregations will be time synchronized

```
result1 = query_agent.vip.rpc.call('platform.historian',
                                   'query',
                                   topic=['Building/device/point1' 'Building/device/
      ↪point2'],
                                   agg_type='avg',
                                   agg_period='1w',
                                   count=20,
                                   order="FIRST_TO_LAST").get(10)
```

Results will be of the format

```
{'values': [
  ['Building/device/point1', '2016-09-06T23:31:27.679910+00:00', 2],
  ['Building/device/point1', '2016-09-15T23:31:27.679910+00:00', 3],
  ['Building/device/point2', '2016-09-06T23:31:27.679910+00:00', 2],
  ['Building/device/point2', '2016-09-15T23:31:27.679910+00:00', 3]],
'metadata': {}}
```

Qurey list of aggregate data collected

```
result = query_agent.vip.rpc.call('platform.historian',
                                  'get_aggregate_topics').get(10)
```

The result will be of the format:

```
[(aggregate topic name, aggregation type, aggregation time period, configured list of ↵
↵topics or topic name pattern), ...]
```

This shows the list of aggregation currently being computed periodically

Qurey list of supported aggregation types

```
result = query_agent.vip.rpc.call(
    AGG_AGENT_VIP,
    'get_supported_aggregations').get(timeout=10)
```

Constraints and Limitations

1. Initial implementation of this agent will not support any data filtering for raw data before computing data aggregation
2. Initial implementation should support all aggregation types directly supported by underlying data store. End user input is needed to figure out what additional aggregation methods are to be supported

MySQL

Name	Description
AVG()	Return the average value of the argument
BIT_AND()	Return bitwise AND
BIT_OR()	Return bitwise OR
BIT_XOR()	Return bitwise XOR
COUNT()	Return a count of the number of rows returned
GROUP_CONCAT()	Return a concatenated string
MAX()	Return the maximum value
MIN()	Return the minimum value
STD()	Return the population standard deviation
STDDEV()	Return the population standard deviation
STDDEV_POP()	Return the population standard deviation
STDDEV_SAMP()	Return the sample standard deviation
SUM()	Return the sum
VAR_POP()	Return the population standard variance
VAR_SAMP()	Return the sample variance
VARIANCE()	Return the population standard variance

SQLite

Name	Description
AVG()	Return the average value of the argument
COUNT()	Return a count of the number of rows returned
GROUP_CONCAT()	Return a concatenated string
MAX()	Return the maximum value
MIN()	Return the minimum value
SUM()	Return sum of all non-NULL values in the group. If there are no non-NULL input rows then returns NULL .
TOTAL()	Return sum of all non-NULL values in the group.If there are no non-NULL input rows returns 0.0

MongoDB

Name	Description
SUM	Returns a sum of numerical values. Ignores non-numeric values
AVG	Returns a average of numerical values. Ignores non-numeric values
MAX	Returns the highest expression value for each group.
MIN	Returns the lowest expression value for each group.
FIRST	Returns a value from the first document for each group. Order is only defined if the documents are in a defined order.
LAST	Returns a value from the last document for each group. Order is only defined if the documents are in a defined order.
PUSH	Returns an array of expression values for each group
AD-DTOSET	Returns an array of unique expression values for each group. Order of the array elements is undefined.
STDDEV-VPOP	Returns the population standard deviation of the input values
STDDEV-VSAMP	Returns the sample standard deviation of the input values

3.6.4 Chargepoint API Driver

Spec Version 1.1

ChargePoint operates the largest independently owned EV charging network in the US. It sells charge stations to businesses and provides a web application to manage and report on these chargestations. Chargepoint offers a [Web Services API](#) that its customers may use to develop applications that integrate with the chargepoint network devices.

The Chargepoint API Driver for VOLTRON will enable real-time monitoring and control of Chargepoint EVSEs within the VOLTRON platform by creating a standard VOLTRON device driver on top of the Chargepoint Web Services API. Each port on each managed chargestation will look like a standard VOLTRON device, monitored and controlled through the VOLTRON device driver interface.

Driver Scope & Functions

This driver will enable VOLTRON to support the following use cases with Chargepoint EVSEs:

- Monitoring of chargestation status, load and energy consumption
- Demand charge reduction
- Time shifted charging
- Demand response program participation

The data and functionality to be made available through the driver interface will be implemented using the following Chargepoint web services:

API Method Name	Key Data/Function Provided
getStationStatus	Port status: AVAILABLE, INUSE, UNREACHABLE, UNKNOWN
shedLoad	Limit station power by percent or max load for some time period.
clearShedState	Clear all shed state and allow normal charging
getLoad	Port load in Kw, shedState, allowedLoad, percentShed
getAlarms	Only the last alarm will be available.
clearAlarms	Clear all alarms.
getStationRights	Name of station rights profile, eg. 'network_manager'
getChargingSessionData	Energy used in last session, start/end timestamps
getStations	Returns description/address/nameplate of chargestation.

The Chargepoint Driver will implement version 5.0 Rev 7 of the Chargepoint API. While the developer's guide is not yet publicly available, the WSDL Schema is. *Note: Station Reservation API has been removed from the 5.0 version of the API.*

WSDL for this API is located here:

https://webservices.chargepoint.com/cp_api_5.0.wsdl

Mapping VOLTRON Device Interface to Chargepoint APIs

The VOLTRON driver interface represents a single device as a list of registers accessed through a simple get_point/set_point API. In contrast, the Chargepoint web services for real-time monitoring and control are spread across eight distinct APIs that return hierarchical XML. The Chargepoint driver is the adaptor that will make a suite of web services look like a single VOLTRON device.

Device Mapping

The chargepoint driver will map a single VOLTRON device (a driver instance) to one chargestation. Since a chargestation can have multiple ports, each with their own set of telemetry, the registry will include a port index column on attributes that are specific to a port. This will allow deployments to use an indexing convention that has been followed with other drivers. (See Registry Configuration for more details)

Driver Configuration

Each device must be configured with its own Driver Configuration File. The Driver Configuration must reference the Registry Configuration File, defining the set of points that will be available from the device. For chargestation devices, the `driver_config` entry of the Driver Configuration file will need to contain all parameters required by the web service API:

Parameter	Purpose
username	Credentials established through Chargepoint account
password	
stationID	Unique station ID assigned by chargepoint

The `driver_type` must be `chargepoint`

A sample driver configuration file for a single device, looks like this:

```
{
  "driver_config": {
    "username" : "1b905c936af141b98f9b0f816087f3605a30c1df1d07f146281b151",
    "password" : "**Put your chargepoint API passqword here**",
    "stationID" : "1:34003",
  },
  "driver_type": "chargepoint",
  "registry_config": "config://chargepoint.csv",
  "interval": 60,
  "heart_beat_point": "heartbeat"
}
```

API Plans & Access Rights

Chargepoint offers API plans that vary in available features and access rights. Some of the API calls to be implemented here are not available across all plans. Furthermore, the attributes returned in response to an API call may be limited by the API plan and access rights associated with the `userid`. Runtime exceptions related to plans and access rights will generate `DriverInterfaceError` exceptions. These can be avoided by using a registry configuration that does not include APIs or attributes that are not available to the `<username>`.

Registry Configuration

The registry file defines the individual points that will be exposed by the Chargepoint driver. It should only reference points that will actually be used since each point is potentially an additional web service call. The driver will be smart and limit API calls to those that are required to satisfy the points found in the CSV.

Naming of points will conform to the conventions established by the Chargepoint Web services API whenever possible. Note that Chargepoint naming conventions are camel-cased with no spaces or hyphens. Multi-word names start with a lowercase letter. Single word names start uppercase.

The available registry entries for each API method name are shown below along with a description of any notable behavior associated with that register. Following that is a sample of the associated XML returned by the API.

getStationStatus

The getStationStatus query returns information for all ports on the chargestation.

Note: In all the registry entries shown below, the **Attribute Name** column defines the unique name within the chargepoint driver that must be used to reference this particular attribute and associated API. The **VOLTTRON point name** usually matches the **Attribute Name** in these examples but may be changed during deployment.

Table 22: getStationStatus

Volttron Point Name	Attribute Name	Register Name	Port #	Type	Units	Starting Value	Writable	Notes
Status	Status	Station-Status-Register	1	string			FALSE	AVAILABLE, INUSE, UNREACHABLE, UNKNOWN
Status.Timestamp	TimeStamp	Station-Status-Register	1	date-time			FALSE	Timestamp of the last communication between the station and Charge-Point

Sample XML returned by getStationStatus.

```
<ns1:getStationStatusResponse xmlns:ns1="urn:dictionary:com.chargepoint.webservices">
  <responseCode>100</responseCode>
  <responseText>API input request executed successfully.</responseText>
  <stationData>
    <stationID>1:33923</stationID>
    <Port>
      <portNumber>1</portNumber>
      <Status>AVAILABLE</Status>
      <TimeStamp>2016-11-07T19:19:19Z</TimeStamp>
    </Port>
    <Port>
      <portNumber>2</portNumber>
      <Status>INUSE</Status>
      <TimeStamp>2016-11-07T19:19:19Z</TimeStamp>
    </Port>
  </stationData>
  <moreFlag>0</moreFlag>
</ns1:getStationStatusResponse>
```

getLoad, shedLoad, clearShedState

Reading any of these values will return the result of a call to getLoad. Writing shedState=True will call shedLoad and pass the last written value of allowedLoad or percentShed. The API allows only one of these two values to be provided. Writing to allowedLoad will simultaneously set percentShed to None and vice versa.

Table 23: getLoad, shedLoad, clearShedState

Volttron Point Name	Attribute Name	Register Name	Port #	Type	Units	Starting Value	Writable	Notes
shedState	shedState	Load-Register	1	integer	0 or 1	0	TRUE	True when load shed limits are in place
portLoad	portLoad	Load-Register	1	float	kw		FALSE	Load in kw
allowedLoad	allowedLoad	Load-Register	1	float	kw		TRUE	Allowed load in kw when shedState is True
percentShed	percentShed	Load-Register	1	integer	percent		TRUE	Percent of max power shed when shedState is True

Sample XML returned by getLoad

```
<ns1:getLoadResponse xmlns:ns1="urn:dictionary:com.chargepoint.webservices">
  <responseCode>100</responseCode>
  <responseText>API input request executed successfully.</responseText>
  <numStations></numStations>
  <groupName></groupName>
  <sgLoad></sgLoad>
  <stationData>
    <stationID>1:33923</stationID>
    <stationName>ALCOGARSTATIONS / ALCOPARK 8 -005</stationName><Address>165 13th
↪St, Oakland, California, 94612, United States</Address>
    <stationLoad>3.314</stationLoad>
    <Port>
      <portNumber>1</portNumber>
      <userID></userID>
      <credentialID></credentialID>
      <shedState>0</shedState>
      <portLoad>0.000</portLoad>
      <allowedLoad>0.000</allowedLoad>
      <percentShed>0</percentShed>
    </Port>
    <Port>
      <portNumber>2</portNumber>
      <userID>664719</userID>
      <credentialID>CNCP0000481668</credentialID>
      <shedState>0</shedState>
      <portLoad>3.314</portLoad>
      <allowedLoad>0.000</allowedLoad>
      <percentShed>0</percentShed>
    </Port>
  </stationData>
</ns1:getLoadResponse>
```

Sample shedLoad XML query to set the allowed load on a port to 3.0kw.

```
<ns1:shedLoad>
  <shedQuery>
    <shedStation>
      <stationID>1:123456</stationID>
      <Ports>
```

(continues on next page)

(continued from previous page)

```

    <Port>
      <portNumber>1</portNumber>
      <allowedLoadPerPort>3.0</allowedLoadPerPort>
    </Port>
  </Ports>
</shedStation>
<timeInterval/>
</shedQuery>
</ns1:shedLoad>

```

getAlarms, clearAlarms

The getAlarms query returns a list of all alarms since last cleared. The driver interface will only return data for the most recent alarm, if present. While the getAlarm query provides various station identifying attributes, these will be made available through registers associated with the getStations API. If an alarm is not specific to a particular port, it will be associated with all chargestation ports and available through any of its device instances.

Write True to clearAlarms to submit the clearAlarms query to the **chargestation**. It will clear alarms across all ports on that chargestation.

Table 24: getAlarms, clearAlarms

Volttron Point Name	Attribute Name	Register Name	Port #	Type	Units	Starting Value	Writable	Notes
alarmType	alarm-Type	Alarm-Register		string			FALSE	eg. 'GFCI Trip'
alarmTime	alarm-Time	Alarm-Register		date-time			FALSE	
clearAlarms	clear-Alarms	Alarm-Register		int		0	TRUE	Sends the clearAlarms query when set to True

```

<Alarms>
  <stationID>1:33973</stationID>
  <stationName>ALCOGARSTATIONS / ALCOPARK 8 -003</stationName>
  <stationModel>CT2100-HD-CCR</stationModel>
  <orgID>1:ORG07225</orgID>
  <organizationName>Alameda County</organizationName>
  <stationManufacturer></stationManufacturer>
  <stationSerialNum>115110013418</stationSerialNum>
  <portNumber></portNumber>
  <alarmType>Reachable</alarmType>
  <alarmTime>2016-09-26T12:19:16Z</alarmTime>
  <recordNumber>1</recordNumber>
</Alarms>

```

getStationRights

Returns the name of the stations rights profile. A station may have multiple station rights profiles, each associated with a different station group ID. For this reason, the stationRightsProfile register will return a dictionary of (sgID, name) pairs. Since this is a chargestation level attribute, it will be returned for all ports.

Table 25: getStationRights

Volttron Point Name	Attribute Name	Register Name	Port #	Type	Units	Starting Value	Writable	Notes
stationRightsProfile	stationRightsProfile	StationRightsRegister		dictionary			FALSE	Dictionary of sgID, rights name tuples.

```

<rightsData>
  <sgID>39491</sgID>
  <sgName>AlcoPark 8</sgName>
  <stationRightsProfile>network_manager</stationRightsProfile>
  <stationData>
    <stationID>1:34003</stationID>
    <stationName>ALCOGARSTATIONS / ALCOPARK 8 -004</stationName>
    <stationSerialNum>115110013369</stationSerialNum>
    <stationMacAddr>000D:6F00:0154:F1FC</stationMacAddr>
  </stationData>
</rightsData>
<rightsData>
  <sgID>58279</sgID>
  <sgName>AlcoGarageStations</sgName>
  <stationRightsProfile>network_manager</stationRightsProfile>
  <stationData>
    <stationID>1:34003</stationID>
    <stationName>ALCOGARSTATIONS / ALCOPARK 8 -004</stationName>
    <stationSerialNum>115110013369</stationSerialNum>
    <stationMacAddr>000D:6F00:0154:F1FC</stationMacAddr>
  </stationData>
</rightsData>

```

getChargingSessionData

Like getAlarms, this query returns a list of session data. The driver interface implementation will make the last session data available.

Table 26: getChargingSessionData

Volttron Point Name	Attribute Name	Register Name	Port #	Type	Units	Starting Value	Writable	Notes
sessionID	sessionID	ChargingSession-Register	1	string			FALSE	
startTime	startTime	ChargingSession-Register	1	date-time			FALSE	
endTime	endTime	ChargingSession-Register	1	date-time			FALSE	
Energy	Energy	ChargingSession-Register	1	float			FALSE	
rfidSerialNumber	rfidSerialNumber	ChargingSession-Register	1	string			FALSE	
driverAccountNumber	driverAccountNumber	ChargingSession-Register	1	string			FALSE	
driverName	driverName	ChargingSession-Register	1	string			FALSE	

```

<ChargingSessionData>
  <stationID>1:34003</stationID>
  <stationName>ALCOGARSTATIONS / ALCOPARK 8 -004</stationName>
  <portNumber>2</portNumber>
  <Address>165 13th St, Oakland, California, 94612, United States</Address>
  <City>Oakland</City>
  <State>California</State>
  <Country>United States</Country>
  <postalCode>94612</postalCode>
  <sessionID>53068029</sessionID>
  <Energy>12.120572</Energy>
  <startTime>2016-10-25T15:53:35Z</startTime>
  <endTime>2016-10-25T20:14:46Z</endTime>
  <userID>452777</userID>
  <recordNumber>1</recordNumber>
  <credentialID>490178743</credentialID>
</ChargingSessionData>

```

getStations

This API call returns a complete description of the chargestation in 40 fields. This information is essentially static and will change infrequently. It should not be scraped on a regular basis. The list of attributes will be included in the registry CSV but are only listed here:

```

stationID, stationManufacturer, stationModel, portNumber, stationName, stationMacAddr,
↪ stationSerialNum, Address, City,
State, Country, postalCode, Lat, Long, Reservable, Level, Mode, Connector, Voltage, ↪
↪ Current, Power, numPorts, Type,
startTime, endTime, minPrice, maxPrice, unitPricePerHour, unitPricePerSession, ↪
↪ unitPricePerKWh, unitPricePerHourThereafter,
sessionTime, Description, mainPhone, orgID, organizationName, sgID, sgName, ↪
↪ currencyCode

```

Engineering Discussion

Questions

- **Allowed python-type** - We propose a register with a *python-type* of dictionary. Is this OK?
- **Scrape Interval** - Scrape all should not return all registers defined in the CSV, we propose fine grained control with a scrape-interval on each register. Response: ok to add extra settings to registry but don't worry about publishing static data with every scrape
- **Data currency** - Since devices are likely to share api calls, at least across ports, we need to think about the currency of the data and possibly allowing this to be a configurable parameter or derived from the scrape interval. Response: add to CSV with default values if not present

Performance

Web service calls across the internet will be significantly slower than typical VOLTRON Bacnet or Modbus devices. It may be prohibitively expensive for each chargepoint sub-agent instance to make individual requests on behalf of its own EVSE+port. We will need to examine the possibility of making a single request for all active chargestations and

sharing that information across driver instances. This could be done through a separate agent that regularly queries the chargepoint network and makes the data available to each sub-agent via an RPC call.

3rd Party Library Dependencies

The chargepoint driver implementation will depend on one additional 3rd part library that is not part of a standard VOLTRON installation:

<https://bitbucket.org/jurko/suds>

Is there a mechanism for drivers to specify their own requirements.txt ?

Driver installation and configuration documentation can reference requirement.txt

3.6.5 Agent Configuration Store

This document describes the configuration store feature and explains how an agent uses it.

The configuration store enables users to store agent configurations on the platform and allows the agent to automatically retrieve them during runtime. Users may update the configurations and the agent will automatically be informed of the changes.

Compatibility

Supporting the configuration store will *not* be required by Agents, however the usage will be strongly encouraged as it should substantially improve user experience.

The previous method for configuring an agent will still be available to agents (and in some cases required). However agents can be created to only work with the configuration store and not support the old method at all.

It will be possible to create an agent to use the traditional method for configuration to establish defaults if no configuration exist in the platform configuration store.

