
BigchainDB Documentation

Release 0.6.0

BigchainDB Contributors

September 29, 2016

1 Table of Contents	1
HTTP Routing Table	87

Table of Contents

1.1 Introduction

BigchainDB is a scalable blockchain database. That is, it's a “big data” database with some blockchain characteristics, including decentralization, immutability and native support for assets.

You can read about the motivations, goals and high-level architecture in the [BigchainDB whitepaper](#).

1.1.1 Is BigchainDB Production-Ready?

No, BigchainDB is not production-ready. You can use it to build a prototype or proof-of-concept (POC); many people are already doing that.

BigchainDB is currently in version 0.X. ([The Releases page on GitHub](#) has the exact version number.) Once we believe that BigchainDB is production-ready, we'll release version 1.0.

[The BigchainDB Roadmap](#) will give you a sense of the things we intend to do with BigchainDB in the near term and the long term.

1.1.2 Some Basic Vocabulary

There is some specialized vocabulary associated with BigchainDB. To get started, you should at least know what we mean by a BigchainDB *node*, *cluster* and *federation*.

A **BigchainDB node** is a machine or set of closely-linked machines running RethinkDB Server, BigchainDB Server, and related software. (A “machine” might be a bare-metal server, a virtual machine or a container.) Each node is controlled by one person or organization.

A set of BigchainDB nodes can connect to each other to form a **cluster**. Each node in the cluster runs the same software. A cluster contains one logical RethinkDB datastore. A cluster may have additional machines to do things such as cluster monitoring.

The people and organizations that run the nodes in a cluster belong to a **federation** (i.e. another organization). A federation must have some sort of governance structure to make decisions. If a cluster is run by a single company, then the federation is just that company.

What's the Difference Between a Cluster and a Federation?

A cluster is just a bunch of connected nodes. A federation is an organization which has a cluster, and where each node in the cluster has a different operator. Confusingly, we sometimes call a federation's cluster its “federation.” You can probably tell what we mean from context.

There are several kinds of nodes:

- A **dev/test node** is a node created by a developer working on BigchainDB Server, e.g. for testing new or changed code. A dev/test node is typically run on the developer's local machine.
- A **bare-bones node** is a node deployed in the cloud, either as part of a testing cluster or as a starting point before upgrading the node to be production-ready. Our cloud deployment starter templates deploy a bare-bones node, as do our scripts for deploying a testing cluster on AWS.
- A **production node** is a node that is part of a federation's BigchainDB cluster. A production node has the most components and requirements.

1.1.3 Setup Instructions for Various Cases

- Set up a local stand-alone BigchainDB node for learning and experimenting: Quickstart
- Set up and run a bare-bones node in the cloud
- Set up and run a local dev/test node for developing and testing BigchainDB Server
- Deploy a testing cluster on AWS
- Set up and run a federation (including production nodes)

Instructions for setting up a client will be provided once there's a public test net.

1.1.4 Can I Help?

Yes! BigchainDB is an open-source project; we welcome contributions of all kinds. If you want to request a feature, file a bug report, make a pull request, or help in some other way, please see [the CONTRIBUTING.md file](#).

1.2 Quickstart

This page has instructions to set up a single stand-alone BigchainDB node for learning or experimenting. Instructions for other cases are elsewhere. We will assume you're using Ubuntu 14.04 or similar. If you're not using Linux, then you might try running BigchainDB with Docker.

A. [Install RethinkDB Server](#)

B. Open a Terminal and run RethinkDB Server with the command:

```
rethinkdb
```

C. Ubuntu 14.04 already has Python 3.4, so you don't need to install it, but you do need to install a couple other things:

```
sudo apt-get update
sudo apt-get install g++ python3-dev
```

D. Get the latest version of pip and setuptools:

```
sudo apt-get install python3-pip
sudo pip3 install --upgrade pip setuptools
```

E. Install the bigchaindb Python package from PyPI:

```
sudo pip3 install bigchaindb
```

F. Configure and run BigchainDB:

```
bigchaindb -y configure
bigchaindb start
```

That's it!

For now, you can get a good sense of how to work with BigchainDB by going through the examples in the section on the Python Server API.

1.3 Cloud Deployment Starter Templates

We have some “starter templates” to deploy a basic, working, but bare-bones BigchainDB node on various cloud providers. They should *not* be used as-is to deploy a node for production. They can be used as a starting point. A full production node should meet the requirements outlined in the section on [production node assumptions, components and requirements](#).

You don't have to use the tools we use in the starter templates (e.g. Terraform and Ansible). You can use whatever tools you prefer.

If you find the cloud deployment starter templates for nodes helpful, then you may also be interested in our scripts for [deploying a testing cluster on AWS](#) (documented in the Clusters & Federations section).

1.3.1 Template: Node Deployment on AWS

If you didn't read the introduction to the cloud deployment starter templates, please do that now. The main point is that they're not for deploying a production node; they can be used as a starting point.

The template documented on this page uses:

- [Terraform](#) to provision infrastructure resources on AWS, and
- [Ansible](#) to manage the software and files on that infrastructure (configuration management).

Install Terraform

The [Terraform documentation](#) has [installation instructions](#) for all common operating systems.

If you don't want to run Terraform on your local machine, you can install it on a cloud machine under your control (e.g. on AWS).

Note: Hashicorp has an enterprise version of Terraform called “Terraform Enterprise.” You can license it by itself or get it as part of Atlas. If you decide to license Terraform Enterprise or Atlas, be sure to install it on your own hosting (i.e. “on premise”), not on the hosting provided by Hashicorp. The reason is that BigchainDB clusters are supposed to be decentralized. If everyone used Hashicorp's hosted Atlas, then that would be a point of centralization.

Ubuntu Installation Tips

If you want to install Terraform on Ubuntu, first [download the .zip file](#). Then install it in /opt:

```
sudo mkdir -p /opt/terraform
sudo unzip path/to/zip-file.zip -d /opt/terraform
```

Why install it in /opt? See [the answers at Ask Ubuntu](#).

Next, add /opt/terraform to your path. If you use bash for your shell, then you could add this line to ~/.bashrc:

```
export PATH="/opt/terraform:$PATH"
```

After doing that, relaunch your shell or force it to read `~/.bashrc` again, e.g. by doing `source ~/.bashrc`. You can verify that terraform is installed and in your path by doing:

```
terraform --version
```

It should say the current version of Terraform.

Use Terraform to Provision a One-Machine Node on AWS

We have an example Terraform configuration (set of files) to provision all the resources needed to run a one-machine BigchainDB node on AWS:

- An instance on EC2 (based on an Ubuntu 14.04 AMI)
- A security group
- An EBS volume
- An elastic IP address

Get Set Up to Use Terraform

First, do the basic AWS setup steps outlined in the Appendices.

Then go to the `.../bigchaindb/ntools/one-m/aws/` directory and open the file `variables.tf`. Most of the variables have sensible default values, but you can change them if you like. In particular, you may want to change `aws_region`. (Terraform looks in `~/.aws/credentials` to get your AWS credentials, so you don't have to enter those anywhere.)

The `ssh_key_name` has no default value, so Terraform will prompt you every time it needs it.

To see what Terraform will do, run:

```
terraform plan
```

It should ask you the value of `ssh_key_name`.

It figured out the plan by reading all the `.tf` Terraform files in the directory.

If you don't want to be asked for the `ssh_key_name`, you can change the default value of `ssh_key_name` (in the file `variables.tf`) or you can set an environment variable named `TF_VAR_ssh_key_name`.

Use Terraform to Provision Resources

To provision all the resources specified in the plan, do the following. **Note: This will provision actual resources on AWS, and those cost money. Be sure to shut down the resources you don't want to keep running later, otherwise the cost will keep growing.**

```
terraform apply
```

Terraform will report its progress as it provisions all the resources. Once it's done, you can go to the Amazon EC2 web console and see the instance, its security group, its elastic IP, and its attached storage volumes (one for the root directory and one for RethinkDB storage).

At this point, there is no software installed on the instance except for Ubuntu 14.04 and whatever else came with the Amazon Machine Image (AMI) specified in the Terraform configuration (files). The next step is to use Ansible to install, configure and run all the necessary software.

Optional: “Destroy” the Resources

If you want to shut down all the resources just provisioned, you must first disable termination protection on the instance:

1. Go to the EC2 console and select the instance you just launched. It should be named `BigchainDB_node`.
2. Click **Actions > Instance Settings > Change Termination Protection > Yes, Disable**
3. Back in your terminal, do `terraform destroy`

Terraform should “destroy” (i.e. terminate or delete) all the AWS resources you provisioned above.

If it fails (e.g. because of an attached and mounted EBS volume), then you can terminate the instance using the EC2 console: **Actions > Instance State > Terminate > Yes, Terminate**. Once the instance is terminated, you should still do `terraform destroy` to make sure that all the other resources are destroyed.

Install Ansible

The Ansible documentation has [installation instructions](#). Note the control machine requirements: at the time of writing, Ansible required Python 2.6 or 2.7. (Support for Python 3 is a goal of [Ansible 2.2](#).)

For example, you could create a special Python 2.x virtualenv named `ansenv` and then install Ansible in it:

```
cd repos/bigchaindb/ntools
virtualenv -p /usr/local/lib/python2.7.11/bin/python ansenv
source ansenv/bin/activate
pip install ansible
```

About Our Example Ansible Playbook

We have an example Ansible playbook to install, configure and run a basic BigchainDB node on Ubuntu 14.04. It’s in `.../bigchaindb/ntools/one-m/ansible/one-m-node.yml`.

When you run the playbook (as per the instructions below), it ensures all the necessary software is installed, configured and running. It can be used to get a BigchainDB node set up on a bare Ubuntu 14.04 machine, but it can also be used to ensure that everything is okay on a running BigchainDB node. (If you run the playbook against a host where everything is okay, then it won’t change anything on that host.)

Create an Ansible Inventory File

An Ansible “inventory” file is a file which lists all the hosts (machines) you want to manage using Ansible. (Ansible will communicate with them via SSH.) Right now, we only want to manage one host.

First, determine the public IP address of the host (i.e. something like `192.0.2.128`).

Then create a one-line text file named `hosts` by doing this:

```
# cd to the directory .../bigchaindb/ntools/one-m/ansible
echo "192.0.2.128" > hosts
```

but replace `192.0.2.128` with the IP address of the host.

Run the Ansible Playbook

The next step is to run the Ansible playbook named `one-m-node.yml`:

```
# cd to the directory ../bigchaindb/ntools/one-m/ansible
ansible-playbook -i hosts --private-key ~/.ssh/<key-name> one-m-node.yml
```

where `<key-name>` should be replaced by the name of the SSH private key you created earlier (for SSHing to the host machine at your cloud hosting provider).

What did you just do? Running that playbook ensures all the software necessary for a one-machine BigchainDB node is installed, configured, and running properly. You can run that playbook on a regular schedule to ensure that the system stays properly configured. If something is okay, it does nothing; it only takes action when something is not as-desired.

Some Notes on the One-Machine Node You Just Got Running

- It ensures that the installed version of RethinkDB is `2.3.4~0trusty`. You can change that by changing the installation task.
- It uses a very basic RethinkDB configuration file based on `bigchaindb/ntools/one-m/ansible/roles/rethinkdb/`
- If you edit the RethinkDB configuration file, then running the Ansible playbook will **not** restart RethinkDB for you. You must do that manually. (You could just stop it using `sudo killall -9 rethinkdb` and then run the playbook to get it started again.)
- It generates and uses a default BigchainDB configuration file, which is stores in `~/.bigchaindb` (the default location).
- If you edit the BigchainDB configuration file, then running the Ansible playbook will ****not*** restart BigchainDB for you. You must do that manually. (You could just stop it using `sudo killall -9 bigchaindb` and then run the playbook to get it started again.)

Optional: Create an Ansible Config File

The above command (`ansible-playbook -i ...`) is fairly long. You can omit the optional arguments if you put their values in an [Ansible configuration file](#) (config file) instead. There are many places where you can put a config file, but to make one specifically for the “one-m” case, you should put it in `../bigchaindb/ntools/one-m/ansible/`. In that directory, create a file named `ansible.cfg` with the following contents:

```
[defaults]
private_key_file = $HOME/.ssh/<key-name>
inventory = hosts
```

where, as before, `<key-name>` must be replaced.

Next Steps

You could make changes to the Terraform configuration and the Ansible playbook to make the node more production-worthy. See the section on production node assumptions, components and requirements.

1.3.2 Template: Node Deployment on Azure

This is a placeholder.

1.4 Production Node Assumptions, Components & Requirements

1.4.1 Production Node Assumptions

If you're not sure what we mean by a BigchainDB *node*, *cluster*, *federation*, or *production node*, then see the section in the Introduction where we defined those terms.

We make some assumptions about production nodes:

1. **Each production node is set up and managed by an experienced professional system administrator (or a team of them).**
2. Each production node in a federation's cluster is managed by a different person or team.

Because of the first assumption, we don't provide a detailed cookbook explaining how to secure a server, or other things that a sysadmin should know. (We do provide some starter templates, but those are just a starting point.)

1.4.2 Production Node Components

A BigchainDB node must include, at least:

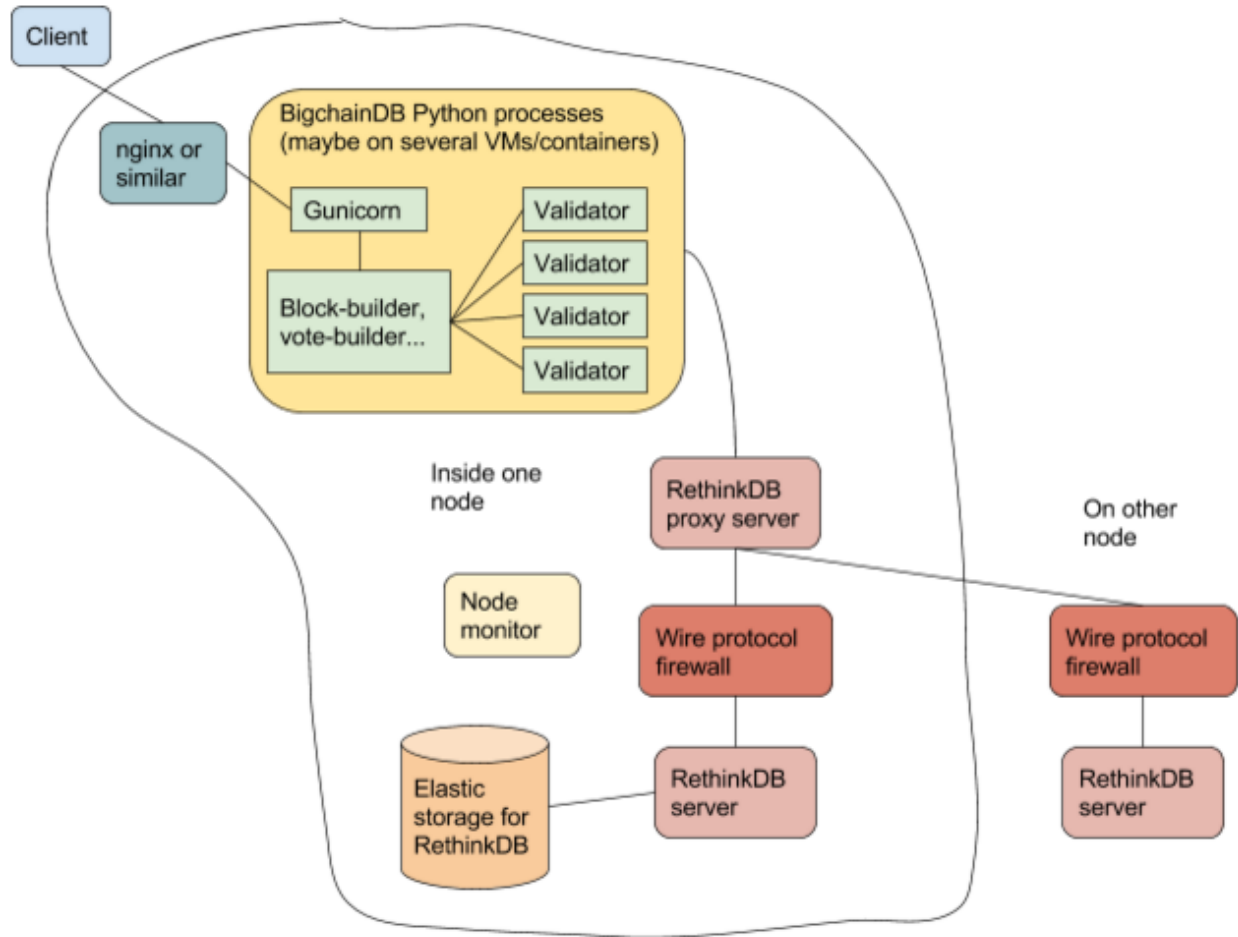
- BigchainDB Server and
- RethinkDB Server.

When doing development and testing, it's common to install both on the same machine, but in a production environment, it may make more sense to install them on separate machines.

In a production environment, a BigchainDB node should have several other components, including:

- nginx or similar, as a reverse proxy and/or load balancer for the Gunicorn server(s) inside the node
- An NTP daemon running on all machines running BigchainDB code, and possibly other machines
- A RethinkDB proxy server
- A RethinkDB "wire protocol firewall" (in the future: this component doesn't exist yet)
- Scalable storage for RethinkDB (e.g. using RAID)
- Monitoring software, to monitor all the machines in the node
- Configuration management agents (if you're using a configuration management system that uses agents)
- Maybe more

The relationship between these components is illustrated below.



1.4.3 Production Node Requirements

Note: This section will be broken apart into several pages, e.g. NTP requirements, RethinkDB requirements, BigchainDB requirements, etc. and those pages will add more details.

OS Requirements

- RethinkDB Server [will run on any modern OS](#). Note that the Fedora package isn't officially supported. Also, official support for Windows is fairly recent ([April 2016](#)).
- BigchainDB Server requires Python 3.4+ and Python 3.4+ [will run on any modern OS](#).
- BigchainDB Server uses the Python multiprocessing package and [some functionality in the multiprocessing package doesn't work on OS X](#). You can still use Mac OS X if you use Docker or a virtual machine.

The BigchainDB core dev team uses Ubuntu 14.04, Ubuntu 16.04, Fedora 23, and Fedora 24.

We don't test BigchainDB on Windows or Mac OS X, but you can try.

- If you run into problems on Windows, then you may want to try using Vagrant. One of our community members ([@Mec-Is](#)) wrote a [page about how to install BigchainDB on a VM with Vagrant](#).
- If you have Mac OS X and want to experiment with BigchainDB, then you could do that using Docker.

Storage Requirements

When it comes to storage for RethinkDB, there are many things that are nice to have (e.g. SSDs, high-speed input/output [IOPS], replication, reliability, scalability, pay-for-what-you-use), but there are few *requirements* other than:

1. have enough storage to store all your data (and its replicas), and
2. make sure your storage solution (hardware and interconnects) can handle your expected read & write rates.

For RethinkDB’s failover mechanisms to work, [every RethinkDB table must have at least three replicas](#) (i.e. a primary replica and two others). For example, if you want to store 10 GB of unique data, then you need at least 30 GB of storage. (Indexes and internal metadata are stored in RAM.)

As for the read & write rates, what do you expect those to be for your situation? It’s not enough for the storage system alone to handle those rates: the interconnects between the nodes must also be able to handle them.

Memory (RAM) Requirements

In their [FAQ](#), RethinkDB recommends that, “RethinkDB servers have at least 2GB of RAM...” ([source](#))

In particular: “RethinkDB requires data structures in RAM on each server proportional to the size of the data on that server’s disk, usually around 1% of the size of the total data set.” ([source](#)) We asked what they meant by “total data set” and [they said](#) it’s “referring to only the data stored on the particular server.”

Also, “The storage engine is used in conjunction with a custom, B-Tree-aware caching engine which allows file sizes many orders of magnitude greater than the amount of available memory. RethinkDB can operate on a terabyte of data with about ten gigabytes of free RAM.” ([source](#)) (In this case, it’s the *cluster* which has a total of one terabyte of data, and it’s the *cluster* which has a total of ten gigabytes of RAM. That is, if you add up the RethinkDB RAM on all the servers, it’s ten gigabytes.)

In reponse to our questions about RAM requirements, [@danielmewes](#) (of RethinkDB) [wrote](#):

... If you replicate the data, the amount of data per server increases accordingly, because multiple copies of the same data will be held by different servers in the cluster.

For example, if you increase the data replication factor from 1 to 2 (i.e. the primary plus one copy), then that will double the RAM needed for metadata. Also from [@danielmewes](#):

For reasonable performance, you should probably aim at something closer to 5-10% of the data size. [Emphasis added] The 1% is the bare minimum and doesn’t include any caching. If you want to run near the minimum, you’ll also need to manually lower RethinkDB’s cache size through the `--cache-size` parameter to free up enough RAM for the metadata overhead...

RethinkDB has [documentation about its memory requirements](#). You can use that page to get a better estimate of how much memory you’ll need. In particular, note that RethinkDB automatically configures the cache size limit to be about half the available memory, but it can be no lower than 100 MB. As [@danielmewes](#) noted, you can manually change the cache size limit (e.g. to free up RAM for queries, metadata, or other things).

If a RethinkDB process (on a server) runs out of RAM, the operating system will start swapping RAM out to disk, slowing everything down. According to [@danielmewes](#):

Going into swap is usually pretty bad for RethinkDB, and RethinkDB servers that have gone into swap often become so slow that other nodes in the cluster consider them unavailable and terminate the connection to them. I recommend adjusting RethinkDB’s cache size conservatively to avoid this scenario. RethinkDB will still make use of additional RAM through the operating system’s block cache (though less efficiently than when it can keep data in its own cache).