Configuration Names and Paths

Any valid OS file path name is a valid configuration name. Any leading or trailing “/”, “” and whitespace is removed by the store.

The canonical name for the main agent configuration is “config”.

The configuration subsystem remembers the case of configuration names. Name matching is case insensitive both on the Agent and platform side. Configuration names are reported to agent callbacks in the original case used when adding them to the configuration. If a new configuration is store with a different case of an existing name the new name case is used.

Configuration Ownership

Each configuration belongs to one agent and one agent only. When an agent refers to a configuration file via it’s path it does not need to supply any information about its identity to the platform in the file path. The only configurations an agent has direct access to are it’s own. The platform will only inform the owning agent configuration changes.

Configuration File Types

Configurations files come in three types: json, csv, and raw. The type of a configuration file is declared when it is added to or changed in the store.

The parser assumes the first row of every CSV file is a header.

Invalid json or csv files are rejected at the time they are added to the store.

Raw files are unparsed and accepted as is.

Other parsed types may be added in the future.

Configuration File Representation to Agents

JSON

A json file is parsed and represented as appropriate data types to the requester.

Consider a file with the following contents:

```
{
  "result": "PREEMPTED",
  "info": null,
  "data": {
    "agentID": "my_agent",
    "taskID": "my_task"
  }
}
```

The file will be parsed and presented as a dictionary with 3 values to the requester.

CSV

A CSV file is represented as a list of objects. Each object represents a row in the CSV file.

For instance this (simplified) CSV file:

Table 27: Example CSV

Volttron Point Name	Modbus Register	Writable	Point Address
ReturnAirCO2	>f	FALSE	1001
ReturnAirCO2Stpt	>f	TRUE	1011

Will be represented like this:

```
[
  {
    "Volttron Point Name": "ReturnAirCO2",
    "Modbus Register": ">f",
    "Writable": "FALSE",
    "Point Address": "1001"
  },
  {
    "Volttron Point Name": "ReturnAirCO2Stpt",
    "Modbus Register": ">f",
```

(continues on next page)

(continued from previous page)

```

    "Writable": "TRUE",
    "Point Address": "1011"
  }
]

```

Raw

Raw files are represented as a string containing the contents of the file.

File references

The Platform Configuration Store supports referencing one configuration file from another. If a referenced file exists the contents of that file will replace the file reference when the file is sent to the owning agent. Otherwise the reference will be replaced with None.

Only configurations that are parsed by the platform (currently “json” or “csv”) will be examined for references. If the file referenced is another parsed file type (json or csv, currently) then the replacement will be the parsed contents of the file.

In a json object the name of a value will never be considered a reference.

A file reference is any value string that starts with “config://”. The rest of the string is the path in the config store to that configuration. The config store path is converted to lower case for comparison purposes.

Consider the following configuration files named “devices/vav1.config” and “registries/vav.csv”, respectively:

```

{
  "driver_config": {"device_address": "10.1.1.5",
                   "device_id": 500},
  "driver_type": "bacnet",
  "registry_config": "config://registries/vav.csv",
  "campus": "pnnl",
  "building": "isb1",
  "unit": "vav1"
}

```

Table 28: vav.csv

Volttron Point Name	Modbus Register	Writable	Point Address
ReturnAirCO2	>f	FALSE	1001
ReturnAirCO2Stpt	>f	TRUE	1011

The resulting configuration returns when an agent asks for “devices/vav1.config”. The python object will have the following configuration:

```

{
  "driver_config": {"device_address": "10.1.1.5",
                   "device_id": 500},
  "driver_type": "bacnet",
  "registry_config": [
    {
      "Volttron Point Name": "ReturnAirCO2",

```

(continues on next page)

(continued from previous page)

```

        "Modbus Register": ">f",
        "Writable": "FALSE",
        "Point Address": "1001"
    },
    {
        "Volttron Point Name": "ReturnAirCO2Stpt",
        "Modbus Register": ">f",
        "Writable": "TRUE",
        "Point Address": "1011"
    }
],
"campus": "pnnl",
"building": "isbl",
"unit": "vav1"
}

```

Circular references are not allowed. Adding a file that creates a circular reference will cause that file to be rejected by the platform.

If a file is changed in anyway (“NEW”, “UPDATE”, or “DELETE”) and that file is referred to by another file then the platform considers the referring configuration as changed. The configuration subsystem on the Agent will call every callback listening to a file or any file referring to that file either directly or indirectly.

Agent Configuration Sub System

The configuration store shall be implemented on the Agent(client) side in the form of a new subsystem called config.

The subsystem caches configurations as the platform updates the state to the agent. Changes to the cache triggered by an RPC call from the platform will trigger callbacks in the agent.

No callback methods are called until the “onconfig” phase of agent startup. A new phase to agent startup called “onconfig” will be added to the Core class. Originally it was planned to have this run after the “onstart” phase has completed but that is currently not possible. Ideally if an agent is using the config store feature it will not need any “onstart” methods.

When the “onconfig” phase is triggered the subsystem will retrieve the current configuration state from the platform and call all callbacks registered to a configuration in the store to the “NEW” action. No callbacks are called before this point in agent startup.

The first time callbacks are called at agent startup any callbacks subscribed to a configuration called “config” are called first.

Configuration Sub System Agent Methods

These methods are part of the interface available to the Agent.

`config.get(config_name="config")` - Get the contents of a configuration. If no name is provided the contents of the main agent configuration “config” is returned. This may not be called before “ONSTART” methods are called. If called during “ONSTART” phase it will trigger the subsystem to initialize early but will not trigger any callbacks.

`config.subscribe(callback, action=("NEW", "UPDATE", "DELETE"), pattern="*")` - Sets up a callback for handling a configuration change. The platform will automatically update the agent when a configuration changes ultimately triggering all callbacks that match the pattern specified. The action argument describes the types of configuration change action that will trigger the callback. Possible actions are “NEW”, “UPDATE”, and “DELETE” or a tuple of any combination of actions. If no action is supplied the callback happens for all changes. A list of actions can be

supplied if desired. If no file name pattern is supplied then the callback is called for all configurations. The pattern is an regex used match the configuration name.

The callback will also be called if any file referenced by a configuration file is changed.

The signature of the callback method is `callback(config_name, action, contents)` where `file_name` is the file that triggered the callback, `action` is the action that triggered the callback, and `contents` are the new contents of the configuration. Contents will be `None` on a “DELETE” action. All callbacks registered for “NEW” events will be called at agent startup after all “ONSTART” methods have been called. Unlike pubsub subscriptions, this may be called at any point in an agent’s lifetime.

`config.unsubscribe(callback=None, config_name_pattern=None)` - Unsubscribe from configuration changes. Specifying a callback only will unsubscribe that callback from all config name patterns they have been bound to. If a pattern only is specified then all callbacks bound to that pattern will be removed. Specifying both will remove that callback from that pattern. Calling with no arguments will remove all subscriptions. This will not be available in the first version of config store.

`config.unsubscribe_all()` - Unsubscribe from all configuration changes.

`config.set(config_name, contents, trigger_callback=False)` - Set the contents of a configuration. This may not be called before “ONSTART” methods are called. This can be used by an agent to store agent state across agent installations. This will *not* trigger any callbacks unless `trigger_callback` is set to `True`. To prevent deadlock with the platform this method may not be called from a configuration callback function. Doing so will raise a `RuntimeError` exception.

This will not modify the local configuration cache the Agent maintains. It will send the configuration change to the platform and rely on the subsequent `update_config` call.

`config.delete(config_name, trigger_callback=False)` - Remove the configuration from the store. This will *not* trigger any callbacks unless `trigger_callback` is `True`. To prevent deadlock with the platform this method may not be called from a configuration callback function. Doing so will raise a `RuntimeError` exception.

`config.list()` - Returns a list of configuration names.

`config.set_default(config_name, contents, trigger_callback=False)` - Set a default value for a configuration. DOES NOT modify the platform’s configuration store but creates a default configuration that is used for agent configuration callbacks if the configuration does not exist in the store or the configuration is deleted from the store. The callback will only be triggered if `trigger_callback` is `true` and the configuration store subsystem on the agent is not aware of a configuration with that name from the platform store.

Typically this will be called in the `__init__` method of an agent with the parsed contents of the packaged configuration file. This may not be called from a configuration callback. Doing so will raise a `RuntimeError`.

`config.delete_default(config_name, trigger_callback=False)` - Delete a default value for a configuration. I have no idea why you would ever call this. It is here for completeness. This may not be called from a configuration callback. Doing so will raise a `RuntimeError`.

Configuration Sub System RPC Methods

These methods are made available on each agent to allow the platform to communicate changes to a configuration to the affected agent.

As these methods are not part of the exposed interface they are subject to change.

`config.update(config_name, action, contents=None, trigger_callback=True)` - called by the platform when a configuration was changed by some method other than the Agent changing the configuration itself. Trigger callback tells the agent whether or not to call any callbacks associate with the configuration.

Notes on trigger_callback

As the configuration subsystem calls all callbacks in the “onconfig” phase and none are called beforehand the trigger_callback setting is effectively ignored if an agent sets a configuration or default configuration before the end of the “onstart” phase.

Platform Configuration Store

The platform configuration store handles the storage and maintenance of configuration states on the platform.

As these methods are not part of the exposed interface they are subject to change.

Platform RPC Methods

Methods for Agents

Agent methods that change configurations do not trigger any callbacks unless trigger_callback is True.

set_config(config_name, contents, trigger_callback=False) - Change/create a configuration file on the platform.

get_configs() - Get all of the configurations for an Agent.

delete_config(config_name, trigger_callback=False) - Delete a configuration.

Methods for Management

manage_store_config(identity, config_name, contents, config_type="raw") - Change/create a configuration on the platform for an agent with the specified identity

manage_delete_config(identity, config_name) - Delete a configuration for an agent with the specified identity. Calls the agent's update_config with the action “DELETE_ALL” and no configuration name.

manage_delete_store(identity) - Delete all configurations for a VIP IDENTITY.

manage_list_config(identity) - Get a list of configurations for an agent with the specified identity.

manage_get_config(identity, config_name, raw=True) - Get the contents of a configuration file. If raw is set to True this function will return the original file, otherwise it will return the parsed representation of the file.

manage_list_stores() - Get a list of all the agents with configurations.

Direct Call Methods

Services local to the platform who wish to use the configuration store may use two helper methods on the agent class created for this purpose. This allows the auth service to use the config store before the router is started.

delete(self, identity, config_name, trigger_callback=False) - Same as functionality as delete_config, but the caller must specify the identity of the config store.

store(self, identity, config_name, contents, trigger_callback=False) - Same functionality as set_config, but the caller must specify the identity of the config store.

Command Line Interface

The command line interface will consist of a new commands for the `volttron-ctl` program called “`config`” with four sub-commands called “`store`”, “`delete`”, “`list`”, “`get`”. These commands will map directly to the management RPC functions in the previous section.

Disabling the Configuration Store

Agents may optionally disable support for the configuration store by passing `enable_store=False` to the `__init__` method of the Agent class. This allows temporary agents to not spin up the subsystem when it is not needed. Platform service agents that do not yet support the configuration store and the temporary agents used by `volttron-ctl` will set this value.

3.6.6 DNP3

DNP3 (Distributed Network Protocol) is a set of communications protocols that are widely used by utilities such as electric power companies, primarily for SCADA purposes. It was adopted in 2010 as IEEE Std 1815-2010, later updated to 1815-2012.

VOLTRON’s DNP3Agent is an implementation of a DNP3 Outstation as specified in IEEE Std 1815-2012. It engages in bidirectional network communications with a DNP3 Master, which might be located at a power utility.

Like some other VOLTRON protocol agents (e.g. SEP2Agent), DNP3Agent can optionally be front-ended by a DNP3 device driver running under VOLTRON’s MasterDriverAgent. This allows a DNP3 Master to be treated like any other device in VOLTRON’s ecosystem.

The VOLTRON DNP3Agent implementation of an Outstation is built on `pydnp3`, an open-source library from Kisen-sum containing Python language bindings for Automatak’s C++ `opendnp3` library, the de facto reference implementation of DNP3.

DNP3Agent exposes DNP3 application-layer functionality, creating an extensible base from which specific custom behavior can be designed and supported. By default, DNP3Agent acts as a simple transfer agent, publishing data received from the Master on the VOLTRON Message Bus, and responding to RPCs from other VOLTRON agents by sending data to the Master.

RPC Calls

DNP3Agent exposes the following VOLTRON RPC calls:

```
def get_point(self, point_name):
    """
        Look up the most-recently-received value for a given output point.

        @param point_name: The point name of a DNP3 PointDefinition.
        @return: The (unwrapped) value of a received point.
    """

def get_point_by_index(self, group, index):
    """
        Look up the most-recently-received value for a given point.

        @param group: The group number of a DNP3 point.
        @param index: The index of a DNP3 point.
        @return: The (unwrapped) value of a received point.
```

(continues on next page)

(continued from previous page)

```

"""
def get_points(self):
    """
        Look up the most-recently-received value of each configured output point.

    @return: A dictionary of point values, indexed by their VOLTRON point names.
    """

def set_point(self, point_name, value):
    """
        Set the value of a given input point.

    @param point_name: The point name of a DNP3 PointDefinition.
    @param value: The value to set. The value's data type must match the one in the
    ↪DNP3 PointDefinition.
    """

def set_points(self, point_list):
    """
        Set point values for a dictionary of points.

    @param point_list: A dictionary of {point_name: value} for a list of DNP3 points
    ↪to set.
    """

def config_points(self, point_map):
    """
        For each of the agent's points, map its VOLTRON point name to its DNP3 group
    ↪and index.

    @param point_map: A dictionary that maps a point's VOLTRON point name to its
    ↪DNP3 group and index.
    """

def get_point_definitions(self, point_name_list):
    """
        For each DNP3 point name in point_name_list, return a dictionary with each of
    ↪the point definitions.

        The returned dictionary looks like this:

        {
            "point_name1": {
                "property1": "property1_value",
                "property2": "property2_value",
                ...
            },
            "point_name2": {
                "property1": "property1_value",
                "property2": "property2_value",
                ...
            }
        }

        If a definition cannot be found for a point name, it is omitted from the
    ↪returned dictionary.

```

(continues on next page)

(continued from previous page)

```
:param point_name_list: A list of point names.
:return: A dictionary of point definitions.
"""
```

Pub/Sub Calls

DNP3Agent uses two topics when publishing data to the VOLTRON message bus:

- **Point Values (default topic: `dnp3/point`):** As DNP3Agent communicates with the Master, it publishes received point values on the VOLTRON message bus.
- **Outstation status (default topic: `dnp3/status`):** If the status of the DNP3Agent outstation changes, for example if it is restarted, it publishes its new status on the VOLTRON message bus.

Data Dictionary of Point Definitions

DNP3Agent loads and uses a data dictionary of point definitions, which are maintained by agreement between the (DNP3Agent) Outstation and the DNP3 Master. The data dictionary is stored in the agent's registry.

Current Point Values

DNP3Agent tracks the most-recently-received value for each point definition in its data dictionary, regardless of whether the point value's source is a VOLTRON RPC call or a message from the DNP3 Master.

Agent Configuration

The DNP3Agent configuration file specifies the following fields:

- **local_ip:** (string) Outstation's host address (DNS resolved). Default: 0.0.0.0.
- **port:** (integer) Outstation's port number - the port that the remote endpoint (Master) is listening on. Default: 20000.
- **point_topic:** (string) VOLTRON message bus topic to use when publishing DNP3 point values. Default: `dnp3/point`.
- **outstation_status_topic:** (string) Message bus topic to use when publishing outstation status. Default: `dnp3/outstation_status`.
- **outstation_config:** (dictionary) Outstation configuration parameters. All are optional. Parameters include:
 - **database_sizes:** (integer) Size of each outstation database buffer. Default: 10.
 - **event_buffers:** (integer) Size of the database event buffers. Default: 10.
 - **allow_unsolicited:** (boolean) Whether to allow unsolicited requests. Default: True.
 - **link_local_addr:** (integer) Link layer local address. Default: 10.
 - **link_remote_addr:** (integer) Link layer remote address. Default: 1.
 - **log_levels:** (list) List of bit field names (OR'd together) that filter what gets logged by DNP3. Default: [NORMAL]. Possible values: ALL, ALL_APP_COMMS, ALL_COMMS, NORMAL, NOTHING.
 - **threads_to_allocate:** (integer) Threads to allocate in the manager's thread pool. Default: 1.

A sample DNP3Agent configuration file is available in *services/core/DNP3Agent/dnp3agent.config*.

VOLTTRON DNP3 Device Driver

VOLTTRON's DNP3 device driver exposes `get_point/set_point` calls, and `scrapes`, for DNP3 points.

The driver periodically issues DNP3Agent RPC calls to refresh its cached representation of DNP3 data. It issues RPC calls to DNP3Agent as needed when responding to `get_point`, `set_point` and `scrape_all` calls.

For information about the DNP3 driver, see *DNP3 Driver Configuration*.

Installing DNP3Agent

To install DNP3Agent, please consult the installation advice in *services/core/DNP3Agent/README.md*. *README.md* specifies a default agent configuration, which can be overridden as needed.

An agent installation script is available:

```
$ export VOLTTRON_ROOT=<volttron github install directory>
$ cd $VOLTTRON_ROOT
$ source services/core/DNP3Agent/install_dnp3_agent.sh
```

When installing MesaAgent, please note that the agent's point definitions must be loaded into the agent's config store. See `install_dnp3_agent.sh` for an example of how to load them.

For Further Information

Questions? Please contact:

- Rob Calvert (rob@kisensum.com)

3.6.7 Driver Override Specification

This document describes the specification for the global override feature. By default, every user is allowed write access to the devices by the master driver. The override feature will allow the user (for example, building administrator) to override this default behavior and enable the user to lock the write access on the devices for a specified duration of time or indefinitely.

Functional Capabilities

1. User shall be able to specify the following when turning on the override behavior on the devices.
 - Override pattern, for example,
 - If pattern is `campus/building1/*` - Override condition is turned on for all the devices under `campus/building1/`.
 - If pattern is `campus/building1/ahu1` - Override condition is turned on for only `campus/building1/ahu1`
 - The pattern matching shall use bash style filename matching semantics.
 - Time duration over which override behavior is applicable. If the time duration is negative, then override condition is applied indefinitely.