Filesystem Requirements

RethinkDB “supports most commonly used file systems” ([source](#)) but it has [issues with BTRFS](#) (B-tree file system).

It’s best to use a filesystem that supports direct I/O, because that will improve RethinkDB performance (if you tell RethinkDB to use direct I/O). Many compressed or encrypted filesystems don’t support direct I/O.

1.4.4 Set Up and Run a Cluster Node

This is a page of general guidelines for setting up a production node. It says nothing about how to upgrade software, storage, processing, etc. or other details of node management. It will be expanded more in the future.

Get a Server

The first step is to get a server (or equivalent) which meets the requirements for a BigchainDB node.

Secure Your Server

The steps that you must take to secure your server depend on your server OS and where your server is physically located. There are many articles and books about how to secure a server. Here we just cover special considerations when securing a BigchainDB node.

There are some notes on BigchainDB-specific firewall setup in the Appendices.

Sync Your System Clock

A BigchainDB node uses its system clock to generate timestamps for blocks and votes, so that clock should be kept in sync with some standard clock(s). The standard way to do that is to run an NTP daemon (Network Time Protocol daemon) on the node. (You could also use `tlsdate`, which uses TLS timestamps rather than NTP, but don’t: it’s not very accurate and it will break with TLS 1.3, which removes the timestamp.)

NTP is a standard protocol. There are many NTP daemons implementing it. We don’t recommend a particular one. On the contrary, we recommend that different nodes in a federation run different NTP daemons, so that a problem with one daemon won’t affect all nodes.

Please see the notes on NTP daemon setup in the Appendices.

Set Up Storage for RethinkDB Data

Below are some things to consider when setting up storage for the RethinkDB data. The Appendices have a section with concrete examples.

We suggest you set up a separate storage “device” (partition, RAID array, or logical volume) to store the RethinkDB data. Here are some questions to ask:

- How easy will it be to add storage in the future? Will I have to shut down my server?
- How big can the storage get? (Remember that [RAID](#) can be used to make several physical drives look like one.)
- How fast can it read & write data? How many input/output operations per second (IOPS)?
- How does IOPS scale as more physical hard drives are added?
- What’s the latency?
- What’s the reliability? Is there replication?

- What's in the Service Level Agreement (SLA), if applicable?
- What's the cost?

There are many options and tradeoffs. Don't forget to look into Amazon Elastic Block Store (EBS) and Amazon Elastic File System (EFS), or their equivalents from other providers.

Storage Notes Specific to RethinkDB

- The RethinkDB storage engine has a number of SSD optimizations, so you *can* benefit from using SSDs. ([source](#))
- If you want a RethinkDB cluster to store an amount of data D , with a replication factor of R (on every table), and the cluster has N nodes, then each node will need to be able to store $R \times D / N$ data.
- RethinkDB tables can have [at most 64 shards](#). For example, if you have only one table and more than 64 nodes, some nodes won't have the primary of any shard, i.e. they will have replicas only. In other words, once you pass 64 nodes, adding more nodes won't provide more storage space for new data. If the biggest single-node storage available is d , then the most you can store in a RethinkDB cluster is $< 64 \times d$: accomplished by putting one primary shard in each of 64 nodes, with all replica shards on other nodes. (This is assuming one table. If there are T tables, then the most you can store is $< 64 \times d \times T$.)
- When you set up storage for your RethinkDB data, you may have to select a filesystem. (Sometimes, the filesystem is already decided by the choice of storage.) We recommend using a filesystem that supports direct I/O (Input/Output). Many compressed or encrypted file systems don't support direct I/O. The ext4 filesystem supports direct I/O (but be careful: if you enable the `data=journal` mode, then direct I/O support will be disabled; the default is `data=ordered`). If your chosen filesystem supports direct I/O and you're using Linux, then you don't need to do anything to request or enable direct I/O. RethinkDB does that.
- RethinkDB stores its data in a specific directory. You can tell RethinkDB *which* directory using the RethinkDB config file, as explained below. In this documentation, we assume the directory is `/data`. If you set up a separate device (partition, RAID array, or logical volume) to store the RethinkDB data, then mount that device on `/data`.

Install RethinkDB Server

If you don't already have RethinkDB Server installed, you must install it. The RethinkDB documentation has instructions for [how to install RethinkDB Server on a variety of operating systems](#).

Configure RethinkDB Server

Create a RethinkDB configuration file (text file) named `instance1.conf` with the following contents (explained below):

```
directory=/data
bind=all
direct-io
# Replace node?_hostname with actual node hostnames below, e.g. rdb.examples.com
join=node0_hostname:29015
join=node1_hostname:29015
join=node2_hostname:29015
# continue until there's a join= line for each node in the federation
```

- `directory=/data` tells the RethinkDB node to store its share of the database data in `/data`.
- `bind=all` binds RethinkDB to all local network interfaces (e.g. loopback, Ethernet, wireless, whatever is available), so it can communicate with the outside world. (The default is to bind only to local interfaces.)

- `direct-io` tells RethinkDB to use direct I/O (explained earlier). Only include this line if your file system supports direct I/O.
- `join=hostname:29015` lines: A cluster node needs to find out the hostnames of all the other nodes somehow. You *could* designate one node to be the one that every other node asks, and put that node's hostname in the config file, but that wouldn't be very decentralized. Instead, we include *every* node in the list of nodes-to-ask.

If you're curious about the RethinkDB config file, there's a [RethinkDB documentation page](#) about it. The [explanations of the RethinkDB command-line options](#) are another useful reference.

See the [RethinkDB documentation on securing your cluster](#).

Install Python 3.4+

If you don't already have it, then you should [install Python 3.4+](#).

If you're testing or developing BigchainDB on a stand-alone node, then you should probably create a Python 3.4+ virtual environment and activate it (e.g. using `virtualenv` or `conda`). Later we will install several Python packages and you probably only want those installed in the virtual environment.

Install BigchainDB Server

BigchainDB Server has some OS-level dependencies that must be installed.

On Ubuntu 14.04, we found that the following was enough:

```
sudo apt-get update
sudo apt-get install g++ python3-dev
```

On Fedora 23, we found that the following was enough (tested in February 2015):

```
sudo dnf update
sudo dnf install gcc-c++ redhat-rpm-config python3-devel
```

(If you're using a version of Fedora before version 22, you may have to use `yum` instead of `dnf`.)

With OS-level dependencies installed, you can install BigchainDB Server with `pip` or from source.

How to Install BigchainDB with pip

BigchainDB (i.e. both the Server and the officially-supported drivers) is distributed as a Python package on PyPI so you can install it using `pip`. First, make sure you have an up-to-date Python 3.4+ version of `pip` installed:

```
pip -V
```

If it says that `pip` isn't installed, or it says `pip` is associated with a Python version less than 3.4, then you must install a `pip` version associated with Python 3.4+. In the following instructions, we call it `pip3` but you may be able to use `pip` if that refers to the same thing. See [the pip installation instructions](#).

On Ubuntu 14.04, we found that this works:

```
sudo apt-get install python3-pip
```

That should install a Python 3 version of `pip` named `pip3`. If that didn't work, then another way to get `pip3` is to do `sudo apt-get install python3-setuptools` followed by `sudo easy_install3 pip`.

You can upgrade `pip` (`pip3`) and `setuptools` to the latest versions using:


```
pip3 install --upgrade pip setuptools
pip3 -V
```

Now you can install BigchainDB Server (and officially-supported BigchainDB drivers) using:

```
pip3 install bigchaindb
```

(If you're not in a virtualenv and you want to install bigchaindb system-wide, then put `sudo` in front.)

Note: You can use `pip3` to upgrade the `bigchaindb` package to the latest version using `pip3 install --upgrade bigchaindb`.

How to Install BigchainDB from Source

If you want to install BitchainDB from source because you want to use the very latest bleeding-edge code, clone the public repository:

```
git clone git@github.com:bigchaindb/bigchaindb.git
python setup.py install
```

Configure BigchainDB Server

Start by creating a default BigchainDB config file:

```
bigchaindb -y configure
```

(There's documentation for the `bigchaindb` command is in the section on the BigchainDB Command Line Interface (CLI).)

Edit the created config file:

- Open `$HOME/.bigchaindb` (the created config file) in your text editor.
- Change `"server": {"bind": "localhost:9984", ... }` to `"server": {"bind": "0.0.0.0:9984", ... }`. This makes it so traffic can come from any IP address to port 9984 (the HTTP Client-Server API port).
- Change `"api_endpoint": "http://localhost:9984/api/v1"` to `"api_endpoint": "http://your_api_hostname:9984/api/v1"`
- Change `"keyring": []` to `"keyring": ["public_key_of_other_node_A", "public_key_of_other_node_B", "..."]` i.e. a list of the public keys of all the other nodes in the federation. The keyring should *not* include your node's public key.

For more information about the BigchainDB config file, see [Configuring a BigchainDB Node](#).

Run RethinkDB Server

Start RethinkDB using:

```
rethinkdb --config-file path/to/instance1.conf
```

except replace the path with the actual path to `instance1.conf`.

Note: It's possible to [make RethinkDB start at system startup](#).

You can verify that RethinkDB is running by opening the RethinkDB web interface in your web browser. It should be at `http://rethinkdb-hostname:8080/`. If you're running RethinkDB on localhost, that would be `http://localhost:8080/`.

Run BigchainDB Server

After all node operators have started RethinkDB, but before they start BigchainDB, one designated node operator must configure the RethinkDB database by running the following commands:

```
bigchaindb init
bigchaindb set-shards numshards
bigchaindb set-replicas numreplicas
```

where:

- `bigchaindb init` creates the database within RethinkDB, the tables, the indexes, and the genesis block.
- `numshards` should be set to the number of nodes in the initial cluster.
- `numreplicas` should be set to the database replication factor decided by the federation. It must be 3 or more for [RethinkDB failover](#) to work.

Once the RethinkDB database is configured, every node operator can start BigchainDB using:

```
bigchaindb start
```

1.5 Develop & Test BigchainDB Server

1.5.1 Set Up & Run a Dev/Test Node

This page explains how to set up a minimal local BigchainDB node for development and testing purposes.

The BigchainDB core dev team develops BigchainDB on recent Ubuntu and Fedora distributions, so we recommend you use one of those. BigchainDB Server doesn't work on Windows and Mac OS X (unless you use a VM or containers).

First, read through the BigchainDB [CONTRIBUTING.md](#) file. It outlines the steps to setup a machine for developing and testing BigchainDB.

Next, create a default BigchainDB config file (in `$HOME/.bigchaindb`):

```
bigchaindb -y configure
```

Note: The BigchainDB CLI and the BigchainDB Configuration Settings are documented elsewhere. (Click the links.)