- Optional revert-to-fail-safe-state flag. If the flag is set, master driver shall set all the set points falling under the override condition to its default state/value immediately. This is to ensure that the devices are in fail-safe state when the override/lock feature is removed. If the flag is not set, the device state/value is untouched.
 - Optional staggered revert flag. If this flag is set, reverting of devices will be staggered.
2. User shall be able to disable/turn off the override behavior on devices by specifying:
 - Pattern on which the override/lock feature has be disabled. (example: campus/building/*)
 3. User shall be able to get a list of all the devices with the override condition set.
 4. User shall be able to get a list of all the override patterns that are currently active.
 5. User shall be able to clear all the overrides.
 6. Any changes to override patterns list shall be stored in the config store. On startup, list of override patterns and corresponding end times are retrieved from the config store. If the end time is indefinite or greater than current time for any pattern, then override is set on the matching devices for remaining duration of time.
 7. Whenever a device is newly configured, a check is made to see if it is part of the overridden patterns. If yes, it is added to list of overridden devices.
 8. When a device is being removed, a check is made to see if it is part of the overridden devices. If yes, it is removed from the list of overridden devices.

Driver RPC Methods

`set_override_on(pattern, duration=0.0, failsafe_revert=True, staggered_revert=True)` - Turn on override condition on all the devices matching the pattern. Time duration for the override condition has to be in seconds. For indefinite duration, the time duration has to be ≤ 0.0 .

`set_override_off(pattern)` - Turn off override condition on all the devices matching the pattern. The specified pattern will be removed from the override patterns list. All the devices falling under the given pattern will be removed from the list of overridden devices.

`get_override_devices()` - Get a list of all the devices with override condition.

`get_override_patterns()` - Get a list of override patterns that are currently active.

`clear_overrides()` - Clear all the overrides.

3.6.8 RPC Communication Between Remote Platforms

This document describes RPC communication between different platforms. In the current setup of VOLTRON, if an agent in one platform wants to make a RPC method call on an agent in a different platform, it responsible for establishing and managing the connection with the target platform. Instead, if allow the VIP routers of each platform to make the connection and manage the RPC communication internally, this will reduce the burden on the agents and enable a more seamless RPC communication between agents on different platforms.

VIP Router

The VIP Router on each platform is responsible for establishing and maintaining the connection with remote platforms.

Functional Capabilities

1. Each VOLTRON platform shall have a list of other VOLTRON platforms that it has to establish connection in a config file.
2. The VIP router of each platform connects to other platforms on startup. It is responsible for maintaining the connection (detects disconnects and initiate reconnects etc).
3. The VIP router routes the external RPC message as described in “Messages for External RPC communication” section.

External RPC Subsystem

External RPC subsystem allows an agent to make RPC method calls on agents running in remote platforms.

Functional Capabilities

1. The agent needs to specify the remote platform name as an additional argument in the original RPC call or notify method.
2. The external RPC subsystem on the agent side adds the remote platform name into its VIP frame and sends to the VIP router for routing to correct destination platform. It is described in detail in the next section.

Messages for External RPC communication

The VIP router and external RPC subsystem on the agent side will be using VIP protocol for communication. The communication between the VIP routers and the external RPC subsystem on the agent side can be best explained with an example. Suppose an agent 1 on platform V1 wants to make RPC method call on agent 2 in platform V2. Then the underlying messages exchanged between the two platforms will look like below.

Message format for external RPC subsystem of agent 1 on platform V1 to send to its VIP router.

```

+++
| |                               Empty recipient frame (implies VIP router is the
↪destination)
+++++++
| VIP1 |                           Signature frame
+++++++
|V1 user id |                       Empty user ID frame
+++++++
| 0001 |                             Method request ID, for example "0001"
+++++++
| external_rpc |                     Subsystem, "external_rpc"
+++++++
| external RPC request message|     Dictionary containing destination platform name,
↪destination agent identity,
|                               |     source agent identity, method name and method
↪arguments
+++++++

```

Message sent by VIP router on platform V1 to VIP router of platform V2.

```

+++++++
| V2 |                               Destination platform ID, "V2" in this case
+++++++

```

(continues on next page)

(continued from previous page)

	Empty recipient frame
+-----+	
VIP1	Signature frame
+-----+	
V1 user id	Empty user ID frame
+-----+	
0001	Method Request ID, for example "0001"
+-----+	
external_rpc	Subsystem, "external_rpc"
+-----+	
external RPC request message	Dictionary containing destination platform name, ↵
↵destination agent identity,	
	source platform name, source agent identity, ↵
↵method and arguments	
+-----+	

When the VIP router of platform V2 receives the message, it extracts the destination agent identity from the external RPC request message frame and routes it to the intended agent.

The result of the RPC method execution needs to be returned back to the calling agent. So the messages for the return path are as follows. The source and destination platforms and agents are interchanged in the reply message.

Message sent by external RPC subsystem of agent 2 on platform V2 to its VIP router.

+--+	
	Empty recipient frame (implies destination is VIP ↵
↵router)	
+-----+	
VIP1	Signature frame
+-----+	
V2 user id	Empty user ID frame
+-----+	
0001	Method Request ID, for example "0001"
+-----+	
external_rpc	Subsystem, "external_rpc"
+-----+	
external rpc reply message	Dictionary containing destination platform name, ↵
↵destination agent identity	
	source platform name, source agent identity and ↵
↵method result	
+-----+	

Message sent by VIP router of platform V2 to VIP router of platform V1.

+-----+	
V1	Source platform ID frame, "V1" in this case
+-----+	
	Empty recipient frame
+-----+	
VIP1	Signature frame
+-----+	
V1 user id	Empty user ID frame
+-----+	
0001	Method Request ID, for example "0001"
+-----+	
external_rpc	Subsystem, "external_rpc"
+-----+	

(continues on next page)

(continued from previous page)

external rpc reply message	Dictionary containing destination platform name, ↵
↵ destination agent identity	
	source platform name, source agent identity and ↵
↵ method result	
+-----+	

The VIP router of platform V1 extracts the destination agent identity from the external RPC reply message frame and routes it to the calling agent.

Methods for External RPC Subsystem

`call(peer, method, *args, **kwargs)` - New 'external_platform' parameter need to be added in kwargs to the original RPC subsystem call. If the platform name of the target platform is passed into the 'external_platform' parameter, the RPC method on the target platform gets executed.

`notify(peer, method, *args, **kwargs)` - New 'external_platform' parameter need to be added in kwargs to the original RPC subsystem notify method. If the platform name of the target platform is passed into the 'external_platform' parameter, the RPC method on the target platform gets executed.

`handle_external_rpc_subsystem(message)` - Handler for the external RPC subsystem messages. It executes the requested RPC method and returns the result to the calling platform.

3.6.9 MesaAgent

MesaAgent is a VOLTRON agent that handles MESA-ESS DNP3 outstation communications. It subclasses and extends the functionality of VOLTRON's DNP3Agent. Like DNP3Agent, MesaAgent models a DNP3 outstation, communicating with a DNP3 master.

DNP3 (Distributed Network Protocol) is a set of communications protocols that are widely used by utilities such as electric power companies, primarily for **SCADA** purposes. It was adopted in 2010 as **IEEE Std 1815-2010**, later updated to **1815-2012**.

VOLTRON's MesaAgent and DNP3Agent are implementations of a DNP3 Outstation as specified in IEEE Std 1815-2012. They engage in bidirectional network communications with a DNP3 Master, which might be located at a power utility.

MESA-ESS is an extension and enhancement to DNP3. It builds on the basic DNP3 communications protocol, adding support for more complex structures, including functions, arrays, curves and schedules. The draft specification for MESA-ESS, as well as a spreadsheet of point definitions, can be found at <http://mesastandards.org/mesa-ess-2016/>.

VOLTRON's DNP3Agent and MesaAgents implementations of an Outstation are built on pydnp3, an open-source library from Kisensum containing Python language bindings for Automatak's C++ [opendnp3](#) library, the de facto reference implementation of DNP3.

MesaAgent exposes DNP3 application-layer functionality, creating an extensible base from which specific custom behavior can be designed and supported, including support for MESA functions, arrays and selector blocks. By default, MesaAgent acts as a simple transfer agent, publishing data received from the Master on the VOLTRON Message Bus, and responding to RPCs from other VOLTRON agents by sending data to the Master. Properties of the point and function definitions also enable the use of more complex controls for point data capture and publication.

MesaAgent was developed by Kisensum for use by 8minutenergy, which provided generous financial support for the open-source contribution to the VOLTRON platform, along with valuable feedback based on experience with the agent in a production context.

RPC Calls

MesaAgent exposes the following VOLTRON RPC calls:

```
def get_point(self, point_name):
    """
        Look up the most-recently-received value for a given output point.

        @param point_name: The point name of a DNP3 PointDefinition.
        @return: The (unwrapped) value of a received point.
    """

def get_point_by_index(self, data_type, index):
    """
        Look up the most-recently-received value for a given point.

        @param data_type: The data_type of a DNP3 point.
        @param index: The index of a DNP3 point.
        @return: The (unwrapped) value of a received point.
    """

def get_points(self):
    """
        Look up the most-recently-received value of each configured output point.

        @return: A dictionary of point values, indexed by their point names.
    """

def get_configured_points(self):
    """
        Look up the most-recently-received value of each configured point.

        @return: A dictionary of point values, indexed by their point names.
    """

def set_point(self, point_name, value):
    """
        Set the value of a given input point.

        @param point_name: The point name of a DNP3 PointDefinition.
        @param value: The value to set. The value's data type must match the one in the
        ↪DNP3 PointDefinition.
    """

def set_points(self, point_dict):
    """
        Set point values for a dictionary of points.

        @param point_dict: A dictionary of {point_name: value} for a list of DNP3 points
        ↪to set.
    """

def config_points(self, point_map):
    """
        For each of the agent's points, map its VOLTRON point name to its DNP3 group
        ↪and index.

        @param point_map: A dictionary that maps a point's VOLTRON point name to its
        ↪DNP3 group and index.
    """
```

(continues on next page)

```

"""
def get_point_definitions(self, point_name_list):
    """
        For each DNP3 point name in point_name_list, return a dictionary with each of
        ↪the point definitions.

        The returned dictionary looks like this:

        {
            "point_name1": {
                "property1": "property1_value",
                "property2": "property2_value",
                ...
            },
            "point_name2": {
                "property1": "property1_value",
                "property2": "property2_value",
                ...
            }
        }

        If a definition cannot be found for a point name, it is omitted from the
        ↪returned dictionary.

        :param point_name_list: A list of point names.
        :return: A dictionary of point definitions.
    """

def get_selector_block(self, point_name, edit_selector):
    """
        Return a dictionary of point values for a given selector block.

        :param point_name: Name of the first point in the selector block.
        :param edit_selector: The index (edit selector) of the block.
        :return: A dictionary of point values.
    """

def reset(self):
    """
        Reset the agent's internal state, emptying point value caches. Used during
        ↪iterative testing.
    """

```

Pub/Sub Calls

MesaAgent uses three topics when publishing data to the VOLTRON message bus:

- **Point Values (default topic: dnp3/point):** As MesaAgent communicates with the Master, it publishes received point values on the VOLTRON message bus.
- **Functions (default topic: mesa/function):** When MesaAgent receives a function step with a “publish” action value, it publishes the current state of the function (all steps received to date) on the VOLTRON message bus.
- **Outstation status (default topic: mesa/status):** If the status of the MesaAgent outstation changes, for example if it is restarted, it publishes its new status on the VOLTRON message bus.

Data Dictionaries of Point and Function Definitions

MesaAgent loads and uses data dictionaries of point and function definitions, which are maintained by agreement between the (MesaAgent) Outstation and the DNP3 Master. The data dictionaries are stored in the agent's registry.

Current Point Values

MesaAgent tracks the most-recently-received value for each point definition in its data dictionary, regardless of whether the point value's source is a VOLTRON RPC call or a message from the DNP3 Master.

Agent Configuration

The MesaAgent configuration specifies the following fields:

- **local_ip:** (string) Outstation's host address (DNS resolved). Default: 0.0.0.0.
- **port:** (integer) Outstation's port number - the port that the remote endpoint (Master) is listening on. Default: 20000.
- **point_topic:** (string) VOLTRON message bus topic to use when publishing DNP3 point values. Default: dnp3/point.
- **function_topic:** (string) Message bus topic to use when publishing MESA-ESS functions. Default: mesa/function.
- **outstation_status_topic:** (string) Message bus topic to use when publishing outstation status. Default: mesa/outstation_status.
- **all_functions_supported_by_default:** (boolean) When deciding whether to reject points for unsupported functions, ignore the values of their 'supported' points: simply treat all functions as supported. Used primarily during testing. Default: False.
- **function_validation:** (boolean) When deciding whether to support sending single points to MesaAgent. If function_validation is True, MesaAgent will raise an exception when receiving any invalid point in current function. If function_validation is False, MesaAgent will reset current function to None instead of raising the exception. Default: False.
- **outstation_config:** (dictionary) Outstation configuration parameters. All are optional. Parameters include:
 - **database_sizes:** (integer) Size of each outstation database buffer. Default: 10.
 - **event_buffers:** (integer) Size of the database event buffers. Default: 10.
 - **allow_unsolicited:** (boolean) Whether to allow unsolicited requests. Default: True.
 - **link_local_addr:** (integer) Link layer local address. Default: 10.
 - **link_remote_addr:** (integer) Link layer remote address. Default: 1.
 - **log_levels:** (list) List of bit field names (OR'd together) that filter what gets logged by DNP3. Default: [NORMAL]. Possible values: ALL, ALL_APP_COMMS, ALL_COMMS, NORMAL, NOTHING.
 - **threads_to_allocate:** (integer) Threads to allocate in the manager's thread pool. Default: 1.

A sample MesaAgent configuration file is available in **services/core/DNP3Agent/mesaagent.config**.

3.6.10 Installing MesaAgent

To install MesaAgent, please consult the installation advice in **services/core/DNP3Agent/README.md**, which includes advice on installing **pydnp3**, a library upon which DNP3Agent depends.

After installing libraries as described in README.md, the agent can be installed from a command-line shell as follows:

```
$ export VOLTRON_ROOT=<volttron github install directory>
$ cd $VOLTRON_ROOT
$ source services/core/DNP3Agent/install_mesa_agent.sh
```

README.md specifies a default agent configuration, which can be overridden as needed.

Here are some things to note when installing MesaAgent:

- MesaAgent source code resides in, and is installed from, a dnp3 subdirectory, thus allowing it to be implemented as a subclass of the base DNP3 agent class. When installing MesaAgent, inform the install script that it should build from the mesa subdirectory by exporting the following environment variable:
 - \$ export AGENT_MODULE=dnp3.mesa.agent
- The agent's point and function definitions must be loaded into the agent's config store. See the install_mesa_agent.sh script for an example of how to load them.

3.6.11 For Further Information

Questions? Please contact:

- Anh Nguyen at ChargePoint (anh.nguyen@chargepoint.com)

3.6.12 Message Bus Visualization and Debugging - Specification

NOTE: This is a planning document, created prior to implementation of the VOLTRON Message Debugger. It describes the tool's general goals, but it's not always accurate about specifics of the ultimate implementation. For a description of Message Debugging as implemented, with advice on how to configure and use it, please see *Message-Debugging*.

Description

VOLTRON agents send messages to each other on the VOLTRON message bus. It can be useful to examine the contents of this message stream while debugging and troubleshooting agents and drivers.

In satisfaction of this specification, a new Message Monitor capability will be implemented allowing VOLTRON agent/driver developers to monitor the message stream, filter it for an interesting set of messages, and display the contents and characteristics of each message.

Some elements below are central to this effort (required), while others are useful improvements (optional) that may be implemented if time permits.

Feature: Capture Messages and Display a Message Summary

When enabled, the Message Monitor will capture details about a stream of routed messages. On demand, it will display a message summary, either in real time as the messages are routed, or retrospectively.

A summary view will convey the high level interactions occurring between VOLTRON agents as conversations that may be expanded for more detail. A simple RPC call that involves 4 message send/rcv segments will be displayed as a single object that can be expanded. In this way, the message viewer will provide a higher-level view of message bus activity than might be gleaned from verbose logs using grep.

Pub/sub interactions will be summarized at the topic level with high-level statistics such as the number of subscribers, # of messages published during the capture period, etc. Drilling into the interaction might show the last message published with the ability to drill deeper into individual messages. A diff display would show how the published data is changing.

Summary view

```
- 11:09:31.0831  RPC      set_point          charge.control  platform.driver
| - params: ('set_load', 10)  return: True
- 11:09:31.5235  Pub/Sub   devices/my_device  platform.driver  2 subscribers
| - Subscriber: charge.control
  | - Last message 11:09:31.1104:
    [
      {
        'Heartbeat': True,
        'PowerState': 0,
        'temperature': 50.0,
        'ValveState': 0
      },
      ...
    ]
  | - Diff to 11:09:21.5431:
      'temperature': 48.7,
```

The summary's contents and format will vary by message subsystem.

RPC request/response pairs will be displayed on a single line:

```
(volttron) dl:volttron myname$ msmon --agent='(Agent1,Agent2) '

Agent1                                     Agent2
2016-11-22T11:09:31.083121+00:00 rpc: devices/my_topic; 2340972387; sent 2016-11-
↪22T11:09:31.277933+00:00 responded: 0.194 sec
2016-11-22T11:09:32.005938+00:00 rpc: devices/my_topic; 2340972388; sent 2016-11-
↪22T11:09:32.282193+00:00 responded: 0.277 sec
2016-11-22T11:09:33.081873+00:00 rpc: devices/my_topic; 2340972389; sent 2016-11-
↪22T11:09:33.271199+00:00 responded: 0.190 sec
2016-11-22T11:09:34.049139+00:00 rpc: devices/my_topic; 2340972390; sent 2016-11-
↪22T11:09:34.285393+00:00 responded: 0.236 sec
2016-11-22T11:09:35.053183+00:00 rpc: devices/my_topic; 2340972391; sent 2016-11-
↪22T11:09:35.279317+00:00 responded: 0.226 sec
2016-11-22T11:09:36.133948+00:00 rpc: devices/my_topic; 2340972392; sent 2016-11-
↪22T11:09:36.133003+00:00 dequeued
```

When PubSub messages are displayed, each message's summary will include its count of subscribers:

```
(volttron) dl:volttron myname$ msmon --agent=(Agent1)

Agent1
2016-11-22T11:09:31.083121+00:00 pubsub: devices/my_topic; 2340972487; sent; 2 subs
2016-11-22T11:09:32.005938+00:00 pubsub: devices/my_topic; 2340972488; sent; 2 subs
2016-11-22T11:09:33.081873+00:00 pubsub: devices/my_topic; 2340972489; sent; 2 subs
2016-11-22T11:09:34.049139+00:00 pubsub: devices/my_topic; 2340972490; sent; 2 subs
2016-11-22T11:09:35.053183+00:00 pubsub: devices/my_topic; 2340972491; sent; 2 subs
```

While streaming output of a message summary, a defined keystroke sequence will “pause” the output, and another keystroke sequence will “resume” displaying the stream.

Feature: Capture and Display Message Details

The Message Monitor will capture a variety of details about each message, including:

1. Sending agent ID
2. Receiving agent ID
3. User ID
4. Message ID
5. Subsystem
6. Topic
7. Message data
8. Message lifecycle timestamps, in UTC (when sent, dequeued, responded)
9. Message status (sent, responded, error, timeout)
10. Message size
11. Other message properties TBD (e.g., queue depth?)

On demand, it will display these details for a single message ID:

```
(volttron)d1:volttron myname$ msmon --id='2340972390'  
  
2016-11-22T11:09:31.053183+00:00 (Agent1)  
INFO:  
  Subsystem: 'pubsub',  
  Sender: 'Agent1',  
  Topic: 'devices/my_topic',  
  ID: '2340972390',  
  Sent: '2016-11-22T11:09:31.004986+00:00',  
  Message:  
  [  
    {  
      'Heartbeat': True,  
      'PowerState': 0,  
      'temperature': 50.0,  
      'ValveState': 0  
    },  
    {  
      'Heartbeat':  
      {  
        'units': 'On/Off',  
        'type': 'integer'  
      },  
      'PowerState':  
      {  
        'units': '1/0',  
        'type': 'integer'  
      },  
      'temperature':  
      {
```

(continues on next page)

(continued from previous page)

```
        'units': 'Fahrenheit',
        'type': 'integer'
    },
    'ValveState':
    {
        'units': '1/0',
        'type': 'integer'
    }
}
```

A VOLTRON message ID is not unique to a single message. A group of messages in a “conversation” may share a common ID, for instance during RPC request/response exchanges. When detailed display of all messages for a single message ID is requested, they will be displayed in chronological order.

Feature: Display Message Statistics

Statistics about the message stream will also be available on demand:

1. Number of messages sent, by agent, subsystem, topic
2. Number of messages received, by agent, subsystem, topic

Feature: Filter the Message Stream

The Message Monitor will be able to filter the message stream display to show only those messages that match a given set of criteria:

1. Sending agent ID(s)
2. Receiving agent ID(s)
3. User ID(s)
4. Subsystem(s)
5. Topic - Specific topic(s)
6. Topic - Prefix(es)
7. Specific data value(s)
8. Sampling start/stop time
9. Other filters TBD

User Interface: Linux Command Line

A Linux command-line interface will enable the following user actions:

1. Enable message tracing
2. Disable message tracing
3. Define message filters
4. Define verbosity of displayed-message output
5. Display message stream

6. Begin recording messages
7. Stop recording messages
8. Display recorded messages
9. Play back (re-send) recorded messages

Feature (not implemented): Watch Most Recent

Optionally, the Message Monitor can be asked to “watch” a specific data element. In that case, it will display the value of that element in the most recent message matching the filters currently in effect. As the data to be displayed changes, the display will be updated in place without scrolling (similar to “top” output):

```
(volttron) dl:volttron myname$ msmon --agent='(Agent1)' --watch='temperature'  
  
Agent1  
2016-11-22T11:09:31.053183+00:00 pubsub: my_topic; 2340972487; sent; 2 subs;   
↔temperature=50
```

Feature (not implemented): Regular Expression Support

It could help for the Message Monitor’s filtering logic to support regular expressions. Regex support has also been requested (Issue #207) when identifying a subscribed pub/sub topic during VOLTTRON message routing.

Optionally, regex support will be implemented in Message Monitor filtering criteria, and also (configurably) during VOLTTRON topic matching.

Feature (not implemented): Message Stream Record and Playback

The Message Monitor will be able to “record” and “play back” a message sequence:

1. Capture a set of messages as a single “recording”
2. Inspect the contents of the “recording”
3. “Play back” the recording – re-send the recording’s message sequence in VOLTTRON

Feature (not implemented): On-the-fly Message Inspection and Modification

VOLTTRON message inspection and modification, on-the-fly, may be supported from the command line. The syntax and implementation would be similar to pdb (Python Debugger), and might be written as an extension to pdb.

Capabilities:

1. Drill-down inspection of message contents.
2. Set a breakpoint based on message properties, halting upon routing a matching message.
3. While halted on a breakpoint, alter a message’s contents.

Feature (not implemented): PyCharm Debugging Plugin

VOLTTRON message debugging may also be published as a PyCharm plugin. The plugin would form a more user-friendly interface for the same set of capabilities described above – on-the-fly message inspection and modification, with the ability to set a breakpoint based on message properties.

User Interface (not implemented): PCAP/Wireshark

Optionally, we may elect to render the message trace as a stream of PCAP data, thereby exploiting Wireshark's filtering and display capabilities. This would be in accord with the enhancement suggested in VOLTRON Issue #260.

User Interface (not implemented): Volttron Central Dashboard Widget

Optionally, the Message Monitor will be integrated as a new Volttron Central dashboard widget, supporting each of the following:

1. Enable/Disable the monitor
2. Filter messages
3. Configure message display details
4. Record/playback messages

User Interface (not implemented): Graphical Display of Message Sequence

Optionally, the Volttron Central dashboard widget will provide graphical display of message sequences, allowing enhanced visualization of request/response patterns.

Related Development: PyCharm Documentation

Also included in this effort will be a contribution to VOLTRON documentation about installing and configuring a PyCharm environment for developing, debugging and testing VOLTRON agents and drivers.

3.6.13 Engineering Design Notes

Grabbing Messages Off the Bus

This tool depends on reading and storing all messages that pass through the VIP router. The Router class already has hooks that allow for the capturing of messages at various points in the routing workflow. The BaseRouter abstract class defines `issue(self, topic, frames, extra)`. This method is called from `BaseRouter.route` and `BaseRouter._send` during the routing of messages. The `topic` parameter (not to be confused with a message topic found in `frames`) identifies the point or state in the routing workflow at which the issue was called.

The defined `topics` are: `INCOMING`, `OUTGOING`, `ERROR` and `UNROUTABLE`. Most messages will result in two calls, one with the `INCOMING` topic as the message enters the router and one with the `OUTGOING` topic as the message is sent on to its destination. Messages without a recipient are intended for the router itself and do not result in an `OUTGOING` call to `issue`.

`Router.issue` contains the concrete implementation of the method. It does two things:

1. It writes the topic, frames and optional extra parameters to the logger using the `FramesFormatter`.
2. It invokes `self._tracker.hit(topic, frames, extra)`. The `Tracker` class collects statistics by topic and counts the messages within a topic by peer, user and subsystem.

The `issue` method can be modified to optionally publish the `issue` messages to an in-process ZMQ address that the message-viewing tool will subscribe to. This will minimize changes to core VOLTRON code and minimize the impact of processing these messages for debugging.

Message Processor

The message processor will subscribe to messages coming out of the Router.issue() method and process these messages based on the current message viewer configuration. Messages will be written to a SQLite db since this is packaged with Python and currently used by other VOLTTRON agents.

Message Viewer

The message viewer will display messages from the SQLite db. We need to consider whether it should also subscribe to receiving messages in real-time. The viewer will be responsible for displaying message statistics and will provide a command line interface to filter and display messages.

Message Db Schema

```
message(id, created_on, issue_topic, extras, sender, recipient, user_id, msg_id, ↵
↵subsystem, data)
```

msg_id will be used to associate pairs of incoming/outgoing messages. ****note:** data will be a jsonified list of frames, alternatively we could add a message_data table with one row per frame.

A session table will track the start and end of a debug session and, at the end of a session, record statistics on the messages in the session.

```
session(id, created_on, name, start_time, end_time, num_messages)
```

The command line tool will allow users to delete old sessions and select a session for review/playback.

3.6.14 PubSub Communication Between Remote Platforms

This document describes pubsub communication between different platforms. The goal of this specification is to improve the current setup of having a forward historian to forward local pubsub messages to remote platforms. So the agents interested in receiving PubSub messages from external platforms will not need to have a forward historian running in source platform to forward pubsub messages to the interested destination platforms. The VIP router will now do all the work; it shall use Routing Service to internally manage connections with external VOLTTRON platforms and use PubSubService for the actual inter platform PubSub communication. For Future: The specification will need to be extended to support PubSub communication between platforms that are multiple hops away. The VIP router of each platform shall need to maintain a routing table and use it to forward pubsub messages to subscribed platforms that are multiple hops away. The routing table shall contain shortest path to each destination platform.

Functional Capabilities

1. Each VOLTTRON platform shall have a list of other VOLTTRON platforms that it has to connect to in a config file.
2. Routing Service of each platform connects to other platforms on startup.
3. The Routing Service in each platform is responsible for connecting to (and also initiate reconnection if required), monitoring and disconnecting from each external platform. The function of Routing Service is explained in detail in Routing Service section.
4. Platform to platform pubsub communication shall be using VIP protocol with the subsystem frame set to “pub-sub”.
5. PubSubService of each VOLTTRON platform shall maintain a list of local and external subscriptions.

6. Each VIP router sends its list of external subscriptions to other connected platforms in the following cases
 - a. On startup
 - b. When a new subscription is added
 - c. When an existing subscription is removed
 - d. When a new platform gets connected
7. When a remote platform disconnection is detected, all stale subscriptions related to that platform shall be removed.
8. Whenever an agent publishes a message to a specific topic, the PubSubService on the local platform first checks the topic against its list of local subscriptions. If a local subscription exists, it sends the publish message to corresponding local subscribers.
9. PubSubService shall also check the topic against list of external subscriptions. If an external subscription exists, it shall use Routing Service to send the publish message to the corresponding external platform.
10. Whenever a router receives messages from other platform, it shall check the destination platform in the incoming message.
 - a. If the destination platform is the local platform, it hand overs the publish message to PubSubService which checks the topic against list of external subscriptions. If an external subscription matches, PubSubService forwards the message to all the local subscribers subscribed to that topic.
 - b. If the destination platform is not the local platform, it discards the message.

Routing Service

1. Routing Service shall maintain connection status (CONNECTING, CONNECTED, DISCONNECTED etc.) for each external platform.
2. In order to establish connection with an external VOLTRON platform, the server key of the remote platform is needed.

The Routing Service shall connect to an external platform once it obtains the server key for that platform from the KeyDiscoveryService.
3. Routing Service shall exchange “hello”/”welcome” handshake messages with the newly connected remote platform to confirm the connection. It shall use VIP protocol with the subsystem frame set to “routing_table” for the handshake messages.
3. Routing Service shall monitor the connection status and inform PubSubService whenever a remote platform gets connected/disconnected.

For Future

1. Each VIP router shall exchange its routing table with its connected platforms on startup and whenever a new platform gets connected or disconnected.
2. The router shall go through each entry in the routing table that it received from other platforms and calculate the shortest, most stable path to each remote platform. It then sends the updated routing table to other platforms for adjustments in the forwarding paths (in their local routing table) if any.
3. Whenever a VIP router detects a new connection, it adds an entry into the routing table and sends updated routing table to its neighboring platforms. Each router in the other platforms shall update and re-calculate the forwarding paths in its local routing table and forward to rest of the platforms.

4. Similarly, whenever a VIP router detects a remote platform disconnection, it deletes the entry in the routing table for that platform and forwards the routing table to other platforms to do the same.

KeyDiscovery Service

1. Each platform tries to obtain the platform discovery information - platform name, VIP address and server key of remote VOLTTRON platforms through HTTP discovery service at startup.
2. If unsuccessful, it shall make regular attempts to obtain discovery information until successful.
3. The platform discovery information shall then be sent to the Routing Service using VIP protocol with subsystem frame set to "routing_table".

Messages for Routing Service

Below shows example messages that are applicable to the Routing Service.

Message sent by KeyDiscovery Service containing the platform discovery information (platform name, VIP address and server key) of a remote platform.

```

+--+
| |                               Empty recipient frame
+-----+
| VIP1 |                           Signature frame
+-----+
| |                               Empty user ID frame
+-----+
| 0001 |                           Request ID, for example "0001"
+-----+
| routing_table |                   Subsystem, "routing_table"
+-----+-----+
| normalmode_platform_connection | Type of operation, "normalmode_platform_connection"
+-----+-----+
| platform discovery information |
| of external platform           | platform name, VIP address and server key of
↪external platform
+-----+-----+
| platform name                 | Remote platform for which the server key belongs to.
+-----+-----+

```

Handshake messages between two newly connected external VOLTTRON platform to confirm successful connection.

Message from initiating platform

```

+--+
| |                               Empty recipient frame
+-----+
| VIP1 |                           Signature frame
+-----+
| |                               Empty user ID frame
+-----+
| 0001 |                           Request ID, for example "0001"
+-----+-----+
| routing_table |                   Subsystem, "routing_table"
+-----+-----+
| hello |                           Operation, "hello"

```

(continues on next page)

(continued from previous page)

```

+-----+
| hello |           Hello handshake request frame
+-----+-----+
| platform name |   Platform initiating a "hello"
+-----+-----+

```

Reply message from the destination platform

```

+-+
| |               Empty recipient frame
+-+-----+
| VIP1 |         Signature frame
+-+-----+
| |             Empty user ID frame
+-+-----+
| 0001 |        Request ID, for example "0001"
+-----+-----+
| routing_table | Subsystem, "routing_table"
+-----+-----+
| hello |       Operation, "hello"
+-+-----+
| welcome |    Welcome handshake reply frame
+-----+-----+
| platform name | Platform sending reply to "hello"
+-----+-----+

```

Messages for PubSub communication

The VIP routers of each platform shall send pubsub messages between platforms using VIP protocol message semantics. Below shows an example of external subscription list message sent by VOLTRON platform V1 router to VOLTRON platform V2.

```

+-+
| |               Empty recipient frame
+-+-----+
| VIP1 |         Signature frame
+-+-----+
| V1 user id |   Empty user ID frame
+-+-----+
| 0001 |        Request ID, for example "0001"
+-----+-----+
| pubsub |       Subsystem, "pubsub"
+-----+-----+
| external_list | Operation, "external_list" in this case
+-----+-----+
| List of       |
| subscriptions | Subscriptions dictionary consisting of VOLTRON platform id and
↳list of topics as
+-----+-----+ key - value pairings, for example: { "V1": ["devices/rtu3"]}

```

This shows an example of external publish message sent by VOLTRON platform V2 router to VOLTRON platform V1.

```

+-+
| |               Empty recipient frame

```

(continues on next page)

```

+-----+
| VIP1 |           Signature frame
+-----+
|V1 user id |       Empty user ID frame
+-----+
| 0001 |           Request ID, for example "0001"
+-----+
| pubsub |         Subsystem, "pubsub"
+-----+
| external_publish | Operation, "external_publish" in this case
+-----+
| topic      |      Message topic
+-----+
| publish message | Actual publish message frame
+-----+

```

API

Methods for Routing Service

`external_route()` - This method receives message frames from external platforms, checks the subsystem frame and redirects to appropriate subsystem (routing table, pubsub) handler. It shall run within a separate thread and get executed whenever there is a new incoming message from other platforms.

`setup()` - This method initiates socket connections with all the external VOLTRON platforms configured in the config file. It also starts monitor thread to monitor connections with external platforms.

`handle_subsystem(frames)` - Routing Service subsystem handler to handle serverkey message from KeyDiscoveryService and “hello/welcome” handshake message from external platforms.

`send_external(instance_name, frames)` - This method sends input message to specified VOLTRON platform/instance.

`register(type, handler)` - Register method for PubSubService to register for connection and disconnection events.

`disconnect_external_instances(instance_name)` - Disconnect from specified VOLTRON platform.

`close_external_connections()` - Disconnect from all external VOLTRON platforms.

`get_connected_platforms()` - Return list of connected platforms.

Methods for PubSubService

`external_platform_add(instance_name)` - Send external subscription list to newly connected external VOLTRON platform.

`external_platform_drop(instance_name)` - Remove all subscriptions for the specified VOLTRON platform

`update_external_subscriptions(frames)` - Store/Update list of external subscriptions as per the subscription list provided in the message frame.

`_distribute_external(frames)` - Publish the message all the external platforms that have subscribed to the topic. It uses `send_external_pubsub_message()` of router to send out the message.

`external_to_local_publish(frames)` - This method retrieves actual message from the message frame, checks the message topic against list of external subscriptions and sends the message to corresponding subscribed agents.

Methods for agent pubsub subsystem

subscribe(peer, prefix, callback, bus=' ', all_platforms=False) - The existing 'subscribe' method is modified to include optional keyword argument - 'all_platforms'. If 'all_platforms' is set to True, the agent is subscribing to topic from local publisher and from external platform publishers.

3.6.15 SEP 2.0 DER Support

Version 1.0

Smart Energy Profile 2.0 (SEP2, IEEE 2030.5) specifies a REST architecture built around the core HTTP verbs: GET, HEAD, PUT, POST and DELETE. A specification for the SEP2 protocol can be found [here](#).

SEP2 EndDevices (clients) POST XML resources representing their state, and GET XML resources containing command and control information from the server. The server never reaches out to the client unless a "subscription" is registered and supported for a particular resource type. This implementation does not use SEP2 registered subscriptions.

The SEP2 specification requires HTTP headers, and it explicitly requires RESTful response codes, for example:

- 201 - "Created"
- 204 - "No Content"
- 301 - "Moved Permanently"
- etc.

SEP2 message encoding may be either XML or EXI. Only XML is supported in this implementation.

SEP2 requires HTTPS/TLS version 1.2 along with support for the cipher suite TLS_ECDHE_ECDSA_WITH_AES_128_CCM_8. Production installation requires a certificate issued by a SEP2 CA. The encryption requirement can be met by using a web server such as Apache to proxy the HTTPs traffic.

SEP2 discovery, if supported, must be implemented by an xmDNS server. Avahi can be modified to perform this function.