Start RethinkDB using:

```
rethinkdb
```

You can verify that RethinkDB is running by opening the RethinkDB web interface in your web browser. It should be at `http://localhost:8080/`.

To run BigchainDB Server, do:

```
bigchaindb start
```

You can run all the unit tests to test your installation.

The BigchainDB [CONTRIBUTING.md](#) file has more details about how to contribute.

1.5.2 Running Unit Tests

Once you’ve installed BigchainDB Server, you may want to run all the unit tests. This section explains how.

First of all, if you installed BigchainDB Server using `pip` (i.e. by getting the package from PyPI), then you didn’t install the tests. **Before you can run all the unit tests, you must install BigchainDB from source.**

To run all the unit tests, first make sure you have RethinkDB running:

```
$ rethinkdb
```

then in another terminal, do:

```
$ python setup.py test
```

(Aside: How does the above command work? The documentation for [pytest-runner](#) explains. We use `pytest` to write all unit tests.)

Using docker-compose to Run the Tests

You can also use `docker-compose` to run the unit tests.

Start RethinkDB in the background:

```
$ docker-compose up -d rdb
```

then run the unit tests using:

```
$ docker-compose run --rm bdb py.test -v
```

1.6 Settings & CLI

1.6.1 BigchainDB Configuration Settings

Note: At the time of writing, BigchainDB Server code and BigchainDB Python driver code are mixed together, so the following settings are the settings used by BigchainDB Server and also by clients written using the Python driver code. Soon, the code will be separated into server, driver and shared modules, so that BigchainDB Server and BigchainDB clients will have different configuration settings.

The value of each configuration setting is determined according to the following rules:

- If it’s set by an environment variable, then use that value
- Otherwise, if it’s set in a local config file, then use that value
- Otherwise, use the default value

For convenience, here’s a list of all the relevant environment variables (documented below):

```
BIGCHAINDB_KEYPAIR_PUBLIC    BIGCHAINDB_KEYPAIR_PRIVATE    BIGCHAINDB_KEYRING
BIGCHAINDB_DATABASE_HOST    BIGCHAINDB_DATABASE_PORT    BIGCHAINDB_DATABASE_NAME
BIGCHAINDB_SERVER_BIND    BIGCHAINDB_SERVER_WORKERS    BIGCHAINDB_SERVER_THREADS
BIGCHAINDB_API_ENDPOINT    BIGCHAINDB_CONSENSUS_PLUGIN    BIGCHAINDB_STATSD_HOST
BIGCHAINDB_STATSD_PORT    BIGCHAINDB_STATSD_RATE    BIGCHAINDB_CONFIG_PATH
```

The local config file is `$HOME/.bigchaindb` by default (a file which might not even exist), but you can tell BigchainDB to use a different file by using the `-c` command-line option, e.g. `bigchaindb -c path/to/config_file.json start` or using the `BIGCHAINDB_CONFIG_PATH` environment variable, e.g.

BIGHAINDB_CONFIG_PATH=.my_bigchaindb_config bigchaindb start. Note that the `-c` command line option will always take precedence if both the BIGHAINDB_CONFIG_PATH and the `-c` command line option are used.

You can read the current default values in the file `bigchaindb/__init__.py`. (The link is to the latest version.)

Running `bigchaindb -y configure` will generate a local config file in `$HOME/.bigchaindb` with all the default values, with two exceptions: It will generate a valid private/public keypair, rather than using the default keypair (None and None).

keypair.public & keypair.private

The cryptographic keypair used by the node. The public key is how the node identifies itself to the world. The private key is used to generate cryptographic signatures. Anyone with the public key can verify that the signature was generated by whoever had the corresponding private key.

Example using environment variables

```
export BIGHAINDB_KEYPAIR_PUBLIC=8wHUvvraRo5yEoJAt66UTZaFq9YZ9tFFwcauKPDtjkGw
export BIGHAINDB_KEYPAIR_PRIVATE=5C5Cknco7YxBRP9AgB1cbUVTL4FAcooxErLygw1DeG2D
```

Example config file snippet

```
"keypair": {
  "public": "8wHUvvraRo5yEoJAt66UTZaFq9YZ9tFFwcauKPDtjkGw",
  "private": "5C5Cknco7YxBRP9AgB1cbUVTL4FAcooxErLygw1DeG2D"
}
```

Internally (i.e. in the Python code), both keys have a default value of None, but that's not a valid key. Therefore you can't rely on the defaults for the keypair. If you want to run BigchainDB, you must provide a valid keypair, either in the environment variables or in the local config file. You can generate a local config file with a valid keypair (and default everything else) using `bigchaindb -y configure`.

keyring

A list of the public keys of all the nodes in the cluster, excluding the public key of this node.

Example using an environment variable

```
export BIGHAINDB_KEYRING=BnCsre9MPBeQK8QZBFznU2dJJ2GwtvnSMdemCmod2XPB:4cYQHoQrvPiut3Sjs8fVR1BMZZpJjMTC4bsMTt9V71aQ
```

Note how the keys in the list are separated by colons.

Example config file snippet

```
"keyring": ["BnCsre9MPBeQK8QZBFznU2dJJ2GwtvnSMdemCmod2XPB",
            "4cYQHoQrvPiut3Sjs8fVR1BMZZpJjMTC4bsMTt9V71aQ"]
```

Default value (from a config file)

```
"keyring": []
```

database.host, database.port & database.name

The RethinkDB database hostname, port and name.

Example using environment variables

```
export BIGCHAINDB_DATABASE_HOST=localhost
export BIGCHAINDB_DATABASE_PORT=28015
export BIGCHAINDB_DATABASE_NAME=bigchain
```

Example config file snippet

```
"database": {
  "host": "localhost",
  "port": 28015,
  "name": "bigchain"
}
```

Default values (a snippet from `bigchaindb/__init__.py`)

```
'database': {
  'host': os.environ.get('BIGCHAINDB_DATABASE_HOST', 'localhost'),
  'port': 28015,
  'name': 'bigchain',
}
```

server.bind, server.workers & server.threads

These settings are for the [Gunicorn HTTP server](#), which is used to serve the HTTP client-server API.

`server.bind` is where to bind the Gunicorn HTTP server socket. It's a string. It can be any valid value for [Gunicorn's bind setting](#). If you want to allow IPv4 connections from anyone, on port 9984, use `'0.0.0.0:9984'`. In a production setting, we recommend you use Gunicorn behind a reverse proxy server. If Gunicorn and the reverse proxy are running on the same machine, then use `'localhost:PORT'` where `PORT` is *not* 9984 (because the reverse proxy needs to listen on port 9984). Maybe use `PORT=9983` in that case because we know 9983 isn't used. If Gunicorn and the reverse proxy are running on different machines, then use `'A.B.C.D:9984'` where `A.B.C.D` is the IP address of the reverse proxy. There's [more information about deploying behind a reverse proxy in the Gunicorn documentation](#). (They call it a proxy.)

`server.workers` is [the number of worker processes](#) for handling requests. If `None` (the default), the value will be `(cpu_count * 2 + 1)`. `server.threads` is [the number of threads-per-worker](#) for handling requests. If `None` (the default), the value will be `(cpu_count * 2 + 1)`. The HTTP server will be able to handle `server.workers * server.threads` requests simultaneously.

Example using environment variables

```
export BIGCHAINDB_SERVER_BIND=0.0.0.0:9984
export BIGCHAINDB_SERVER_WORKERS=5
export BIGCHAINDB_SERVER_THREADS=5
```

Example config file snippet

```
"server": {
  "bind": "0.0.0.0:9984",
  "workers": 5,
  "threads": 5
}
```

Default values (from a config file)

```
"server": {
  "bind": "localhost:9984",
  "workers": null,
  "threads": null
}
```

api_endpoint

api_endpoint is the URL where a BigchainDB client can get access to the HTTP client-server API.

Example using an environment variable

```
export BIGCHAINDB_API_ENDPOINT="http://localhost:9984/api/v1"
```

Example config file snippet

```
"api_endpoint": "http://webserver.blocks587.net:9984/api/v1"
```

Default value (from a config file)

```
"api_endpoint": "http://localhost:9984/api/v1"
```

consensus_plugin

The consensus plugin to use.

Example using an environment variable

```
export BIGCHAINDB_CONSENSUS_PLUGIN=default
```

Example config file snippet: the default

```
"consensus_plugin": "default"
```

statsd.host, statsd.port & statsd.rate

These settings are used to configure where, and how often, [StatsD](#) should send data for cluster monitoring purposes. statsd.host is the hostname of the monitoring server, where StatsD should send its data. stats.port is the port. statsd.rate is the fraction of transaction operations that should be sampled. It's a float between 0.0 and 1.0.

Example using environment variables

```
export BIGCHAINDB_STATSD_HOST="http://monitor.monitors-r-us.io"
export BIGCHAINDB_STATSD_PORT=8125
export BIGCHAINDB_STATSD_RATE=0.01
```

Example config file snippet: the default

```
"statsd": {"host": "localhost", "port": 8125, "rate": 0.01}
```

1.6.2 BigchainDB Command Line Interface (CLI)

Note: At the time of writing, BigchainDB Server and our BigchainDB client are combined, so the BigchainDB CLI includes some server-specific commands and some client-specific commands (e.g. `bigchaindb load`). Soon, BigchainDB Server will be separate from all BigchainDB clients, and they'll all have different CLIs.

The command-line command to interact with BigchainDB is `bigchaindb`.

bigchaindb -help

Show help for the `bigchaindb` command. `bigchaindb -h` does the same thing.

bigchaindb -version

Show the version number. `bigchaindb -v` does the same thing.

bigchaindb configure

Generate a local config file (which can be used to set some or all BigchainDB node configuration settings). It will auto-generate a public-private keypair and then ask you for the values of other configuration settings. If you press Enter for a value, it will use the default value.

If you use the `-c` command-line option, it will generate the file at the specified path:

```
bigchaindb -c path/to/new_config.json configure
```

If you don't use the `-c` command-line option, the file will be written to `$HOME/.bigchaindb` (the default location where BigchainDB looks for a config file, if one isn't specified).

If you use the `-y` command-line option, then there won't be any interactive prompts: it will just generate a keypair and use the default values for all the other configuration settings.

```
bigchaindb -y configure
```

bigchaindb show-config

Show the values of the BigchainDB node configuration settings.

bigchaindb export-my-pubkey

Write the node's public key (i.e. one of its configuration values) to standard output (stdout).

bigchaindb init

Create a RethinkDB database, all RethinkDB database tables, various RethinkDB database indexes, and the genesis block.