Function Sets

SEP2 groups XML resources into "Function Sets." Some of these function sets provide a core set of functionality used across higher-level function sets. This implementation implements resources from the following function sets:

- Time
- Device Information
- Device Capabilities
- End Device
- Function Set Assignments
- Power Status
- Distributed Energy Resources

Distributed Energy Resources

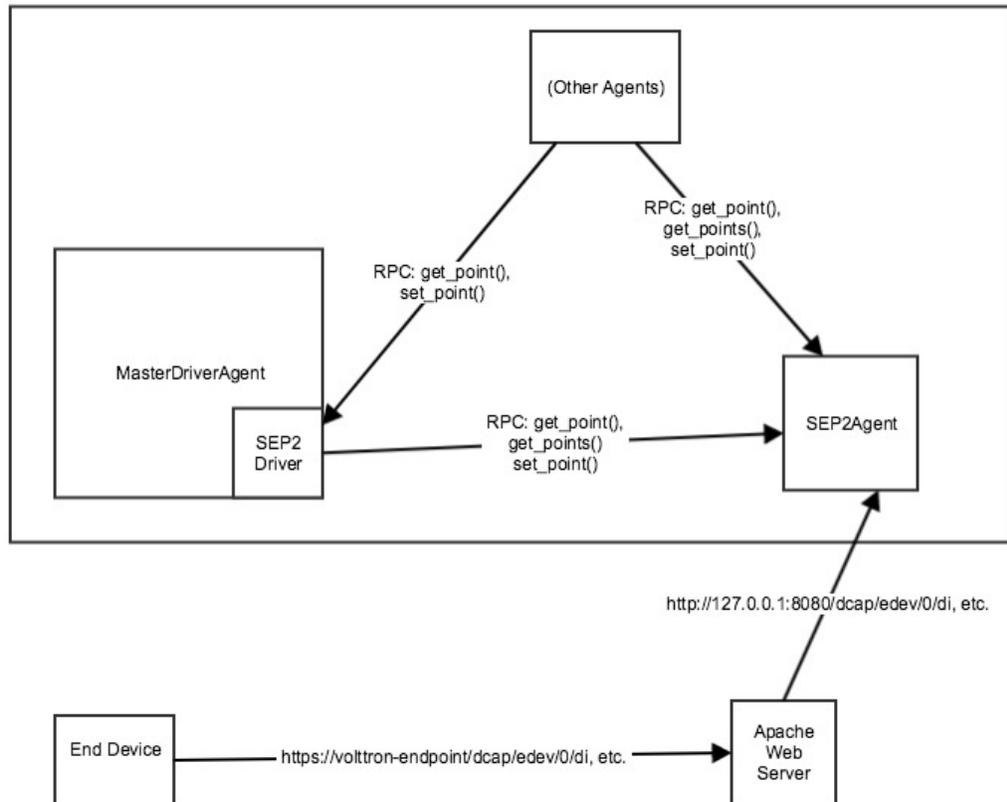
Distributed energy resources (DERs) are devices that generate energy, e.g., solar inverters, or store energy, e.g., battery storage systems, electric vehicle supply equipment (EVSEs). These devices are managed by a SEP2 DER server using DERPrograms which are described by the SEP2 specification as follows:

Servers host one or more DERPrograms, which in turn expose DERControl events to DER clients. DERControl instances contain attributes that allow DER clients to respond to events that are targeted to their device type. A DERControl instance also includes scheduling attributes that allow DER clients to store and process future events. These attributes include start time and duration, as well an indication of the need for randomization of the start and / or duration of the event. The SEP2 DER client model is based on the SunSpec Alliance Inverter Control Model [SunSpec] which is derived from IEC 61850-90-7 [61850] and [EPRI].

EndDevices post multiple SEP2 resources describing their status. The following is an example of a Power Status resource that might be posted by an EVSE (vehicle charging station):

```
<PowerStatus xmlns="http://zigbee.org/sep" xmlns:xsi="http://www.w3.org/2001/
↪XMLSchema-instance" href="/sep2/edev/96/ps">
  <batteryStatus>4</batteryStatus>
  <changedTime>1487812095</changedTime>
  <currentPowerSource>1</currentPowerSource>
  <estimatedChargeRemaining>9300</estimatedChargeRemaining>
  <PEVInfo>
    <chargingPowerNow>
      <multiplier>3</multiplier>
      <value>-5</value>
    </chargingPowerNow>
    <energyRequestNow>
      <multiplier>3</multiplier>
      <value>22</value>
    </energyRequestNow>
    <maxForwardPower>
      <multiplier>3</multiplier>
      <value>7</value>
    </maxForwardPower>
    <minimumChargingDuration>11280</minimumChargingDuration>
    <targetStateOfCharge>10000</targetStateOfCharge>
    <timeChargeIsNeeded>9223372036854775807</timeChargeIsNeeded>
    <timeChargingStatusPEV>1487812095</timeChargingStatusPEV>
  </PEVInfo>
</PowerStatus>
```

Design Details



VOLTRON's SEP2 implementation includes a SEP2Agent and a SEP2 device driver, as described below.

VOLTRON SEP2Agent

SEP2Agent implements a SEP2 server that receives HTTP POST/PUT requests from SEP2 devices. The requests are routed to SEP2Agent over the VOLTRON message bus by VOLTRON's MasterWebService. SEP2Agent returns an appropriate HTTP response. In some cases (e.g., DERControl requests), this response includes a data payload.

SEP2Agent maps SEP2 resource data to a VOLTRON SEP2 data model based on SunSpec, using block numbers and point names as defined in the SunSpec Information Model, which in turn is harmonized with 61850. The data model is given in detail below.

Each device's data is stored by SEP2Agent in an EndDevice memory structure. This structure is not persisted to a database. Each EndDevice retains only the most recently received value for each field.

SEP2Agent exposes RPC calls for getting and setting EndDevice data.

VOLTRON SEP2 Device Driver

The SEP2 device driver is a new addition to VOLTRON MasterDriverAgent's family of standard device drivers. It exposes `get_point/set_point` calls for SEP2 EndDevice fields.

The SEP2 device driver periodically issues SEP2Agent RPC calls to refresh its cached representation of EndDevice data. It issues RPC calls to SEP2Agent as needed when responding to `get_point`, `set_point` and `scrape_all` calls.

Field Definitions

These field IDs correspond to the ones in the SEP2 device driver's configuration file, sep2.csv. They have been used in that file's "Voltron Point Name" column and also in its "Point Name" column.

Field ID	SEP2 Resource/Property	Description	Units	Type
b1_Md	device_information mfModel	Model (32 char lim).		string
b1_Opt	device_information lfdi	Long-form device identifier (32 char lim).		string
b1_SN	abstract_device sfdi	Short-form device identifier (32 char lim).		string
b1_Vr	device_information mfHwVer	Version (16 char lim).		string
b113_A	mirror_meter_reading PhaseCurrentAvg	AC current.	A	float
b113_DCA	mirror_meter_reading InstantPackCurrent	DC current.	A	float
b113_DCV	mirror_meter_reading LineVoltageAvg	DC voltage.	V	float
b113_DCW	mirror_meter_reading PhasePowerAvg	DC power.	W	float
b113_PF	mirror_meter_reading PhasePFA	AC power factor.	%	float
b113_WH	mirror_meter_reading EnergyIMP	AC energy.	Wh	float
b120_AhrRtg	der_capability rtgAh	Usable capacity of the battery. Maximum charge minus minimum charge.	Ah	float
b120_ARtg	der_capability rtgA	Maximum RMS AC current level capability of the inverter.	A	float
b120_MaxChaRte	der_capability rtgMaxChargeRate	Maximum rate of energy transfer into the device.	W	float
b120_MaxDisChaRte	der_capability rtgMaxDischargeRate	Maximum rate of energy transfer out of the device.	W	float
b120_WHRtg	der_capability rtgWh	Nominal energy rating of the storage device.	Wh	float
3.6. Platform Specifications				509
b120_WRtg	der_capability rtgW	Continuous power output capability of the inverter.	W	float

Revising and Expanding the Field Definitions

The SEP2-to-SunSpec field mappings in this implementation are a relatively thin subset of all possible field definitions. Developers are encouraged to expand the definitions.

The procedure for expanding the field mappings requires you to make changes in two places:

1. Update the driver's point definitions in `services/core/MasterDriverAgent/master_driver/sep2.csv`
2. Update the SEP2-to-SunSpec field mappings in `services/core/SEP2Agent/sep2/end_device.py` and `__init__.py`

When updating VOLTTRON's SEP2 data model, please use field IDs that conform to the SunSpec block-number-and-field-name model outlined in the SunSpec Information Model Reference (see the link below).

For Further Information

SunSpec References:

- Information model specification: <http://sunspec.org/wp-content/uploads/2015/06/SunSpec-Information-Models-12041.pdf>
- Information model reference spreadsheet: <http://sunspec.org/wp-content/uploads/2015/06/SunSpec-Information-Model-Reference.xlsx>
- Inverter models: <http://sunspec.org/wp-content/uploads/2015/06/SunSpec-Inverter-Models-12020.pdf>
- Energy storage models: <http://sunspec.org/wp-content/uploads/2015/06/SunSpec-Energy-Storage-Models-12032.pdf>

Questions? Please contact:

- Rob Calvert (rob@kisensum.com) or James Sheridan (james@kisensum.com)

3.6.16 Tagging agent specification

Description

Tagging service provides VOLTTRON users the ability to add semantic tags to different topics so that topic can be queried by tags instead of specific topic name or topic name pattern.

Taxonomy

VOLLTTRON will use tags from [Project Haystack](#). Tags defined in haystack will be imported into VOLTTRON and grouped by categories to tag topics and topic name prefix.

Dependency

Once data in VOLTTRON has been tagged, users will be able to query topics based on tags and use the resultant topics to query the historian

Features

1. User should be able to tag individual components of a topic such as campus, building, device, point etc.
2. Using the tagging service users should only be able to add tags already defined in the volttron tagging schema. New tags should be explicitly added to the tagging schema before it can be used to tag topics or topic prefix

3. Users should be able batch process and tag multiple topic names or topic prefix using a template. At the end of this, users should be notified about the list of topics that did not confirm to the template. This will help users to individually add or edit tags for those specific topics
4. When users query for topics based on a tag, the results would correspond to the current metadata values. It is up to the calling agent/application to periodically query for latest updates if needed.
5. Users should be able query based on tags on a specific topic or its topic prefix/parents
6. Allow for count and skip parameters in queries to restrict count and allow pagination

API

1. Get the list of tag categories available

rpc call to tagging service method **'get_categories'** with optional parameters:

1. **include_description** - set to True to return available description for each category. Default = False
2. **skip** - number of categories to skip. this parameter along with count can be used for paginating results
3. **count** - limit the total number of tag categories returned to given count
4. **order** - ASCENDING or DESCENDING. By default, it will be sorted in ascending order

2. Get the list of tags for a specific category

rpc call to tagging service method **'get_tags_by_category'** with parameter:

1. **category** - <category name>

and optional parameters:

2. **include_kind - indicate if result should include the** kind/data type for tags returned. Defaults to False
3. **include_description - indicate if result should include** available description for tags returned. Defaults to False
4. **skip** - number of tags to skip. this parameter along with count can be used for paginating results
5. **count** - limit the total number of tags returned to given count
6. **order** - ASCENDING or DESCENDING. By default, it will be sorted in ascending order

3. Get the list of tags for a topic_name or topic_name_prefix

rpc call to tagging service method **get_tags_by_topic**

with parameter

1. **topic_prefix** - topic name or topic name prefix

and optional parameters:

2. **include_kind - indicate if result should include the** kind/data type for tags returned. Defaults to False
3. **include_description - indicate if result should include** available description for tags returned. Defaults to False
4. **skip** - number of tags to skip. this parameter along with count can be used for paginating results

5. **count** - limit the total number of tags returned to given count
6. **order** - ASCENDING or DESCENDING. By default, it will be sorted in ascending order

4. Find topic names by tags

rpc call to tagging service method **'get_topics_by_tags'** with the one or more of the following parameters

1. **and_condition** - dictionary of tag and its corresponding values that should be matched using equality operator or a list of tags that should exist/be true. Tag conditions are combined with AND condition. Only topics that match all the tags in the list would be returned
2. **or_condition** - dictionary of tag and its corresponding values that should be matched using equality operator or a list of tags that should exist/be true. Tag conditions are combined with OR condition. Topics that match any of the tags in the list would be returned. If both **and_condition** and **or_condition** are provided then they are combined using AND operator.
3. **condition** - conditional statement to be used for matching tags. If this parameter is provided the above two parameters are ignored. The value for this parameter should be an expression that contains one or more query conditions combined together with an "AND" or "OR". Query conditions can be grouped together using parenthesis. Each condition in the expression should conform to one of the following format:
 1. <tag name/ parent.tag_name> <binary_operator> <value>
 2. <tag name/ parent.tag_name>
 3. <tag name/ parent.tag_name> LIKE <regular expression within single quotes
 4. the word NOT can be prefixed before any of the above three to negate the condition.
 5. expressions can be grouped with parenthesis.

For example

```
condition="tag1 = 1 and not (tag2 < ' ' and tag2 > ' ') and tag3_
↳and NOT tag4 LIKE '^a.*b$'"
condition="NOT (tag5='US' OR tag5='UK') AND NOT tag3 AND NOT_
↳(tag4 LIKE 'a.*')"
condition="campusRef.geoPostalCode='20500' and equip and boiler"
```

6. **skip** - number of topics to skip. this parameter along with count can be used for paginating results
7. **count** - limit the total number of tag topics returned to given count
8. **order** - ASCENDING or DESCENDING. By default, it will be sorted in ascending order

5. Query data based on tags

Use above api to get topics by tags and then use the result to query historian's query api.

6. Add tags to specific topic name or topic name prefix

rpc call to tagging service method **'add_topic_tags'** with parameters:

1. **topic_prefix** - topic name or topic name prefix
2. **tags** - {<valid tag>:value, <valid_tag>: value,... }

3. **update_version** - True/False. Default to False. If set to True and if any of the tags update an existing tag value the older value would be preserved as part of tag version history. **NOTE:** This is a placeholder. Current version does not support versioning.

7. Add tags to multiple topics

rpc call to tagging service method **'add_tags'** with parameters:

1. **tags** - dictionary object containing the topic and the tag details. format:

```
<topic_name or prefix or topic_name pattern>: {<valid tag>:<value>, ... }, ... }
```

2. **update_version** - True/False. Default to False. If set to True and if any of the tags update an existing tag value the older value would be preserved as part of tag version history

Use case examples

1. Loading news tags for an existing VOLTRON instance

Current topic names:

```
/campus1/building1/deviceA1/point1
/campus1/building1/deviceA1/point2
/campus1/building1/deviceA1/point3
/campus1/building1/deviceA2/point1
/campus1/building1/deviceA2/point2
/campus1/building1/deviceA2/point3
/campus1/building1/deviceB1/point1
/campus1/building1/deviceB1/point2
/campus1/building1/deviceB2/point1
/campus1/building1/deviceB1/point2
```

Step 1:

Create a python dictionary object contains topic name pattern and its corresponding tag/value pair. Use topic pattern names to fill out tags that can be applied to more than one topic or topic prefix. Use specific topic name and topic prefix for tags that apply only to a single entity. For example:

```
{
  # tags specific to building1
  '/campus1/building1':
    {
      'site': true,
      'dis': ": 'some building description'",
      'yearBuilt': 2015,
      'area': '24000sqft'
    },
  # tags that apply to all device of a specific type
  '/campus1/building1/deviceA*':
```

(continues on next page)

(continued from previous page)

```
{
  'dis': "building1 chilled water system - CHW",
  'equip': true,
  'campusRef': 'campus1',
  'siteRef': 'campus1/building1',
  'chilled': true,
  'water' : true,
  'secondaryLoop': true
}
# tags that apply to point1 of all device of a specific type
'/campus1/building1/deviceA*/point1':
{
  'dis': "building1 chilled water system - point1",
  'point': true,
  'kind': 'Bool',
  'campusRef': 'campus1',
  'siteRef': 'campus1/building1'
}
# tags that apply to point2 of all device of a specific type
'/campus1/building1/deviceA*/point2':
{
  'dis': "building1 chilled water system - point2",
  'point': true,
  'kind': 'Number',
  'campusRef': 'campus1',
  'siteRef': 'campus1/building1'
}
# tags that apply to point3 of all device of a specific type
'/campus1/building1/deviceA*/point3':
{
  'dis': "building1 chilled water system - point3",
  'point': true,
  'kind': 'Number',
  'campusRef': 'campus1',
  'siteRef': 'campus1/building1'
}
# tags that apply to all device of a specific type
'/campus1/building1/deviceB*':
{
  'dis': "building1 device of type B",
  'equip': true,
  'chilled': true,
  'water' : true,
  'secondaryLoop': true,
  'campusRef': 'campus1',
  'siteRef': 'campus1/building1'
}
# tags that apply to point1 of all device of a specific type
'/campus1/building1/deviceB*/point1':
{
  'dis': "building1 device B - point1",
  'point': true,
  'kind': 'Bool',
  'campusRef': 'campus1',
  'siteRef': 'campus1/building1',
  'command': true
}
}
```

(continues on next page)

(continued from previous page)

```
# tags that apply to point1 of all device of a specific type
'/campus1/building1/deviceB*/point2':
{
  'dis': "building1 device B - point2",
  'point': true,
  'kind': 'Number',
  'campusRef': 'campus1',
  'siteRef': 'campus1/building1'
}
}
```

Step 2: Create tags using template above

Make an RPC call to the `add_tags` method and pass the python dictionary object

Step 3: Create tags specific to a point or device

Any tags that were not included in step one and needs to be added later can be added using the rpc call to tagging service either the method `'add_topic_tags'` `'add_tags'`

For example:

```
agent.vip.rpc.call(
    'platform.tagging',
    'add_topic_tags',
    topic_prefix='/campus1/building1/deviceA1',
    tags={'tag1': 'value'})
```

```
agent.vip.rpc.call(
    'platform.tagging',
    'add_topic_tags',
    tags={
        '/campus1/building1/deviceA2':
            {'tag1': 'value'},
        '/campus1/building1/deviceA2/point1':
            {'equipRef': 'campus1/building1/deviceA2'}
    }
)
```

2. Querying based on a topic's tag and it parent's tags

Query - Find all points that has the tag `'command'` and belong to a device/unit that has a tag `'chilled'`

```
agent.vip.rpc.call(
    'platform.tagging',
    'get_topics_by_tags',
    condition='temperature and equip.chilled)
```

In the above code block `'command'` and `'chilled'` are the tag names that would be searched, but since the tag `'chilled'` is prefixed with `'equip.'` the tag in a parent topic

The above query would match the topic `'/campus1/building1/deviceB1/point1'` if tags in the system are as follows

‘/campus1/building1/deviceB1/point1’ tags:

```
{
'dis': "building1 device B - point1",
'point': true,
'kind': 'Bool',
'campusRef': 'campus1',
'siteRef': 'campus1/building1',
'equipRef': 'campus1/building1/deviceB1',
'command': true
}
```

‘/campus1/building1/deviceB1’ tags

```
{
'dis': "building1 device of type B",
'equip': true,
'chilled': true,
'water' : true,
'secondaryLoop': true,
'campusRef': 'campus1',
'siteRef': 'campus1/building1'
}
```

Possible future improvements

1. Versioning - When a value of a tag is changed, users should be prompted to verify if this change denotes a new version or a value correction. If this value denotes a new version, then older value of the tag should be preserved in a history/audit store
2. Validation of tag values based on data type
3. Support for units validation and conversions
4. Processing and saving geologic coordinates that can enable users to do geo-spatial queries in databases that support it.

3.6.17 VOLTTRON Web Framework

This document describes the interaction between web enabled agents and the MasterWebService agent.

The web framework enables agent developers to expose JSON, static, and websocket endpoints.

Web SubSystem

Enabling

The web subsystem is not enabled by default as it is only required by a small subset of agents. To enable the web subsystem the platform instance must have an enabled the web server and the agent must pass `enable_web=True` to the agent constructor.