Note: The `bigchaindb start` command (see below) always starts by trying a `bigchaindb init` first. If it sees that the RethinkDB database already exists, then it doesn't re-initialize the database. One doesn't have to do `bigchaindb init` before `bigchaindb start`. `bigchaindb init` is useful if you only want to initialize (but not start).

bigchaindb drop

Drop (erase) the RethinkDB database. You will be prompted to make sure. If you want to force-drop the database (i.e. skipping the yes/no prompt), then use `bigchaindb -y drop`

bigchaindb start

Start BigchainDB. It always begins by trying a `bigchaindb init` first. See the note in the documentation for `bigchaindb init`. You can also use the `--experimental-start-rethinkdb` command line option to automatically start rethinkdb with bigchaindb if rethinkdb is not already running, e.g. `bigchaindb --experimental-start-rethinkdb start`. Note that this will also shutdown rethinkdb when the bigchaindb process stops.

bigchaindb load

Write transactions to the backlog (for benchmarking tests). You can learn more about it using:

```
$ bigchaindb load -h
```

bigchaindb set-shards

Set the number of shards in the underlying datastore. For example, the following command will set the number of shards to four:

```
$ bigchaindb set-shards 4
```

bigchaindb set-replicas

Set the number of replicas (of each shard) in the underlying datastore. For example, the following command will set the number of replicas to three (i.e. it will set the replication factor to three):

```
$ bigchaindb set-replicas 3
```

1.7 Drivers & Clients

1.7.1 The Python Server API by Example

Currently, the HTTP Client-Server API is very rudimentary, so you may want to use the Python Server API to develop prototype clients and applications, for now. Keep in mind that in the future, clients will only be able to use the HTTP Client-Server API (and possibly other Client-Server APIs) to communicate with BigchainDB nodes.

This section has examples of using the Python Server API to interact *directly* with a BigchainDB node running BigchainDB Server. That is, in these examples, the Python code and BigchainDB Server run on the same machine.

One can also interact with a BigchainDB node via other APIs, including the HTTP Client-Server API.

Getting Started

First, make sure you have RethinkDB and BigchainDB *installed and running*, i.e. you installed them and you ran:

```
$ rethinkdb
$ bigchaindb configure
$ bigchaindb start
```

Don't shut them down! In a new terminal, open a Python shell:

```
$ python
```

Now we can import the `Bigchain` class and create an instance:

```
from bigchaindb import Bigchain
b = Bigchain()
```


This instantiates an object `b` of class `Bigchain`. When instantiating a `Bigchain` object without arguments (as above), it reads the configurations stored in `$HOME/.bigchaindb`.

In a federation of BigchainDB nodes, each node has its own `Bigchain` instance.

The `Bigchain` class is the main API for all BigchainDB interactions, right now. It does things that BigchainDB nodes do, but it also does things that BigchainDB clients do. In the future, it will be refactored into different parts. The `Bigchain` class is documented elsewhere ([link](#)).

Create a Digital Asset

At a high level, a “digital asset” is something which can be represented digitally and can be assigned to a user. In BigchainDB, users are identified by their public key, and the data payload in a digital asset is represented using a generic [Python dict](#).

In BigchainDB, only the federation nodes are allowed to create digital assets, by doing a special kind of transaction: a `CREATE` transaction.

```
from bigchaindb import crypto

# Create a test user
testuser1_priv, testuser1_pub = crypto.generate_key_pair()

# Define a digital asset data payload
digital_asset_payload = {'msg': 'Hello BigchainDB!'}

# A create transaction uses the operation `CREATE` and has no inputs
tx = b.create_transaction(b.me, testuser1_pub, None, 'CREATE', payload=digital_asset_payload)

# All transactions need to be signed by the user creating the transaction
tx_signed = b.sign_transaction(tx, b.me_private)

# Write the transaction to the bigchain.
# The transaction will be stored in a backlog where it will be validated,
# included in a block, and written to the bigchain
b.write_transaction(tx_signed)
```

Read the Creation Transaction from the DB

After a couple of seconds, we can check if the transactions was included in the bigchain:

```
# Retrieve a transaction from the bigchain
tx_retrieved = b.get_transaction(tx_signed['id'])
tx_retrieved
```

```
{
  "id": "933cd83a419d2735822a2154c84176a2f419cbd449a74b94e592ab807af23861",
  "transaction": {
    "conditions": [
      {
        "cid": 0,
        "condition": {
          "details": {
            "bitmask": 32,
            "public_key": "BwuhqQX8FPsmqYiRV2CSZYWwsSWgSSQqFHjqxKEuqkPs",
            "signature": None,
            "type": "fulfillment",
```

```
        "type_id":4
      },
      "uri":"cc:4:20:oqXTWvR3afHHX8OaOO84kZxS6nH4GEBXD4Vw8Mc5iBo:96"
    },
    "owners_after":[
      "BwuhqQX8FPsmqYiRV2CSZYWWsSWgSSQQFHjqxKEuqkPs"
    ]
  },
  "data":{
    "hash":"872fa6e6f46246cd44afdb2ee9cfae0e72885fb0910e2bcf9a5a2a4eadb417b8",
    "payload":{
      "msg":"Hello BigchainDB!"
    }
  },
  "fulfillments":[
    {
      "owners_before":[
        "3LQ5dTiddXymDhNzETB1rEkp4mA7fEV1Qeiu5ghHiJm9"
      ],
      "fid":0,
      "fulfillment":"cf:4:Iq-BcczwraM2UpF-TDPdwK8fQ6IXkD_6uJaxBZd984yxCGX7Csx-S2FBVe8LVyW2",
      "input":None
    }
  ],
  "operation":"CREATE",
  "timestamp":"1460981667.449279"
},
"version":1
}
```

The new owner of the digital asset is now BwuhqQX8FPsmqYiRV2CSZYWWsSWgSSQQFHjqxKEuqkPs, which is the public key of testuser1.

Note that the current owner (with public key 3LQ5dTiddXymDhNzETB1rEkp4mA7fEV1Qeiu5ghHiJm9) is the federation node which created the asset and assigned it to testuser1.

Transfer the Digital Asset

Now that testuser1 has a digital asset assigned to him, he can transfer it to another user. Transfer transactions require an input. The input will be the transaction id of a digital asset that was assigned to testuser1, which in our case is 933cd83a419d2735822a2154c84176a2f419cbd449a74b94e592ab807af23861.

BigchainDB makes use of the crypto-conditions library to both cryptographically lock and unlock transactions. The locking script is referred to as a condition and a corresponding fulfillment unlocks the condition of the input_tx.

Since a transaction can have multiple outputs with each its own (crypto)condition, each transaction input should also refer to the condition index cid.

```
# Create a second testuser
testuser2_priv, testuser2_pub = crypto.generate_key_pair()

# Retrieve the transaction with condition id
tx_retrieved_id = b.get_owned_ids(testuser1_pub).pop()
tx_retrieved_id
```

```
{
  "cid":0,
  "txid":"933cd83a419d2735822a2154c84176a2f419cbd449a74b94e592ab807af23861"
}
```

```
# Create a transfer transaction
tx_transfer = b.create_transaction(testuser1_pub, testuser2_pub, tx_retrieved_id, 'TRANSFER')

# Sign the transaction
tx_transfer_signed = b.sign_transaction(tx_transfer, testuser1_priv)

# Write the transaction
b.write_transaction(tx_transfer_signed)

# Check if the transaction is already in the bigchain
tx_transfer_retrieved = b.get_transaction(tx_transfer_signed['id'])
tx_transfer_retrieved
```

```
{
  "id":"aall365317cb89bfdae2375bae76d6b8232008f8672507080e3766ca06976dcd",
  "transaction":{
    "conditions":[
      {
        "cid":0,
        "condition":{
          "details":{
            "bitmask":32,
            "public_key":"qv8DvdNG5nZHWCP5aPSggqxAvAPJpQj19abRvFCntor",
            "signature":None,
            "type":"fulfillment",
            "type_id":4
          },
          "uri":"cc:4:20:DIfyalZvV_9uko001mxmK3nxsfAWSKYFF33XDYkby4E:96"
        },
        "owners_after":[
          "qv8DvdNG5nZHWCP5aPSggqxAvAPJpQj19abRvFCntor"
        ]
      }
    ],
    "data":None,
    "fulfillments":[
      {
        "owners_before":[
          "BwuhqQX8FPsmqYiRV2CSZYWWsSWgSSQQFHjqxKEuqkPs"
        ],
        "fid":0,
        "fulfillment":"cf:4:oqXTWvR3afHHX8Oa0084kZxS6nH4GEBXD4Vw8Mc5iBqzkVR6cFJhRvMGKa-Lc81s",
        "input":{
          "cid":0,
          "txid":"933cd83a419d2735822a2154c84176a2f419cbd449a74b94e592ab807af23861"
        }
      }
    ],
    "operation":"TRANSFER",
    "timestamp":"1460981677.472037"
  },
  "version":1
}
```

Double Spends

BigchainDB makes sure that a user can't transfer the same digital asset two or more times (i.e. it prevents double spends).

If we try to create another transaction with the same input as before, the transaction will be marked invalid and the validation will throw a double spend exception:

```
# Create another transfer transaction with the same input
tx_transfer2 = b.create_transaction(testuser1_pub, testuser2_pub, tx_retrieved_id, 'TRANSFER')

# Sign the transaction
tx_transfer_signed2 = b.sign_transaction(tx_transfer2, testuser1_priv)

# Check if the transaction is valid
b.validate_transaction(tx_transfer_signed2)
```

```
DoubleSpend: input `{'cid': 0, 'txid': '933cd83a419d2735822a2154c84176a2f419cbd449a74b94e592ab807af23...
```

Multiple Owners

To create a new digital asset with *multiple* owners, one can simply provide a list of owners_after:

```
# Create a new asset and assign it to multiple owners
tx_multisig = b.create_transaction(b.me, [testuser1_pub, testuser2_pub], None, 'CREATE')

# Have the federation node sign the transaction
tx_multisig_signed = b.sign_transaction(tx_multisig, b.me_private)

# Write the transaction
b.write_transaction(tx_multisig_signed)

# Check if the transaction is already in the bigchain
tx_multisig_retrieved = b.get_transaction(tx_multisig_signed['id'])
tx_multisig_retrieved
```

```
{
  "id": "a9a6e5c74ea02b8885c83125f1b74a2ba8ca42236ec5e1c358aa1053ec721ccb",
  "transaction": {
    "conditions": [
      {
        "cid": 0,
        "condition": {
          "details": {
            "bitmask": 41,
            "subfulfillments": [
              {
                "bitmask": 32,
                "public_key": "BwuhqQX8FPsmqYiRV2CSZYWWsSWgSSQQFHjqxKEuqkPs",
                "signature": None,
                "type": "fulfillment",
                "type_id": 4,
                "weight": 1
              },
              {
                "bitmask": 32,
                "public_key": "qv8DvdNG5nZHWCP5aPSqgqxAvAPJpQj19abRvFCntor",

```

```

        "signature":None,
        "type":"fulfillment",
        "type_id":4,
        "weight":1
    },
    ],
    "threshold":2,
    "type":"fulfillment",
    "type_id":2
},
"uri":"cc:2:29:DpflJzUSlnTUBx81D8QUolOA-M9nQnrGwvWSk7f3REc:206"
},
"owners_after":[
    "BwuhqQX8FPsmqYiRV2CSZYWWsSWgSSQQFHjqxKEuqkPs",
    "qv8DvdNG5nZHWCP5aPSqgqxAvaPJpQj19abRvFCntor"
]
}
],
"data":None,
"fulfillments":[
    {
        "owners_before":[
            "3LQ5dTiddXymDhNzETB1rEkp4mA7fEVlQeiu5ghHiJm9"
        ],
        "fid":0,
        "fulfillment":"cf:4:Iq-BcczwraM2UpF-TDPdwK8fQ6IXkD_6uJaxBZd984z5qdHRz9Jag68dkOyZS5_Y",
        "input":None
    }
],
"operation":"CREATE",
"timestamp":"1460981687.501433"
},
"version":1
}

```

The asset can be transferred as soon as each of the `owners_after` signs the transaction.