Methods

The web subsystem allows an agent to register three different types of endpoints; path based, JSON and websocket. A path based endpoint allows the agent to specify a prefix and a static path on the file system to serve static files. The prefix can be a regular expression.

Note: The web subsystem is only available when the constructor contains `enable_web=True`.

The below examples are within the context of an object that has extended the `volttron.platform.vip.Agent` base class.

Note: For all endpoint methods the first match wins. Therefore ordering which endpoints are registered first becomes important.

```
@Core.receiver('onstart')
def onstart(self, sender, **kwargs):
    """
    Allow serving of static content from /var/www
    """
    self.vip.web.register_path(r'^/vc/.*', '/var/www')
```

JSON endpoints allows an agent to serve data responses to specific queries from a web client.non-static responses. The agent will pass a callback to the subsystem which will be called when the endpoint is triggered.

```
def jsonrpc(env, data):
    """
    The main entry point for jsonrpc data
    """
    return {'dynamic': 'data'}

@Core.receiver('onstart')
def onstart(self, sender, **kwargs):
    """
    Register the /vc/jsonrpc endpoint for doing json-rpc based methods
    """
    self.vip.web.register_endpoint(r'/vc/jsonrpc', self.jsonrpc)
```

Websocket endpoints allow bi-directional communication between the client and the server. Client connections can be authenticated during the opening of a websocket through the response of an open callback.

```
def _open_authenticate_ws_endpoint(self, fromip, endpoint):
    """
    A client attempted to open an endpoint to the server.

    Return True or False if the endpoint should be allowed.

    :rtype: bool
    """
    return True

def _ws_closed(self, endpoint):
    _log.debug("CLOSED endpoint: {}".format(endpoint))

def _ws_received(self, endpoint, message):
```

(continues on next page)

(continued from previous page)

```

_log.debug("RECEIVED endpoint: {} message: {}".format(endpoint,
                                                    message))

@Core.receiver('onstart')
def onstart(self, sender, **kwargs):
    self.vip.web.register_websocket(r'/vc/ws', self.open_authenticate_ws_endpoint,
    ↪self._ws_closed, self._ws_received)

```

3.7 Applications

Community-contributed applications, agents and drivers that are not directly integrated into the VOLTRON core platform reside in a separate github repository, <https://github.com/VOLTRON/voltron-applications>. This section provides user guides and other documents for those contributions.

3.7.1 Simulation Subsystem

The simulation subsystem includes a set of device simulators and a clock that can run faster (or slower) than real time. It can be used to test VOLTRON agents or drivers. It could be particularly useful when simulating multi-agent and/or multi-driver scenarios.

The source code for the agents and drivers comprising this subsystem resides in the <https://github.com/VOLTRON/voltron-applications> github repository.

This subsystem is designed to be extended easily. Its initial delivery includes a set of simulated energy devices that report status primarily in terms of power (kilowatts) produced and consumed. It could easily be adapted, though, to simulate and report data for devices that produce, consume and manage resources other than energy.

Three agents work together to run a simulation:

1. **SimulationClockAgent.** This agent manages the simulation's clock. After it has been supplied with a start time, a stop time, and a clock-speed multiplier, and it has been asked to start a simulation, it provides the current simulated time in response to requests. If no stop time has been provided, the SimulationClockAgent continues to manage the simulation clock until the agent is stopped. If no clock-speed multiplier has been provided, the simulation clock runs at normal wallclock speed.
2. **SimulationDriverAgent.** Like MasterDriverAgent, this agent is a front-end manager for device drivers. It handles `get_point/set_point` requests from other agents, and it periodically "scrapes" and publishes each driver's points. If a device driver has been built to run under MasterDriverAgent, with a few minor modifications (detailed below) it can be adapted to run under SimulationDriverAgent.
3. **SimulationAgent.** This agent configures, starts, and reports on a simulation. It furnishes a variety of configuration parameters to the other simulation agents, starts the clock, subscribes to scraped driver points, and generates a CSV output file.

Four device drivers have been provided:

1. **storage (simstorage).** The storage driver simulates an energy storage device (i.e., a battery). When it receives a power dispatch value (positive to charge the battery, negative to discharge it), it adjusts its charging behavior accordingly. Its reported power doesn't necessarily match the dispatch value, since (like an actual battery) it stays within configured max-charge/max-discharge limits, and its power dwindles as its state of charge approaches a full or empty state.
2. **pv (simpv).** The PV driver simulates a photovoltaic array (solar panels), reporting the quantity of solar power produced. Solar power is calculated as a function of (simulated) time, using a data file of incident-sunlight metrics. A year's worth of solar data has been provided as a sample resource.

3. **load (simload)**. The load driver simulates the behavior of a power consumer such as a building, reporting the quantity of power consumed. It gets its power metrics as a function of (simulated) time from a data file of power readings. A year's worth of building-load data has been provided as a sample resource.
4. **meter (simmeter)**. The meter driver simulates the behavior of a circuit's power meter. This driver, as delivered, is actually just a shell of a simulated device. It's able to report power as a function of (simulated) time, but it has no built-in default logic for deciding what particular power metrics to report.

Linux Installation

The following steps describe how to set up and run a simulation. They assume that VOLTRON / `volttron` and VOLTRON / `volttron-applications` repositories have been downloaded from github, and that Linux shell variables `$VOLTRON_ROOT` and `$VOLTRON_APPLICATIONS_ROOT` point at the root directories of these repositories.

First, create a soft link to the applications directory from the `volttron` directory, if that hasn't been done already:

```
$ cd $VOLTRON_ROOT
$ ln -s $VOLTRON_APPLICATIONS_ROOT applications
```

With VOLTRON running, load each simulation driver's configuration into a "simulation.driver" config store:

```
$ export SIMULATION_DRIVER_ROOT=$VOLTRON_ROOT/applications/kisensum/Simulation/
↪SimulationDriverAgent

$ volttron-ctl config store simulation.driver simload.csv $SIMULATION_DRIVER_ROOT/
↪simload.csv --csv
$ volttron-ctl config store simulation.driver devices/simload $SIMULATION_DRIVER_ROOT/
↪simload.config

$ volttron-ctl config store simulation.driver simmeter.csv $SIMULATION_DRIVER_ROOT/
↪simmeter.csv --csv
$ volttron-ctl config store simulation.driver devices/simmeter $SIMULATION_DRIVER_
↪ROOT/simmeter.config

$ volttron-ctl config store simulation.driver simpv.csv $SIMULATION_DRIVER_ROOT/simpv.
↪csv --csv
$ volttron-ctl config store simulation.driver devices/simpv $SIMULATION_DRIVER_ROOT/
↪simpv.config

$ volttron-ctl config store simulation.driver simstorage.csv $SIMULATION_DRIVER_ROOT/
↪simstorage.csv --csv
$ volttron-ctl config store simulation.driver devices/simstorage $SIMULATION_DRIVER_
↪ROOT/simstorage.config
```

Install and start each simulation agent:

```
$ export SIMULATION_ROOT=$VOLTRON_ROOT/applications/kisensum/Simulation
$ export VIP_SOCKET="ipc://$VOLTRON_HOME/run/vip.socket"

$ python scripts/install-agent.py \
  --vip-identity simulation.driver \
  --tag simulation.driver \
  --agent-source $SIMULATION_ROOT/SimulationDriverAgent \
  --config $SIMULATION_ROOT/SimulationDriverAgent/simulationdriver.config \
  --force \
  --start
```

(continues on next page)

(continued from previous page)

```
$ python scripts/install-agent.py \  
  --vip-identity simulationclock \  
  --tag          simulationclock \  
  --agent-source $SIMULATION_ROOT/SimulationClockAgent \  
  --config       $SIMULATION_ROOT/SimulationClockAgent/simulationclock.config \  
  --force \  
  --start  
  
$ python scripts/install-agent.py \  
  --vip-identity simulationagent \  
  --tag          simulationagent \  
  --agent-source $SIMULATION_ROOT/SimulationAgent \  
  --config       $SIMULATION_ROOT/SimulationAgent/simulationagent.config \  
  --force \  
  --start
```

SimulationAgent Configuration Parameters

This section describes SimulationAgent’s configurable parameters. Each of these has a default value and behavior, allowing the simulation to be run “out of the box” without configuring any parameters.

Type	Param Name	Data Type	Default	Comments
General	agent_id	str	simulation	
General	heartbeat_period	int sec	5	
General	sim_driver_list	list of str	[simload, simmeter, simpv, simstorage]	Allowed keywords are simload, simmeter, simpv, simstorage.
Clock	sim_start	date- time str	2017-02-02 13:00:00	
Clock	sim_end	date- time str	None	If None, sim doesn't stop.
Clock	sim_speed	float sec	180.0	This is a multiplier, e.g. 1 sec actual time = 180 sec sim time.
Load	load_timestamp_column_header	str	actual_date	
Load	load_power_column_header	str	load_kw	
Load	load_data_frequency	int min	15	
Load	load_data_year	str	2015	
Load	load_csv_file_path	str	~/repos/voltron-applications/kisensum/Simulation/SimulationAgent/data/load_and_pv.csv	~ and shell variables in the pathname will be expanded. The file must exist.
PV	pv_panel_area	float m2	50.0	
PV	pv_efficiency	float 0.0- 1.0	0.75	
PV	pv_data_frequency	int min	30	
PV	pv_data_year	str	2015	
PV	pv_csv_file_path	str	~/repos/voltron-applications/kisensum/Simulation/SimulationAgent/data/nrel_pv_readings.csv	~ and shell variables in the pathname will be expanded. The file must exist.
Storage	storage_soc_kwh	float kWh	30.0	
Storage	storage_max_soc_kwh	float kWh	50.0	
Storage	storage_max_charge_kw	float kW	15.0	
Storage	storage_max_discharge_kw	float kW	12.0	
Storage	storage_reduced_charge_threshold	float 0.0- 1.0	0.80	Charging will be reduced when SOC % > this value.
Storage	storage_reduced_discharge_soc_threshold	float 0.0- 1.0	0.20	Discharging will be reduced when SOC % < this value.
Dispatch	storage_setpoint_rule	str key- word	oscillation	See below.
Dispatch	positive_dispatch_kw	float kW >= 0.0	15.0	
Dispatch	negative_dispatch_kw	float kW <= 0.0	-15.0	

The **oscillation** setpoint rule slowly oscillates between charging and discharging based on the storage device's state of charge (SOC):

```
If SOC < (`go_positive_if_below` * `storage_max_soc_kwh`):
    dispatch power = `positive_dispatch_kw`

If SOC > (`go_negative_if_above` * `storage_max_soc_kwh`)
    dispatch power = `negative_dispatch_kw`

Otherwise:
    dispatch power is unchanged from its previous value.
```

The **alternate** setpoint rule is used when `storage_setpoint_rule` has been configured with any value other than **oscillation**. It simply charges at the dispatched charging value (subject to the constraints of the other parameters, e.g. `storage_max_discharge_kw`):

```
dispatch power = `positive_dispatch_kw`
```

Driver Parameters and Points

Load Driver

The load driver's parameters specify how to look up power metrics in its data file.

Type	Name	Data Type	Default	Comments
Param/Point	csv_file_path	string		This parameter must be supplied by the agent.
Param/Point	times-tamp_column_header	string	local_date	
Param/Point	power_column_header	string	load_kw	
Param/Point	data_frequency_min	int	15	
Param/Point	data_year	string	2015	
Point	power_kw	float	0.0	
Point	last_timestamp	datetime		

Meter Driver

Type	Name	Data Type	Default	Comments
Point	power_kw	float	0.0	
Point	last_timestamp	datetime		

PV Driver

The PV driver's parameters specify how to look up sunlight metrics in its data file, and how to calculate the power generated from that sunlight.

Type	Name	Data Type	Default	Comments
Param/Point	csv_file_path	string		This parameter must be supplied by the agent.
Param/Point	max_power_kw	float	10.0	
Param/Point	panel_area	float	50.0	
Param/Point	efficiency	float	0.75	
Param/Point	data_frequency_min	int	30	
Param/Point	data_year	string	2015	
Point	power_kw	float	0.0	
Point	last_timestamp	datetime		

Storage Driver

The storage driver's parameters describe the device's power and SOC limits, its initial SOC, and the SOC thresholds at which charging and discharging start to be reduced as its SOC approaches a full or empty state. This reduced power is calculated as a straight-line reduction: charging power is reduced in a straight line from `reduced_charge_soc_threshold` to 100% SOC, and discharging power is reduced in a straight line from `reduced_discharge_soc_threshold` to 0% SOC.

Type	Name	Data Type	Default	Comments
Param/Point	max_charge_kw	float	15.0	
Param/Point	max_discharge_kw	float	15.0	
Param/Point	max_soc_kwh	float	50.0	
Param/Point	soc_kwh	float	25.0	
Param/Point	reduced_charge_soc_threshold	float	0.8	
Param/Point	reduced_discharge_soc_threshold	float	0.2	
Point	dispatch_kw	float	0.0	
Point	power_kw	float	0.0	
Point	last_timestamp	datetime		

Working with the Sample Data Files

The Load and PV simulation drivers report power readings that are based on metrics from sample data files. The software distribution includes sample Load and PV files containing at least a year's worth of building-load and sunlight data.

CSV files containing different data sets of load and PV data can be substituted by specifying their paths in SimulationAgent's configuration, altering its other parameters if the file structures and/or contents are different.

Load Data File

`load_and_pv.csv` contains building-load and PV power readings at 15-minute increments from 01/01/2014 - 12/31/2015. The data comes from a location in central Texas. The file's data columns are: `utc_date`, `local_date`, `time_offset`, `load_kw`, `pv_kw`. The load driver looks up the row with a matching `local_date` and returns its `load_kw` value.

Adjust the following SimulationAgent configuration parameters to change how load power is derived from the data file:

- Use `load_csv_file_path` to set the path of the sample data file
- Use `load_data_frequency_min` to set the frequency of the sample data

- Use `load_data_year` to set the year of the sample data
- Use `load_timestamp_column_header` to indicate the header name of the timestamp column
- Use `load_power_column_header` to indicate the header name of the power column

PV Data File

`nrel_pv_readings.csv` contains irradiance data at 30-minute increments from 01/01/2015 - 12/31/2015, downloaded from NREL's National Solar Radiation Database, <https://nsrdb.nrel.gov>. The file's data columns are: Year, Month, Day, Hour, Minute, DHI, DNI, Temperature. The PV driver looks up the row with a matching date/time and uses its DHI (diffuse horizontal irradiance) to calculate the resulting solar power produced:

```
power_kw = irradiance * panel_area * efficiency / elapsed_time_hrs
```

Adjust the following SimulationAgent configuration parameters to change how solar power is derived from the data file:

- Use `pv_csv_file_path` to set the path of the sample data file
- Use `pv_data_frequency_min` to set the frequency of the sample data
- Use `pv_data_year` to set the year of the sample data
- Use `pv_panel_area` and `pv_efficiency` to indicate how to transform an irradiance measurement in wh/m2 into a power reading in kw.

If a PV data file will be used that has a column structure which differs from the one in the supplied sample, an adjustment may need to be made to the `simpv` driver software.

Running the Simulation

One way to monitor the simulation's progress is to look at debug trace in VOLTTRON's log output, for example:

```
2017-05-01 15:05:42,815 (simulationagent-1.0 9635) simulation.agent DEBUG: 2017-05-01_
↳15:05:42.815484 Initializing drivers
2017-05-01 15:05:42,815 (simulationagent-1.0 9635) simulation.agent DEBUG: _
↳Initializing Load: timestamp_column_header=local_date, power_column_header=load_kw, _
↳data_frequency_min=15, data_year=2015, csv_file_path=/Users/robcalvert/repos/
↳volttron-applications/kisensum/Simulation/SimulationAgent/data/load_and_pv.csv
2017-05-01 15:05:42,823 (simulationagent-1.0 9635) simulation.agent DEBUG: _
↳Initializing PV: panel_area=50, efficiency=0.75, data_frequency_min=30, data_
↳year=2015, csv_file_path=/Users/robcalvert/repos/volttron-applications/kisensum/
↳Simulation/SimulationAgent/data/nrel_pv_readings.csv
2017-05-01 15:05:42,832 (simulationagent-1.0 9635) simulation.agent DEBUG: _
↳Initializing Storage: soc_kwh=30.0, max_soc_kwh=50.0, max_charge_kw=15.0, max_
↳discharge_kw=12.0, reduced_charge_soc_threshold = 0.8, reduced_discharge_soc_
↳threshold = 0.2
2017-05-01 15:05:42,844 (simulationagent-1.0 9635) simulation.agent DEBUG: 2017-05-01_
↳15:05:42.842162 Started clock at sim time 2017-02-02 13:00:00, end at 2017-02-02_
↳16:00:00, speed multiplier = 180.0
2017-05-01 15:05:57,861 (simulationagent-1.0 9635) simulation.agent DEBUG: 2017-05-01_
↳15:05:57.842164 Reporting at sim time 2017-02-02 13:42:00
2017-05-01 15:05:57,862 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/
↳simload/power_kw = 486.1
2017-05-01 15:05:57,862 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/
↳simpv/power_kw = -0.975
```

(continues on next page)

(continued from previous page)

```

2017-05-01 15:05:57,862 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/
↳simstorage/dispatch_kw = 0.0
2017-05-01 15:05:57,862 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/
↳simstorage/last_timestamp = 2017-02-02 13:33:00
2017-05-01 15:05:57,862 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/
↳simstorage/power_kw = 0.0
2017-05-01 15:05:57,862 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/
↳simstorage/soc_kwh = 30.0
2017-05-01 15:05:57,862 (simulationagent-1.0 9635) simulation.agent DEBUG: net_power_
↳kw = 485.125
2017-05-01 15:05:57,862 (simulationagent-1.0 9635) simulation.agent DEBUG: report_
↳time = 2017-02-02 13:42:00
2017-05-01 15:05:57,862 (simulationagent-1.0 9635) simulation.agent DEBUG:
↳Setting storage dispatch to 15.0 kW
2017-05-01 15:06:12,901 (simulationagent-1.0 9635) simulation.agent DEBUG: 2017-05-01_
↳15:06:12.869471 Reporting at sim time 2017-02-02 14:30:00
2017-05-01 15:06:12,901 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/
↳simload/power_kw = 467.5
2017-05-01 15:06:12,901 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/
↳simpv/power_kw = -5.925
2017-05-01 15:06:12,901 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/
↳simstorage/dispatch_kw = 15.0
2017-05-01 15:06:12,901 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/
↳simstorage/last_timestamp = 2017-02-02 14:27:00
2017-05-01 15:06:12,901 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/
↳simstorage/power_kw = 15.0
2017-05-01 15:06:12,901 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/
↳simstorage/soc_kwh = 43.5
2017-05-01 15:06:12,901 (simulationagent-1.0 9635) simulation.agent DEBUG: net_power_
↳kw = 476.575
2017-05-01 15:06:12,901 (simulationagent-1.0 9635) simulation.agent DEBUG: report_
↳time = 2017-02-02 14:30:00
2017-05-01 15:06:12,901 (simulationagent-1.0 9635) simulation.agent DEBUG:
↳Setting storage dispatch to 15.0 kW
2017-05-01 15:06:27,931 (simulationagent-1.0 9635) simulation.agent DEBUG: 2017-05-01_
↳15:06:27.907951 Reporting at sim time 2017-02-02 15:15:00
2017-05-01 15:06:27,931 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/
↳simload/power_kw = 474.2
2017-05-01 15:06:27,931 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/
↳simpv/power_kw = -11.7
2017-05-01 15:06:27,932 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/
↳simstorage/dispatch_kw = 15.0
2017-05-01 15:06:27,932 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/
↳simstorage/last_timestamp = 2017-02-02 15:03:00
2017-05-01 15:06:27,932 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/
↳simstorage/power_kw = 5.362
2017-05-01 15:06:27,932 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/
↳simstorage/soc_kwh = 48.033
2017-05-01 15:06:27,932 (simulationagent-1.0 9635) simulation.agent DEBUG: net_power_
↳kw = 467.862
2017-05-01 15:06:27,932 (simulationagent-1.0 9635) simulation.agent DEBUG: report_
↳time = 2017-02-02 15:15:00
2017-05-01 15:06:27,932 (simulationagent-1.0 9635) simulation.agent DEBUG:
↳Setting storage dispatch to -15.0 kW
2017-05-01 15:06:42,971 (simulationagent-1.0 9635) simulation.agent DEBUG: 2017-05-01_
↳15:06:42.939181 Reporting at sim time 2017-02-02 16:00:00
2017-05-01 15:06:42,971 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/
↳simload/power_kw = 469.5

```

(continues on next page)

(continued from previous page)

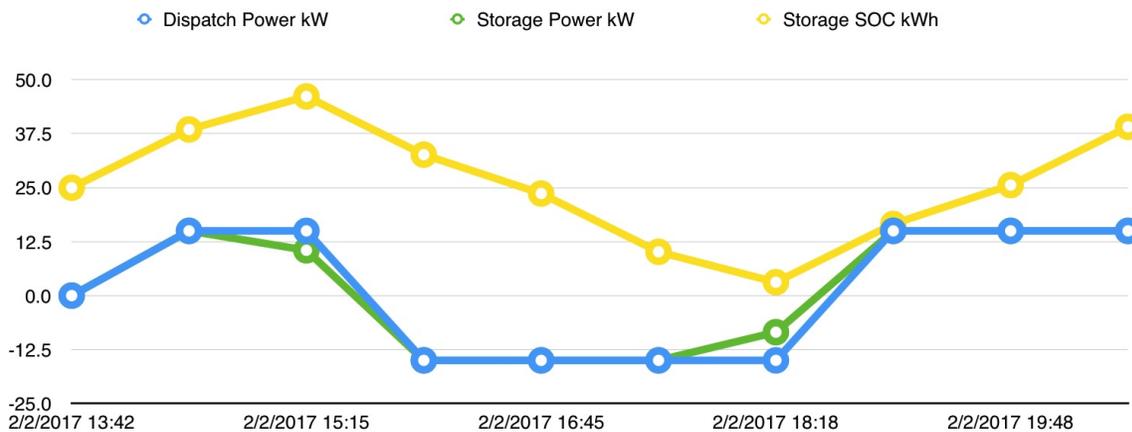
```
2017-05-01 15:06:42,971 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/  
↳simpv/power_kw = -9.375  
2017-05-01 15:06:42,971 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/  
↳simstorage/dispatch_kw = -15.0  
2017-05-01 15:06:42,971 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/  
↳simstorage/last_timestamp = 2017-02-02 15:57:00  
2017-05-01 15:06:42,971 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/  
↳simstorage/power_kw = -12.0  
2017-05-01 15:06:42,971 (simulationagent-1.0 9635) simulation.agent DEBUG: devices/  
↳simstorage/soc_kwh = 37.233  
2017-05-01 15:06:42,971 (simulationagent-1.0 9635) simulation.agent DEBUG: net_power_  
↳kw = 448.125  
2017-05-01 15:06:42,971 (simulationagent-1.0 9635) simulation.agent DEBUG: report_  
↳time = 2017-02-02 16:00:00  
2017-05-01 15:06:42,971 (simulationagent-1.0 9635) simulation.agent DEBUG:      ↳  
↳Setting storage dispatch to -15.0 kW  
2017-05-01 15:06:58,001 (simulationagent-1.0 9635) simulation.agent DEBUG: The_↳  
↳simulation has ended.
```

Report Output

The SimulationAgent also writes a CSV output file so that simulation results can be reported by spreadsheets, for example this graph of the simulated storage device following an oscillating dispatch:

Simulation Output

Time	Dispatch Power kW	Storage Power kW	Storage SOC kWh	PV Power kW	Net Power kW	Load Power kW
2/2/2017 13:42	0.0	0.0	25.0	-1.0	485.1	486.1
2/2/2017 14:30	15.0	15.0	38.5	-5.9	476.6	467.5
2/2/2017 15:15	15.0	10.5	46.2	-11.7	473.0	474.2
2/2/2017 16:00	-15.0	-15.0	32.7	-9.4	445.1	469.5
2/2/2017 16:45	-15.0	-15.0	23.7	-14.7	473.6	503.3
2/2/2017 17:33	-15.0	-15.0	10.1	-16.1	475.5	506.5
2/2/2017 18:18	-15.0	-8.5	3.1	-18.5	439.8	466.8
2/2/2017 19:03	15.0	15.0	16.6	-18.3	476.8	480.1
2/2/2017 19:48	15.0	15.0	25.6	-21.8	460.6	467.4
2/2/2017 20:33	15.0	15.0	39.1	-19.8	434.1	438.9



Using the Simulation Framework to Test a Driver

If you're developing a VOLTRON driver, and you intend to add it to the drivers managed by `MasterDriverAgent`, then with a few tweaks, you can adapt it so that it's testable from this simulation framework.

As with drivers under `MasterDriverAgent`, your driver should go in a `.py` module that implements a `Register` class and an `Interface` class. In order to work within the simulation framework, simulation drivers need to be adjusted as follows:

- Place the module in the `interfaces` directory under `SimulationDriverAgent`.
- The module's `Register` class should inherit from `SimulationRegister`.
- The module's `Interface` class should inherit from `SimulationInterface`.
- If the driver has logic that depends on time, get the simulated time by calling `self.sim_time()`.

Add files with your driver's config and point definitions, and load them into the config store:

```
$ volttron-ctl config store simulation.driver \
  yourdriver.csv \
  $VOLTRON_ROOT/applications/kisensum/Simulation/SimulationDriverAgent/yourdriver.
↪ csv --csv
$ volttron-ctl config store simulation.driver \
```

(continues on next page)

(continued from previous page)

```
devices/yourdriver \  
$VOLTRON_ROOT/applications/kisensum/Simulation/SimulationDriverAgent/yourdriver.  
↪config
```

To manage your driver from the SimulationAgent, first add the driver to the `sim_driver_list` in that agent's config:

```
"sim_driver_list": ["simload", "simpv", "simstorage", "yourdriver"]
```

Then, if you choose, you can also revise SimulationAgent's config and logic to scrape and report your driver's points, and/or send RPC requests to your driver.

For Further Information

If you have comments or questions about this simulation support, please contact Rob Calvert at Kisensum, Inc.:

- (github) @rob-calvert
- (email) rob@kisensum.com

3.7.2 VEN Agent: OpenADR 2.0b Interface Specification

OpenADR (Automated Demand Response) is a standard for alerting and responding to the need to adjust electric power consumption in response to fluctuations in grid demand.

OpenADR communications are conducted between Virtual Top Nodes (VTNs) and Virtual End Nodes (VENs). In this implementation a VOLTRON agent, VEN agent, acts as a VEN, communicating with its VTN by means of EIEvent and EIReport services in conformance with a subset of the OpenADR 2.0b specification. This document's "VOLTRON Interface" section defines how the VEN agent relays information to, and receives data from, other VOLTRON agents.

The OpenADR 2.0b specification (<http://www.openadr.org/specification>) is available from the OpenADR Alliance. This implementation also generally follows the DR program characteristics of the Capacity Program described in Section 9.2 of the OpenADR Program Guide (http://www.openadr.org/assets/openadr_drprogramguide_v1.0.pdf).

DR Capacity Bidding and Events

The OpenADR Capacity Bidding program relies on a pre-committed agreement about the VEN's load shed capacity. This agreement is reached in a bidding process transacted outside of the OpenADR interaction, typically with a long-term scope, perhaps a month or longer. The VTN can "call an event," indicating that a load-shed event should occur in conformance with this agreement. The VTN indicates the level of load shedding desired, when the event should occur, and for how long. The VEN responds with an "optIn" acknowledgment. (It can also "optOut," but since it has been pre-committed, an "optOut" may incur penalties.)

Reporting

The VEN agent reports device status and usage telemetry to the VTN, relying on information received periodically from other VOLTRON agents.

General Approach

Events:

- The VEN agent maintains a persistent record of DR events.
- Event updates (including creation) trigger publication of event JSON on the VOLTRON message bus.
- Other VOLTRON agents can also call a `get_events()` RPC to retrieve the current status of particular events, or of all active events.

Reporting:

- The VEN agent configuration defines telemetry values (data points) that can be reported to the VTN.
- The VEN agent maintains a persistent record of telemetry values over time.
- Other VOLTRON agents are expected to call `report_telemetry()` to supply the VEN agent with a regular stream of telemetry values for reporting.
- Other VOLTRON agents can receive notification of changes in telemetry reporting requirements by subscribing to publication of telemetry parameters.

VEN Agent VOLTRON Interface

The VEN agent implements the following VOLTRON PubSub and RPC calls.

PubSub: event update

```
.. code-block:: python
```

```
def publish_event(self, an_event):
```

```
    """ Publish an event.
```

```
    When an event is created/updated, it is published to the VOLTRON bus with a topic that includes 'openadr/event_update'.
```

```
    Event JSON structure:
```

```
    { "event_id" : String, "creation_time" : DateTime, "start_time" : DateTime, "end_time"
      : DateTime or None, "signals" : String, # Values: json string describing one or more
      signals. "status" : String, # Values: unresponded, far, near, active,
      # completed, canceled.
      "opt_type" : String # Values: optIn, optOut, none.
    }
```

```
    If an event status is 'unresponded', the VEN agent is awaiting a decision on whether to optIn or optOut. The downstream agent that subscribes to this PubSub message should communicate that choice to the VEN agent by calling respond_to_event() (see below). The VEN agent then relays the choice to the VTN.
```

```
    @param an_event: an EiEvent. """
```

PubSub: telemetry parameters update

```
.. code-block:: python
```

```
def publish_telemetry_parameters_for_report(self, report):
```

```
    """ Publish telemetry parameters.
```

```
    When the VEN agent telemetry reporting parameters have been updated (by the VTN), they are published with a topic that includes 'openadr/telemetry_parameters'. If a particular report has been updated, the reported parameters are for that report.
```

Telemetry parameters JSON example: {

```

“telemetry”: {
  “baseline_power_kw”: { “r_id”: “baseline_power”, “frequency”: “30”,
    “report_type”: “baseline”, “reading_type”: “Mean”, “method_name”:
    “get_baseline_power”
  } “current_power_kw”: {
    “r_id”: “actual_power”, “frequency”: “30”, “report_type”: “reading”, “read-
    ing_type”: “Mean”, “method_name”: “get_current_power”
  } “manual_override”: “False”, “report_status”: “active”, “online”: “False”,
}
}

```

The above example indicates that, for reporting purposes, telemetry values for `baseline_power` and `actual_power` should be updated – via `report_telemetry()` – at least once every 30 seconds.

Telemetry value definitions such as `baseline_power` and `actual_power` come from the agent configuration.

@param report: (EiReport) The report whose parameters should be published. “”“

RPC calls:

```

@RPC.export
def respond_to_event(self, event_id, opt_in_choice=None):
    """
    Respond to an event, opting in or opting out.

    If an event's status=unresponded, it is awaiting this call.
    When this RPC is received, the VENAgent sends an eventResponse to
    the VTN, indicating whether optIn or optOut has been chosen.
    If an event remains unresponded for a set period of time,
    it times out and automatically optsIn to the event.

    Since this call causes a change in the event's status, it triggers
    a PubSub call for the event update, as described above.

    @param event_id: (String) ID of an event.
    @param opt_in_choice: (String) 'OptIn' to opt into the event, anything else is_
    →treated as 'OptOut'.
    """

```

```

@RPC.export
def get_events(self, event_id=None, in_progress_only=True, started_after=None, end_
    →time_before=None):
    """
    Return a list of events as a JSON string.

    Sample request:
    self.get_events(started_after=utils.get_aware_utc_now() -_
    →timedelta(hours=1),
                                end_time_before=utils.get_aware_utc_now())

    Return a list of events.

```

(continues on next page)

(continued from previous page)

```

    By default, return only event requests with status=active or
    ↪status=unresponded.

```

```

    If an event's status=active, a DR event is currently in progress.

```

```

@param event_id: (String) Default None.
@param in_progress_only: (Boolean) Default True.
@param started_after: (DateTime) Default None.
@param end_time_before: (DateTime) Default None.
@return: (JSON) A list of events -- see 'PubSub: event update'.
"""

```

```
@RPC.export
```

```
def get_telemetry_parameters(self):
```

```
    """
```

```
        Return the VEN agent's current set of telemetry parameters.

```

```

    @return: (JSON) Current telemetry parameters -- see 'PubSub: telemetry parameters_
    ↪update'.
    """

```

```
@RPC.export
```

```
def set_telemetry_status(self, online, manual_override):
```

```
    """
```

```
        Update the VEN agent's reporting status.

```

```
        Set these properties to either 'TRUE' or 'FALSE'.

```

```

@param online: (Boolean) Whether the VEN agent's resource is online.
@param manual_override: (Boolean) Whether resource control has been overridden.
    """

```

```
@RPC.export
```

```
def report_telemetry(self, telemetry):
```

```
    """
```

```

    Receive an update of the VENAgent's report metrics, and store them in the_
    ↪agent's database.

```

```

    Examples of telemetry are:

```

```

    {
        'baseline_power_kw': '15.2',
        'current_power_kw': '371.1',
        'start_time': '2017-11-21T23:41:46.051405',
        'end_time': '2017-11-21T23:42:45.951405'
    }

```

```

    @param telemetry_values: (JSON) Current value of each report metric, with_
    ↪reporting-interval start/end.
    """

```

For Further Information

Please contact Rob Calvert at Kisensum, rob@kisensum.com

3.7.3 Reference Application

This reference application for VOLTRON's OpenADR Virtual End Node (VEN) and its Simulation Subsystem demonstrates interactions between the VOLTRON VEN agent and simulated devices. It employs a Virtual Top Node (VTN) server, demonstrating the full range of interaction and communication in a VOLTRON implementation of the OpenADR (Automated Demand Response) standard.

The simulation subsystem, described in more detail in *Simulated Subsystem*, includes a set of device simulators and a clock that can run faster (or slower) than real time (using ReferenceApp's default configuration, the clock runs at normal speed).

Eight VOLTRON agents work together to run this simulation:

1. **ReferenceAppAgent.** This agent configures, starts, and reports on a simulation. It furnishes a variety of configuration parameters to the other simulation agents, starts the clock, subscribes to scraped driver points, and generates a CSV output file. The ReferenceApp also serves as the mediator between the simulated device drivers and the VEN, adjusting driver behavior (particularly the behavior of the "simstorage" battery) while an OpenADR event is in progress, and aggregating and relaying relevant driver metrics to the VEN for reporting to the VTN.
2. **SimulationClockAgent.** This agent manages the simulation's clock. After it has been supplied with a start time, a stop time, and a clock-speed multiplier, and it has been asked to start a simulation, it provides the current simulated time in response to requests. If no stop time has been provided (this is the default behavior while the ReferenceApp is managing the clock), the SimulationClockAgent runs the simulation until the agent is stopped. If no clock-speed multiplier has been provided, the simulation clock runs at normal wallclock speed.
3. **SimulationDriverAgent.** Like MasterDriverAgent, this agent is a front-end manager for device drivers. It handles get_point/set_point requests from other agents, and it periodically "scrapes" and publishes each driver's points. If a device driver has been built to run under MasterDriverAgent, with a few minor modifications (detailed below) it can be adapted to run under SimulationDriverAgent.
4. **ActuatorAgent.** This agent manages write access to device drivers. Another agent may request a scheduled time period, called a Task, during which it controls a device.
5. **OpenADRVenAgent.** This agent implements an OpenADR Virtual End Node (VEN). It receives demand-response event notifications from a Virtual Top Node (VTN), making the event information available to the ReferenceAppAgent and other interested VOLTRON agents. It also reports metrics to the VTN based on information furnished by the ReferenceAppAgent.
6. **SQLHistorian.** This agent, a "platform historian," captures metrics reported by the simulated devices, storing them in a SQLite database.
7. **VolttronCentralPlatform.** This agent makes the platform historian's device metrics available for reporting by the VolttronCentralAgent.
8. **VolttronCentralAgent.** This agent manages a web user interface that can produce graphical displays of the simulated device metrics captured by the SQLHistorian.

Three simulated device drivers are used:

1. **storage (simstorage).** The storage driver simulates an energy storage device (i.e., a battery). When it receives a power dispatch value (positive to charge the battery, negative to discharge it), it adjusts the storage unit's charging behavior accordingly. Its reported power doesn't necessarily match the dispatch value, since (like an actual battery) it stays within configured max-charge/max-discharge limits, and power dwindles as its state of charge approaches a full or empty state.
2. **pv (simpv).** The PV driver simulates a photovoltaic array (solar panels), reporting the quantity of solar power produced. Solar power is calculated as a function of (simulated) time, using a data file of incident-sunlight metrics. A year's worth of solar data has been provided as a sample resource.

3. **load (simload).** The load driver simulates the behavior of a power consumer such as a building, reporting the quantity of power consumed. It gets its power metrics as a function of (simulated) time from a data file of power readings. A year's worth of building-load data has been provided as a sample resource.

Linux Installation

The following steps describe how to set up and run a simulation. They assume that the VOLTRON / `volttron` and VOLTRON / `volttron-applications` repositories have been downloaded from github.