To do so, simply provide a list of all private keys to the signing routine:

```

# Create a third testuser
testuser3_priv, testuser3_pub = crypto.generate_key_pair()

# Retrieve the multisig transaction
tx_multisig_retrieved_id = b.get_owned_ids(testuser2_pub).pop()

# Transfer the asset from the 2 owners to the third testuser
tx_multisig_transfer = b.create_transaction([testuser1_pub, testuser2_pub], testuser3_pub, tx_multisig_retrieved_id)

# Sign with both private keys
tx_multisig_transfer_signed = b.sign_transaction(tx_multisig_transfer, [testuser1_priv, testuser2_priv])

# Write the transaction
b.write_transaction(tx_multisig_transfer_signed)

# Check if the transaction is already in the bigchain
tx_multisig_transfer_retrieved = b.get_transaction(tx_multisig_transfer_signed['id'])
tx_multisig_transfer_retrieved

```

```
{
  "id": "e689e23f774e7c562eeb310c7c712b34fb6210bea5deb9175e48b68810029150",
  "transaction": {
    "conditions": [
      {
        "cid": 0,
        "condition": {
          "details": {
            "bitmask": 32,
            "public_key": "8YN9fALMj9CkeCcmTiM2kxwurpkMzHg9RkwSLJKMasvG",
            "signature": None,
            "type": "fulfillment",
            "type_id": 4
          },
          "uri": "cc:4:20:cAq6JQJXtwlxURqrksiyqLThB9zh08ZxSPLTDSaReYE:96"
        },
        "owners_after": [
          "8YN9fALMj9CkeCcmTiM2kxwurpkMzHg9RkwSLJKMasvG"
        ]
      },
      {
        "data": None,
        "fulfillments": [
          {
            "owners_before": [
              "BwuhqQX8FPsmqYiRV2CSZYWWsSWgSSQQFHjqxKEuqkPs",
              "qv8DvdNG5nZHWCP5aPSggqxAvaPJpQj19abRvFCntor"
            ],
            "fid": 0,
            "fulfillment": "cf:4:0qXTWvR3afHHX8OaOO84kZxS6nH4GEBXD4Vw8Mc5iBrcuiGDNVgpH9SwiuNeYZ-n",
            "input": {
              "cid": 0,
              "txid": "aa11365317cb89bfdae2375bae76d6b8232008f8672507080e3766ca06976dcd"
            }
          }
        ],
        "operation": "TRANSFER",
        "timestamp": "1460981697.526878"
      },
      {
        "version": 1
      }
    ]
  }
}
```

Multiple Inputs and Outputs

With BigchainDB it is possible to send multiple assets to someone in a single transfer.

The transaction will create a fulfillment - condition pair for each input, which can be referred to by `fid` and `cid` respectively.

```
# Create some assets for bulk transfer
for i in range(3):
    tx_mimo_asset = b.create_transaction(b.me, testuser1_pub, None, 'CREATE')
    tx_mimo_asset_signed = b.sign_transaction(tx_mimo_asset, b.me_private)
    b.write_transaction(tx_mimo_asset_signed)

# Wait until they appear on the bigchain and get the inputs
owned_mimo_inputs = b.get_owned_ids(testuser1_pub)
```

```

# Check the number of assets
print(len(owned_mimo_inputs))

# Create a signed TRANSFER transaction with all the assets
tx_mimo = b.create_transaction(testuser1_pub, testuser2_pub, owned_mimo_inputs, 'TRANSFER')
tx_mimo_signed = b.sign_transaction(tx_mimo, testuser1_priv)

# Write the transaction
b.write_transaction(tx_mimo_signed)

# Check if the transaction is already in the bigchain
tx_mimo_retrieved = b.get_transaction(tx_mimo_signed['id'])
tx_mimo_retrieved

```

```

{
  "id": "8b63689691a3c2e8faba89c6efe3caa0661f862c14d88d1e63ebd65d49484de2",
  "transaction": {
    "conditions": [
      {
        "cid": 0,
        "condition": {
          "details": {
            "bitmask": 32,
            "public_key": "qv8DvdNG5nZHWCP5aPSqgqxAvaPJpQj19abRvFCntor",
            "signature": None,
            "type": "fulfillment",
            "type_id": 4
          },
          "uri": "cc:4:20:2AXg2JJ7mQ8o2Q9-hafP-XmFh3YR7I2_Sz55AubfxIc:96"
        },
        "owners_after": [
          "qv8DvdNG5nZHWCP5aPSqgqxAvaPJpQj19abRvFCntor"
        ]
      },
      {
        "cid": 1,
        "condition": {
          "details": {
            "bitmask": 32,
            "public_key": "qv8DvdNG5nZHWCP5aPSqgqxAvaPJpQj19abRvFCntor",
            "signature": None,
            "type": "fulfillment",
            "type_id": 4
          },
          "uri": "cc:4:20:2AXg2JJ7mQ8o2Q9-hafP-XmFh3YR7I2_Sz55AubfxIc:96"
        },
        "owners_after": [
          "qv8DvdNG5nZHWCP5aPSqgqxAvaPJpQj19abRvFCntor"
        ]
      },
      {
        "cid": 2,
        "condition": {
          "details": {
            "bitmask": 32,
            "public_key": "qv8DvdNG5nZHWCP5aPSqgqxAvaPJpQj19abRvFCntor",
            "signature": None,
            "type": "fulfillment",

```

```
        "type_id":4
      },
      "uri":"cc:4:20:2AXg2JJ7mQ8o2Q9-hafP-XmFh3YR7I2_Sz55AubfxIc:96"
    },
    "owners_after":[
      "qv8DvdNG5nZHWCP5aPSggqxAvaPJpQj19abRvFCntor"
    ]
  }
],
"data":None,
"fulfillments":[
  {
    "owners_before":[
      "BwuhqQX8FPsmqYiRV2CSZYWWsSWgSSQQFHjqxKEuqkPs"
    ],
    "fid":0,
    "fulfillment":"cf:4:sTzo4fvm8U8XrlXcgcGkNZgkfS9QHg2grgrJiX-c0LT_a83V0wbNRVbmb0eOy6tL",
    "input":{
      "cid":0,
      "txid":"9a99f3c82aea23fb344acb1505926365e2c6b722761c4be6ab8916702c94c024"
    }
  },
  {
    "owners_before":[
      "BwuhqQX8FPsmqYiRV2CSZYWWsSWgSSQQFHjqxKEuqkPs"
    ],
    "fid":1,
    "fulfillment":"cf:4:sTzo4fvm8U8XrlXcgcGkNZgkfS9QHg2grgrJiX-c0LSJe3B_yjgXd1JHPBJhAdyw",
    "input":{
      "cid":0,
      "txid":"783014b92f35da0c2526e1db6f81452c61853d29eda50d057fd043d507d03ef9"
    }
  },
  {
    "owners_before":[
      "BwuhqQX8FPsmqYiRV2CSZYWWsSWgSSQQFHjqxKEuqkPs"
    ],
    "fid":2,
    "fulfillment":"cf:4:sTzo4fvm8U8XrlXcgcGkNZgkfS9QHg2grgrJiX-c0LReUQd-vDMseuVi03qY5Fxe",
    "input":{
      "cid":0,
      "txid":"9ab6151334b06f3f3aab282597ee8a7c12b9d7a0c43f356713f7ef9663375f50"
    }
  }
],
"operation":"TRANSFER",
"timestamp":"1461049149.568927"
},
"version":1
}
```


Crypto-Conditions (Advanced)

Introduction

Crypto-conditions provide a mechanism to describe a signed message such that multiple actors in a distributed system can all verify the same signed message and agree on whether it matches the description.

This provides a useful primitive for event-based systems that are distributed on the Internet since we can describe events in a standard deterministic manner (represented by signed messages) and therefore define generic authenticated event handlers.

Crypto-conditions are part of the Interledger protocol and the full specification can be found [here](#).

Implementations of the crypto-conditions are available in [Python](#) and [JavaScript](#).

Threshold Conditions

Threshold conditions introduce multi-signatures, m-of-n signatures or even more complex binary Merkle trees to BigchainDB.

Setting up a generic threshold condition is a bit more elaborate than regular transaction signing but allow for flexible signing between multiple parties or groups.

The basic workflow for creating a more complex cryptocondition is the following:

1. Create a transaction template that include the public key of all (nested) parties as `owners_after`
2. Set up the threshold condition using the [cryptocondition library](#)
3. Update the condition and hash in the transaction template

We'll illustrate this by a threshold condition where 2 out of 3 `owners_after` need to sign the transaction:

```
import copy

import cryptoconditions as cc
from bigchaindb import util, crypto

# Create some new testusers
thresholduser1_priv, thresholduser1_pub = crypto.generate_key_pair()
thresholduser2_priv, thresholduser2_pub = crypto.generate_key_pair()
thresholduser3_priv, thresholduser3_pub = crypto.generate_key_pair()

# Retrieve the last transaction of testuser2
tx_retrieved_id = b.get_owned_ids(testuser2_pub).pop()

# Create a base template for a 1-input/2-output transaction
threshold_tx = b.create_transaction(testuser2_pub, [thresholduser1_pub, thresholduser2_pub, thresholduser3_pub])

# Create a Threshold Cryptocondition
threshold_condition = cc.ThresholdSha256Fulfillment(threshold=2)
threshold_condition.add_subfulfillment(cc.Ed25519Fulfillment(public_key=thresholduser1_pub))
threshold_condition.add_subfulfillment(cc.Ed25519Fulfillment(public_key=thresholduser2_pub))
threshold_condition.add_subfulfillment(cc.Ed25519Fulfillment(public_key=thresholduser3_pub))

# Update the condition in the newly created transaction
threshold_tx['transaction']['conditions'][0]['condition'] = {
    'details': threshold_condition.to_dict(),
    'uri': threshold_condition.condition.serialize_uri()
}
```

```
}

# Conditions have been updated, so the transaction hash (ID) needs updating
threshold_tx['id'] = util.get_hash_data(threshold_tx)

# Sign the transaction
threshold_tx_signed = b.sign_transaction(threshold_tx, testuser2_priv)

# Write the transaction
b.write_transaction(threshold_tx_signed)

# Check if the transaction is already in the bigchain
tx_threshold_retrieved = b.get_transaction(threshold_tx_signed['id'])
tx_threshold_retrieved
```

```
{
  "id": "0057d29ff735d91505decf5e7195ea8da675b01676165abf23ea774bbb469383",
  "transaction": {
    "conditions": [
      {
        "cid": 0,
        "condition": {
          "details": {
            "bitmask": 41,
            "subfulfillments": [
              {
                "bitmask": 32,
                "public_key": "8NaGq26YMcEvj8Sc5MnqspKzFTQdleZBAuuPDw4ERHpz",
                "signature": None,
                "type": "fulfillment",
                "type_id": 4,
                "weight": 1
              },
              {
                "bitmask": 32,
                "public_key": "ALE9Agojob28D1fHWCxFXJwpqrYPkcsUs26YksBVj27z",
                "signature": None,
                "type": "fulfillment",
                "type_id": 4,
                "weight": 1
              },
              {
                "bitmask": 32,
                "public_key": "Cx4jWSGci7fw6z5QyeApCijbwnMpyuhp4C1kzuFc3XrM",
                "signature": None,
                "type": "fulfillment",
                "type_id": 4,
                "weight": 1
              }
            ],
            "threshold": 2,
            "type": "fulfillment",
            "type_id": 2
          },
          "uri": "cc:2:29:FoElId4TE5TU2loonT7sayXhxcmaJVOCeIduh56Dxw:246"
        },
        "owners_after": [
          "8NaGq26YMcEvj8Sc5MnqspKzFTQdleZBAuuPDw4ERHpz",

```

```

        "ALE9Agojob28D1fHWCxFXJwpqrYPkcsUs26YksBVj27z",
        "Cx4jWSGci7fw6z5QyeApCijbwnMpyuhp4C1kzuFc3XrM"
    ]
    },
    ],
    "data":None,
    "fulfillments":[
        {
            "owners_before":[
                "qv8DvdNG5nZHWCP5aPSqgqxAvaPJpQj19abRvFCntor"
            ],
            "fid":0,
            "fulfillment":"cf:4:DIfyalZvV_9uko001mxmK3nxsfAWSKYF33XDYkbY4EbD7-_neXJJEE_tVTDcl_Eo",
            "input":{
                "cid":0,
                "txid":"aa11365317cb89bfdae2375bae76d6b8232008f8672507080e3766ca06976dcd"
            }
        }
    ],
    "operation":"TRANSFER",
    "timestamp":"1460981707.559401"
},
"version":1
}

```

The transaction can now be transferred by fulfilling the threshold condition.

The fulfillment involves:

1. Create a transaction template that include the public key of all (nested) parties as `owners_before`
2. Parsing the threshold condition into a fulfillment using the `cryptocondition` library
3. Signing all necessary subfulfillments and updating the fulfillment field in the transaction

```

# Create a new testuser to receive
thresholduser4_priv, thresholduser4_pub = crypto.generate_key_pair()

# Retrieve the last transaction of thresholduser1_pub
tx_retrieved_id = b.get_owned_ids(thresholduser1_pub).pop()

# Create a base template for a 2-input/1-output transaction
threshold_tx_transfer = b.create_transaction([thresholduser1_pub, thresholduser2_pub, thresholduser3

# Parse the threshold cryptocondition
threshold_fulfillment = cc.Fulfillment.from_dict(threshold_tx['transaction']['conditions'][0]['conditi

subfulfillment1 = threshold_fulfillment.get_subcondition_from_vk(thresholduser1_pub)[0]
subfulfillment2 = threshold_fulfillment.get_subcondition_from_vk(thresholduser2_pub)[0]
subfulfillment3 = threshold_fulfillment.get_subcondition_from_vk(thresholduser3_pub)[0]

# Get the fulfillment message to sign
threshold_tx_fulfillment_message = util.get_fulfillment_message(threshold_tx_transfer,
                                                                threshold_tx_transfer['transaction']
                                                                serialized=True)

# Clear the subconditions of the threshold fulfillment, they will be added again after signing
threshold_fulfillment.subconditions = []

```

```
# Sign and add the subconditions until threshold of 2 is reached
subfulfillment1.sign(threshold_tx_fulfillment_message, crypto.SigningKey(thresholduser1_priv))
threshold_fulfillment.add_subfulfillment(subfulfillment1)
subfulfillment2.sign(threshold_tx_fulfillment_message, crypto.SigningKey(thresholduser2_priv))
threshold_fulfillment.add_subfulfillment(subfulfillment2)

# Add remaining (unfulfilled) fulfillment as a condition
threshold_fulfillment.add_subcondition(subfulfillment3.condition)

# Update the fulfillment
threshold_tx_transfer['transaction']['fulfillments'][0]['fulfillment'] = threshold_fulfillment.serialize()

# Optional validation checks
assert threshold_fulfillment.validate(threshold_tx_fulfillment_message) == True
assert b.validate_fulfillments(threshold_tx_transfer) == True
assert b.validate_transaction(threshold_tx_transfer)

b.write_transaction(threshold_tx_transfer)
threshold_tx_transfer
```

```
{
  "id": "a45b2340c59df7422a5788b3c462dee708a18cdf09d1a10bd26be3f31af4b8d7",
  "transaction": {
    "conditions": [
      {
        "cid": 0,
        "condition": {
          "details": {
            "bitmask": 32,
            "public_key": "ED2pyPfsbNRTHkdMnaFkAwCSpZWRmbaM1h8fYzgRRMmc",
            "signature": None,
            "type": "fulfillment",
            "type_id": 4
          },
          "uri": "cc:4:20:xDz3NhRG-3eVzIB9sgnd99LKjOyDF-KlxWuf1TgNT0s:96"
        },
        "owners_after": [
          "ED2pyPfsbNRTHkdMnaFkAwCSpZWRmbaM1h8fYzgRRMmc"
        ]
      }
    ],
    "data": None,
    "fulfillments": [
      {
        "owners_before": [
          "8NaGq26YMcEvj8Sc5MnqspKzFTQdleZBAuuPDw4ERHpz",
          "ALE9Agojob28D1fHWCxFXJwpqrYPkcsUs26YksBVj27z",
          "Cx4jWSGci7fw6z5QyeApCijbwnMpyuhp4C1kzuFc3XrM"
        ],
        "fid": 0,
        "fulfillment": "cf:2:AQIBAwEBACcABAEgILGLuLLaNHo-KE59tkrpYmlVeucul6Eg9TcSuBqnMVwmAWAB",
        "input": {
          "cid": 0,
          "txid": "0057d29ff735d91505decf5e7195ea8da675b01676165abf23ea774bbb469383"
        }
      }
    ],
    "operation": "TRANSFER",
```

```

        "timestamp": "1460981717.579700"
    },
    "version": 1
}

```

Hash-locked Conditions

A hash-lock condition on an asset is like a password condition: anyone with the secret preimage (like a password) can fulfill the hash-lock condition and transfer the asset to themselves.

Under the hood, fulfilling a hash-lock condition amounts to finding a string (a “preimage”) which, when hashed, results in a given value. It’s easy to verify that a given preimage hashes to the given value, but it’s computationally difficult to *find* a string which hashes to the given value. The only practical way to get a valid preimage is to get it from the original creator (possibly via intermediaries).

One possible use case is to distribute preimages as “digital vouchers.” The first person to redeem a voucher will get the associated asset.

A federation node can create an asset with a hash-lock condition and no `owners_after`. Anyone who can fulfill the hash-lock condition can transfer the asset to themselves.

```

# Create a hash-locked asset without any owners_after
hashlock_tx = b.create_transaction(b.me, None, None, 'CREATE')

# Define a secret that will be hashed - fulfillments need to guess the secret
secret = b'much secret! wow!'
first_tx_condition = cc.PreimageSha256Fulfillment(preimage=secret)

# The conditions list is empty, so we need to append a new condition
hashlock_tx['transaction']['conditions'].append({
    'condition': {
        'uri': first_tx_condition.condition.serialize_uri()
    },
    'cid': 0,
    'owners_after': None
})

# Conditions have been updated, so the hash needs updating
hashlock_tx['id'] = util.get_hash_data(hashlock_tx)

# The asset needs to be signed by the owner_before
hashlock_tx_signed = b.sign_transaction(hashlock_tx, b.me_private)

# Some validations
assert b.validate_transaction(hashlock_tx_signed) == hashlock_tx_signed

b.write_transaction(hashlock_tx_signed)
hashlock_tx_signed

```

```

{
  "id": "604c520244b7ff63604527baf269e0cbfb887122f503703120fd347d6b99a237",
  "transaction": {
    "conditions": [
      {
        "cid": 0,
        "condition": {
          "uri": "cc:0:3:nsW2IiYgk9EUtsg4uBe3pBnOgRoAEX2IIsPgjqZz47U:17"
        }
      }
    ]
  }
}

```

```
        },
        "owners_after":None
    }
],
"data":None,
"fulfillments":[
    {
        "owners_before":[
            "FmLm6MxCABc8TsiZKdeYaZKo5yZWMM6Vty7Q1B6EgcP2"
        ],
        "fid":0,
        "fulfillment":"cf:4:21-D-LfNhIQhvY5914ArFTUGpgPKc7EVC1ZtJqqOTHGxlp9FuRr9tRfkbdtX2MZ",
        "input":None
    }
],
"operation":"CREATE",
"timestamp":"1461250387.910102"
},
"version":1
}
```

In order to redeem the asset, one needs to create a fulfillment with the correct secret:

```
hashlockuser_priv, hashlockuser_pub = crypto.generate_key_pair()

# Create hashlock fulfillment tx
hashlock_fulfill_tx = b.create_transaction(None, hashlockuser_pub, {'txid': hashlock_tx['id'], 'cid':

# Provide a wrong secret
hashlock_fulfill_tx_fulfillment = cc.PreimageSha256Fulfillment(preimage=b'')
hashlock_fulfill_tx['transaction']['fulfillments'][0]['fulfillment'] = \
    hashlock_fulfill_tx_fulfillment.serialize_uri()

assert b.is_valid_transaction(hashlock_fulfill_tx) == False

# Provide the right secret
hashlock_fulfill_tx_fulfillment = cc.PreimageSha256Fulfillment(preimage=secret)
hashlock_fulfill_tx['transaction']['fulfillments'][0]['fulfillment'] = \
    hashlock_fulfill_tx_fulfillment.serialize_uri()

assert b.validate_transaction(hashlock_fulfill_tx) == hashlock_fulfill_tx

b.write_transaction(hashlock_fulfill_tx)
hashlock_fulfill_tx
```

```
{
  "id":"fe6871bf3ca62eb61c52c5555cec2e07af51df817723f0cb76e5cf6248f449d2",
  "transaction":{
    "conditions":[
      {
        "cid":0,
        "condition":{
          "details":{
            "bitmask":32,
            "public_key":"EiqCKxnBCmmNb83qyGch48tULK9RLaEt4xFA43UVCVDb",
            "signature":None,
            "type":"fulfillment",
            "type_id":4
          },

```

```

        "uri": "cc:4:20:y9884Md2YI_wdnGSTJGhwvFaNsKLe8sqwimqk-2JLSI:96"
    },
    "owners_after": [
        "EiqCKxnBCmmNb83qyGch48tULK9RLaEt4xFA43UVCVDb"
    ]
  },
  ],
  "data": None,
  "fulfillments": [
    {
      "owners_before": [],
      "fid": 0,
      "fulfillment": "cf:0:bXVjaCBzZWNyZXQhIHdvdyE",
      "input": {
        "cid": 0,
        "txid": "604c520244b7ff63604527baf269e0cbfb887122f503703120fd347d6b99a237"
      }
    }
  ],
  "operation": "TRANSFER",
  "timestamp": "1461250397.944510"
},
"version": 1
}

```

Timeout Conditions

Timeout conditions allow assets to expire after a certain time. The primary use case of timeout conditions is to enable *Escrow*.