Installing and running a simulation is walked through in the Jupyter notebook in `$(VOLTRON_ROOT)/examples/JupyterNotebooks/ReferenceAppAgent.ipynb`. In order to run this notebook, install Jupyter and start the Jupyter server:

```
$ cd $VOLTRON_ROOT
$ source env/bin/activate
$ pip install jupyter
$ jupyter notebook
```

By default, a browser will open with the Jupyter Notebook dashboard at <http://localhost:8888>. Run the notebook by navigating in the Jupyter Notebook dashboard to <http://localhost:8888/tree/examples/JupyterNotebooks/ReferenceAppAgent.ipynb>.

ReferenceAppAgent Configuration Parameters

This section describes ReferenceAppAgents's configurable parameters. Each of these has a default value and behavior, allowing the simulation to be run "out of the box" without configuring any parameters.

Type	Param Name	Data Type	Default
General	<code>agent_id</code>	str	<code>reference_app</code>
General	<code>heartbeat_period</code>	int sec	5
General	<code>sim_driver_list</code>	list of str	[<code>simload</code> , <code>simpv</code> , <code>simstorage</code>]
General	<code>opt_type</code>	str	<code>optIn</code>
General	<code>report_interval_secs</code>	int sec	30
General	<code>baseline_power_kw</code>	int kw	500
Clock	<code>sim_start</code>	datetime str	2017-04-30 13:00:00
Clock	<code>sim_end</code>	datetime str	None
Clock	<code>sim_speed</code>	float sec	1.0
Load	<code>load_timestamp_column_header</code>	str	<code>local_date</code>
Load	<code>load_power_column_header</code>	str	<code>load_kw</code>
Load	<code>load_data_frequency_min</code>	int min	15
Load	<code>load_data_year</code>	str	2015
Load	<code>load_csv_file_path</code>	str	<code>~/repos/volttron-applications/kisensum/ReferenceAppA</code>
PV	<code>pv_panel_area</code>	float m2	1000.0
PV	<code>pv_efficiency</code>	float 0.0-1.0	0.75
PV	<code>pv_data_frequency_min</code>	int min	30
PV	<code>pv_data_year</code>	str	2015
PV	<code>pv_csv_file_path</code>	str	<code>~/repos/volttron-applications/kisensum/ReferenceAppA</code>
Storage	<code>storage_soc_kwh</code>	float kWh	450.0
Storage	<code>storage_max_soc_kwh</code>	float kWh	500.0
Storage	<code>storage_max_charge_kw</code>	float kW	150.0
Storage	<code>storage_max_discharge_kw</code>	float kW	150.0

Type	Param Name	Data Type	Default
Storage	storage_reduced_charge_soc_threshold	float 0.0-1.0	0.80
Storage	storage_reduced_discharge_soc_threshold	float 0.0-1.0	0.20
Dispatch	positive_dispatch_kw	float kW >= 0.0	150.0
Dispatch	negative_dispatch_kw	float kW <= 0.0	-150.0
Dispatch	go_positive_if_below	float 0.0-1.0	0.1
Dispatch	go_negative_if_above	float 0.0-1.0	0.9
Report	report_interval	int seconds	15
Report	report_file_path	str	\$VOLTTRON_HOME/run/simulation_out.csv
Actuator	actuator_id	str	simulation.actuator
VEN	venagent_id	str	venagent

Driver Parameters and Points

Load Driver

The load driver's parameters specify how to look up power metrics in its data file.

Type	Name	Data Type	Default	Comments
Param/Point	csv_file_path	string		This parameter must be supplied by the agent.
Param/Point	times-tamp_column_header	string	local_date	
Param/Point	power_column_header	string	load_kw	
Param/Point	data_frequency_min	int	15	
Param/Point	data_year	string	2015	
Point	power_kw	float	0.0	
Point	last_timestamp	datetime		

PV Driver

The PV driver's parameters specify how to look up sunlight metrics in its data file, and how to calculate the power generated from that sunlight.

Type	Name	Data Type	Default	Comments
Param/Point	csv_file_path	string		This parameter must be supplied by the agent.
Param/Point	max_power_kw	float	10.0	
Param/Point	panel_area	float	50.0	
Param/Point	efficiency	float	0.75	
Param/Point	data_frequency_min	int	30	
Param/Point	data_year	string	2015	
Point	power_kw	float	0.0	
Point	last_timestamp	datetime		

Storage Driver

The storage driver's parameters describe the device's power and SOC limits, its initial SOC, and the SOC thresholds at which charging and discharging start to be reduced as its SOC approaches a full or empty state. This reduced power is calculated as a straight-line reduction: charging power is reduced in a straight line from `reduced_charge_soc_threshold` to 100% SOC, and discharging power is reduced in a straight line from `reduced_discharge_soc_threshold` to 0% SOC.

Type	Name	Data Type	Default	Comments
Param/Point	<code>max_charge_kw</code>	float	15.0	
Param/Point	<code>max_discharge_kw</code>	float	15.0	
Param/Point	<code>max_soc_kwh</code>	float	50.0	
Param/Point	<code>soc_kwh</code>	float	25.0	
Param/Point	<code>reduced_charge_soc_threshold</code>	float	0.8	
Param/Point	<code>reduced_discharge_soc_threshold</code>	float	0.2	
Point	<code>dispatch_kw</code>	float	0.0	
Point	<code>power_kw</code>	float	0.0	
Point	<code>last_timestamp</code>	datetime		

VEN Configuration

The VEN may be configured according to its documentation [here](#).

Running the Simulation

There are three main ways to monitor the ReferenceApp simulation's progress.

One way is to look at debug trace in VOLTRON's log output, for example:

```
2018-01-08 17:41:30,333 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: 2018-
↳ 01-08 17:41:30.333260 Initializing drivers
2018-01-08 17:41:30,333 (referenceappagent-1.0 23842) referenceapp.agent DEBUG:
↳ Initializing Load: timestamp_column_header=local_date, power_column_header=load
↳ kw, data_frequency_min=15, data_year=2015, csv_file_path=/home/ubuntu/repos/
↳ volttron-applications/kisensum/ReferenceAppAgent/data/load_and_pv.csv
2018-01-08 17:41:30,379 (referenceappagent-1.0 23842) referenceapp.agent DEBUG:
↳ Initializing PV: panel_area=50.0, efficiency=0.75, data_frequency_min=30, data_
↳ year=2015, csv_file_path=/home/ubuntu/repos/volttron-applications/kisensum/
↳ ReferenceAppAgent/data/nrel_pv_readings.csv
2018-01-08 17:41:30,423 (referenceappagent-1.0 23842) referenceapp.agent DEBUG:
↳ Initializing Storage: soc_kwh=25.0, max_soc_kwh=50.0, max_charge_kw=15.0, max_
↳ discharge_kw=15.0, reduced_charge_soc_threshold = 0.8, reduced_discharge_soc_
↳ threshold = 0.2
2018-01-08 17:41:32,331 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: 2018-
↳ 01-08 17:41:32.328390 Reporting at sim time 2018-01-08 17:41:31.328388
2018-01-08 17:41:32,331 (referenceappagent-1.0 23842) referenceapp.agent DEBUG:
↳ net_power_kw = 0
2018-01-08 17:41:32,331 (referenceappagent-1.0 23842) referenceapp.agent DEBUG:
↳ report_time = 2018-01-08 17:41:31.328388
2018-01-08 17:41:32,338 (referenceappagent-1.0 23842) referenceapp.agent DEBUG:
↳ Setting storage dispatch to 15.0 kW
2018-01-08 17:41:46,577 (referenceappagent-1.0 23842) referenceapp.agent DEBUG:
↳ Received event: ID=4, status=far, start=2017-12-01 18:40:55+00:00, end=2017-12-02
↳ 18:37:56+00:00, opt_type=none, all params={"status": "far", "signals": "{}\null\":
↳ {"intervals\": {"0\": {"duration\": "\PT23H57M1S\","uid\": "\0\","payloads\
↳ ": {}}, "currentLevel\": null, "signalID\": null}}, "event_id": "4", "start_time
↳ ": "2017-12-01 18:40:55+00:00", "creation_time": "2018-01-08 17:41:45.774548", "opt
↳ type": "none", "priority": 1, "end_time": "2017-12-02 18:37:56+00:00"}
```

(continues on next page)

(continued from previous page)

```

2018-01-08 17:41:46,577 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ Sending an optIn response for event ID 4
2018-01-08 17:41:46,583 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: 2018-
↳ 01-08 17:41:46.576130 Reporting at sim time 2018-01-08 17:41:46.328388
2018-01-08 17:41:46,583 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ devices/simload/power_kw = 519.3
2018-01-08 17:41:46,583 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ devices/simpv/power_kw = -17.175
2018-01-08 17:41:46,583 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ devices/simstorage/dispatch_kw = 15.0
2018-01-08 17:41:46,584 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ devices/simstorage/power_kw = 15.0
2018-01-08 17:41:46,584 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ devices/simstorage/soc_kwh = 25.025
2018-01-08 17:41:46,584 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ net_power_kw = 49.755
2018-01-08 17:41:46,584 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ report_time = 2018-01-08 17:41:46.328388
2018-01-08 17:41:46,596 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ Setting storage dispatch to 15.0 kW
2018-01-08 17:41:48,617 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ Received event: ID=4, status=completed, start=2017-12-01 18:40:55+00:00, end=2017-
↳ 12-02 18:37:56+00:00, opt_type=optIn, all params={"status": "completed", "signals":
↳ {"\"null\": {\"intervals\": {\"0\": {\"duration\": \"PT23H57M1S\", \"uid\": \"0\", \
↳ \"payloads\": {}}}, \"currentLevel\": null, \"signalID\": null}}, \"event_id\": \"4\",
↳ \"start_time\": \"2017-12-01 18:40:55+00:00\", \"creation_time\": \"2018-01-08 17:41:45.
↳ 774548\", \"opt_type\": \"optIn\", \"priority\": 1, \"end_time\": \"2017-12-02 18:37:56+00:00
↳ }"}
2018-01-08 17:42:59,563 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: 2018-
↳ 01-08 17:42:59.559264 Reporting at sim time 2018-01-08 17:42:59.328388
2018-01-08 17:42:59,563 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ devices/simload/power_kw = 519.3
2018-01-08 17:42:59,563 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ devices/simpv/power_kw = -17.175
2018-01-08 17:42:59,563 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ devices/simstorage/dispatch_kw = 15.0
2018-01-08 17:42:59,563 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ devices/simstorage/power_kw = 15.0
2018-01-08 17:42:59,563 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ devices/simstorage/soc_kwh = 25.238
2018-01-08 17:42:59,563 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ net_power_kw = 49.755
2018-01-08 17:42:59,563 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ report_time = 2018-01-08 17:42:59.328388
2018-01-08 17:42:59,578 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ Setting storage dispatch to -1.05158333333 kW
2018-01-08 17:43:01,596 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: 2018-
↳ 01-08 17:43:01.589877 Reporting at sim time 2018-01-08 17:43:01.328388
2018-01-08 17:43:01,596 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ devices/simload/power_kw = 519.3
2018-01-08 17:43:01,596 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ devices/simpv/power_kw = -17.175
2018-01-08 17:43:01,597 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ devices/simstorage/dispatch_kw = -1.05158333333
2018-01-08 17:43:01,597 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ devices/simstorage/power_kw = -1.051
2018-01-08 17:43:01,597 (referenceappagent-1.0 23842) referenceapp.agent DEBUG: ↵
↳ devices/simstorage/soc_kwh = 25.236

```

(continues on next page)

(continued from previous page)

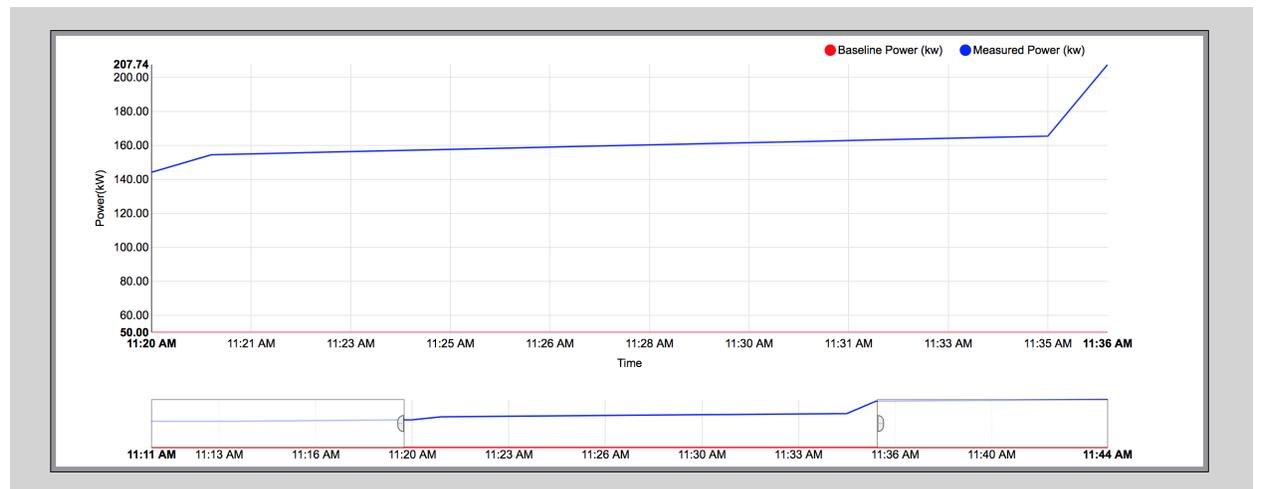
```

2018-01-08 17:43:01,597 (referenceappagent-1.0 23842) referenceapp.agent DEBUG:
↳ net_power_kw = 33.704
2018-01-08 17:43:01,597 (referenceappagent-1.0 23842) referenceapp.agent DEBUG:
↳ report_time = 2018-01-08 17:43:01.328388
2018-01-08 17:43:01,598 (referenceappagent-1.0 23842) referenceapp.agent DEBUG:
↳ Reporting telemetry: {'start_time': '2018-01-08 17:42:31.598889+00:00', 'baseline_
↳ power_kw': '50', 'current_power_kw': '33.704', 'end_time': '2018-01-08 17:43:01.
↳ 598889+00:00'}
2018-01-08 17:43:01,611 (referenceappagent-1.0 23842) referenceapp.agent DEBUG:
↳ Setting storage dispatch to -1.0515 kW
    
```

Another way to monitor progress is to launch the VolttronCentral web UI, which can be found at <http://127.0.0.1:8080/vc/index.html>. Here, in addition to checking agent status, one can track metrics reported by the simulated device drivers. For example, these graphs track the simstorage battery’s power consumption and state of charge over time. The abrupt shift from charging to discharging happens because an OpenADR event has just started:



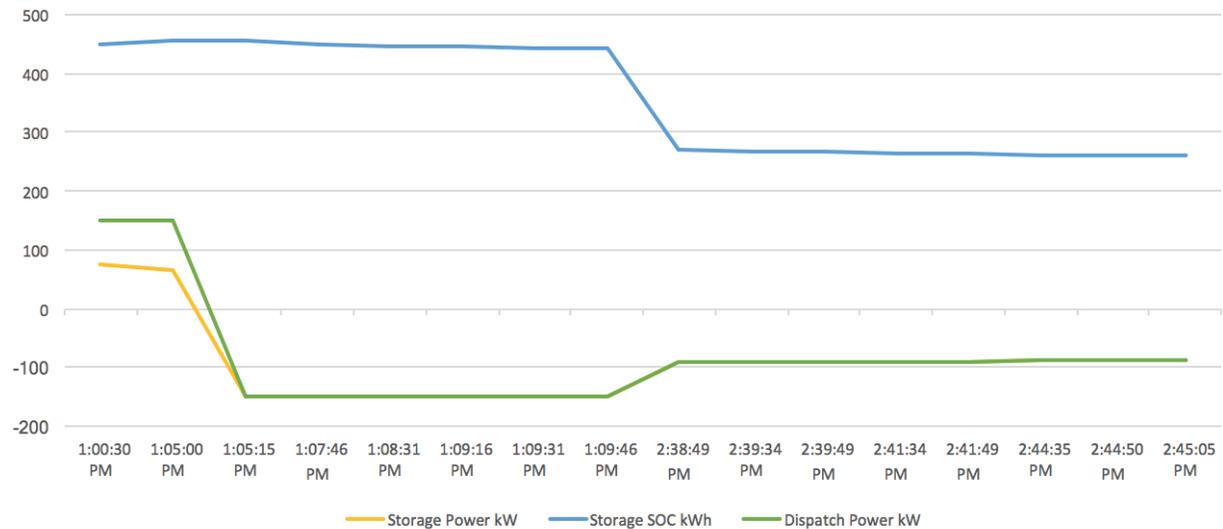
A third way to monitor progress, while there is an active DR event, is to examine the event’s graph in the VTN web UI. This displays the VEN’s power consumption, which is an aggregate of the consumption reported by each simulated device driver:



Report Output

The ReferenceAppAgent also writes a CSV output file so that simulation results can be reported in a spreadsheet, for example this graph of the simulated storage device:

Simulation Output							
Time	Net Power kW	Load Power kW	PV Power kW	Storage Power kW	Storage SOC kWh	Dispatch Power kW	
1:00:30 PM	364.312	460.5	-171	74.812	450.249	150	
1:05:00 PM	356.374	460.5	-171	66.874	455.528	150	
1:05:15 PM	139.5	460.5	-171	-150	454.778	-150	
1:07:46 PM	139.5	460.5	-171	-150	448.528	-150	
1:08:31 PM	139.561	460.5	-171	-149.939	446.778	-149.9395474	
1:09:16 PM	140.148	460.5	-171	-149.352	444.781	-149.3529006	
1:09:31 PM	140.398	460.5	-171	-149.102	444.283	-149.1024863	
1:09:46 PM	140.565	460.5	-171	-148.935	443.536	-148.9355433	
2:38:49 PM	173.918	463.8	-199.5	-90.382	269.315	-90.38282894	
2:39:34 PM	174.272	463.8	-199.5	-90.028	268.107	-90.02849427	
2:39:49 PM	174.424	463.8	-199.5	-89.876	267.807	-89.87663656	
2:41:34 PM	175.276	463.8	-199.5	-89.024	265.117	-89.02415495	
2:41:49 PM	175.426	463.8	-199.5	-88.874	264.819	-88.87430859	
2:44:35 PM	176.812	463.8	-199.5	-87.488	260.692	-87.48881646	
2:44:50 PM	176.91	463.8	-199.5	-87.39	260.254	-87.39093025	
2:45:05 PM	190.256	477	-199.5	-87.244	259.962	-87.24410094	



For Further Information

If you have comments or questions about this simulation support, please contact Rob Calvert or Nate Hill at Kisensum, Inc.:

- (github) @rob-calvert
- (email) rob@kisensum.com
- (github) @hillnate
- (github) nate@kisensum.com

3.8 Indices and tables

- [genindex](#)
- [modindex](#)
- [search](#)