The condition can only be fulfilled before the expiry time. Once expired, the asset is lost and cannot be fulfilled by anyone.

Note: The timeout conditions are BigchainDB-specific and not (yet) supported by the ILP standard.

Caveat: The times between nodes in a BigchainDB federation may (and will) differ slightly. In this case, the majority of the nodes will decide.

```

# Create a timeout asset without any owners_after
tx_timeout = b.create_transaction(b.me, None, None, 'CREATE')

# Set expiry time - the asset needs to be transferred before expiration
time_sleep = 12
time_expire = str(float(util.timestamp()) + time_sleep) # 12 secs from now
condition_timeout = cc.TimeoutFulfillment(expire_time=time_expire)

# The conditions list is empty, so we need to append a new condition
tx_timeout['transaction']['conditions'].append({
    'condition': {
        'details': condition_timeout.to_dict(),
        'uri': condition_timeout.condition.serialize_uri()
    },
    'cid': 0,
    'owners_after': None
})

# Conditions have been updated, so the hash needs updating
tx_timeout['id'] = util.get_hash_data(tx_timeout)

```

```
# The asset needs to be signed by the owner_before
tx_timeout_signed = b.sign_transaction(tx_timeout, b.me_private)

# Some validations
assert b.validate_transaction(tx_timeout_signed) == tx_timeout_signed

b.write_transaction(tx_timeout_signed)
tx_timeout_signed
```

```
{
  "id": "78145396cd368f7168fb01c97aaf1df6f85244d7b544073dfcb42397dae38f90",
  "transaction": {
    "conditions": [
      {
        "cid": 0,
        "condition": {
          "details": {
            "bitmask": 9,
            "expire_time": "1464167910.643431",
            "type": "fulfillment",
            "type_id": 99
          },
          "uri": "cc:63:9:sceU_NZc3cAjAvaR1TVmgj7am5y8hJEBoqLm-tbqGbQ:17"
        },
        "owners_after": null
      }
    ],
    "data": null,
    "fulfillments": [
      {
        "owners_before": [
          "FmLm6MxCABc8TsiZKdeYaZKo5yZWMM6Vty7Q1B6EgcP2"
        ],
        "fid": 0,
        "fulfillment": null,
        "input": null
      }
    ],
    "operation": "CREATE",
    "timestamp": "1464167898.643353"
  },
  "version": 1
}
```

The following demonstrates that the transaction invalidates once the timeout occurs:

```
from time import sleep

# Create a timeout fulfillment tx
tx_timeout_transfer = b.create_transaction(None, testuser1_pub, {'txid': tx_timeout['id'], 'cid': 0},

# Parse the timeout condition and create the corresponding fulfillment
timeout_fulfillment = cc.Fulfillment.from_dict(
    tx_timeout['transaction']['conditions'][0]['condition']['details'])
tx_timeout_transfer['transaction']['fulfillments'][0]['fulfillment'] = timeout_fulfillment.serialize

# No need to sign transaction, like with hashlocks
```



```
# Small test to see the state change
for i in range(time_sleep - 4):
    tx_timeout_valid = b.is_valid_transaction(tx_timeout_transfer) == tx_timeout_transfer
    seconds_to_timeout = int(float(time_expire) - float(util.timestamp()))
    print('tx_timeout valid: {} ({}s to timeout)'.format(tx_timeout_valid, seconds_to_timeout))
    sleep(1)
```

If you were fast enough, you should see the following output:

```
tx_timeout valid: True (3s to timeout)
tx_timeout valid: True (2s to timeout)
tx_timeout valid: True (1s to timeout)
tx_timeout valid: True (0s to timeout)
tx_timeout valid: False (0s to timeout)
tx_timeout valid: False (-1s to timeout)
tx_timeout valid: False (-2s to timeout)
tx_timeout valid: False (-3s to timeout)
```

Escrow

Escrow is a mechanism for conditional release of assets.

This means that the assets are locked up by a trusted party until an `execute` condition is presented. In order not to tie up the assets forever, the escrow foresees an `abort` condition, which is typically an expiry time.

BigchainDB and cryptoconditions provides escrow out-of-the-box, without the need of a trusted party.

A threshold condition is used to represent the escrow, since BigchainDB transactions cannot have a *pending* state.

The logic for switching between `execute` and `abort` conditions is conceptually simple:

```
if timeout_condition.validate(utcnow()):
    execute_fulfillment.validate(msg) == True
    abort_fulfillment.validate(msg) == False
else:
    execute_fulfillment.validate(msg) == False
    abort_fulfillment.validate(msg) == True
```

The above switch can be implemented as follows using threshold cryptoconditions:

The inverted timeout is denoted by a -1 threshold, which negates the output of the fulfillment.

```
inverted_fulfillment.validate(msg) == not fulfillment.validate(msg)
```

Note: inverted thresholds are BigchainDB-specific and not supported by the ILP standard. The main reason is that it's difficult to tell whether the fulfillment was negated, or just omitted.

The following code snippet shows how to create an escrow condition:

```
# Retrieve the last transaction of testuser2_pub (or create a new asset)
tx_retrieved_id = b.get_owned_ids(testuser2_pub).pop()

# Create a base template with the execute and abort address
tx_escrow = b.create_transaction(testuser2_pub, [testuser2_pub, testuser1_pub], tx_retrieved_id, 'TR

# Set expiry time - the execute address needs to fulfill before expiration
time_sleep = 12
time_expire = str(float(util.timestamp()) + time_sleep) # 12 secs from now

# Create the escrow and timeout condition
```

```

condition_escrow = cc.ThresholdSha256Fulfillment(threshold=1) # OR Gate
condition_timeout = cc.TimeoutFulfillment(expire_time=time_expire) # only valid if now() <= time_expire
condition_timeout_inverted = cc.InvertedThresholdSha256Fulfillment(threshold=1)
condition_timeout_inverted.add_subfulfillment(condition_timeout) # invert the timeout condition

# Create the execute branch
condition_execute = cc.ThresholdSha256Fulfillment(threshold=2) # AND gate
condition_execute.add_subfulfillment(cc.Ed25519Fulfillment(public_key=testuser1_pub)) # execute add
condition_execute.add_subfulfillment(condition_timeout) # federation checks on expiry
condition_escrow.add_subfulfillment(condition_execute)

# Create the abort branch
condition_abort = cc.ThresholdSha256Fulfillment(threshold=2) # AND gate
condition_abort.add_subfulfillment(cc.Ed25519Fulfillment(public_key=testuser2_pub)) # abort address
condition_abort.add_subfulfillment(condition_timeout_inverted)
condition_escrow.add_subfulfillment(condition_abort)

# Update the condition in the newly created transaction
tx_escrow['transaction']['conditions'][0]['condition'] = {
    'details': condition_escrow.to_dict(),
    'uri': condition_escrow.condition.serialize_uri()
}

# Conditions have been updated, so the hash needs updating
tx_escrow['id'] = util.get_hash_data(tx_escrow)

# The asset needs to be signed by the owner before
tx_escrow_signed = b.sign_transaction(tx_escrow, testuser2_priv)

# Some validations
assert b.validate_transaction(tx_escrow_signed) == tx_escrow_signed

b.write_transaction(tx_escrow_signed)
tx_escrow_signed

```

```

{
  "id": "1a281da2b9bc3d2beba92479058d440de3353427fd64045a61737bad0d0c809c",
  "transaction": {
    "conditions": [
      {
        "cid": 0,
        "condition": {
          "details": {
            "bitmask": 41,
            "subfulfillments": [
              {
                "bitmask": 41,
                "subfulfillments": [
                  {
                    "bitmask": 32,
                    "public_key": "qv8DvdNG5nZHWCP5aPSggqxAvapJpQj19abRvFCntor",
                    "signature": null,
                    "type": "fulfillment",
                    "type_id": 4,
                    "weight": 1
                  },
                  {
                    "bitmask": 9,

```

```

        "expire_time": "1464242352.227917",
        "type": "fulfillment",
        "type_id": 99,
        "weight": 1
    },
    {
        "threshold": 2,
        "type": "fulfillment",
        "type_id": 2,
        "weight": 1
    },
    {
        "bitmask": 41,
        "subfulfillments": [
            {
                "bitmask": 32,
                "public_key": "BwuhqQX8FPsmqYiRV2CSZYWWsSWgSSQQFHjqxKEuqkPs",
                "signature": null,
                "type": "fulfillment",
                "type_id": 4,
                "weight": 1
            },
            {
                "bitmask": 9,
                "subfulfillments": [
                    {
                        "bitmask": 9,
                        "expire_time": "1464242352.227917",
                        "type": "fulfillment",
                        "type_id": 99,
                        "weight": 1
                    }
                ],
                "threshold": 1,
                "type": "fulfillment",
                "type_id": 98,
                "weight": 1
            }
        ],
        "threshold": 2,
        "type": "fulfillment",
        "type_id": 2,
        "weight": 1
    },
    {
        "threshold": 1,
        "type": "fulfillment",
        "type_id": 2
    },
    {
        "uri": "cc:2:29:sg08ERtpqRgxot7mu7XMdNkZTc29xCbWE1r8DgxuL8:181"
    },
    {
        "owners_after": [
            "BwuhqQX8FPsmqYiRV2CSZYWWsSWgSSQQFHjqxKEuqkPs",
            "qv8DvdNG5nZHWCP5aPSqgqxAvAPJpJ19abRvFCntor"
        ]
    },
    {
        "data": null,

```

```
    "fulfillments": [
      {
        "owners_before": [
          "qv8DvdNG5nZHWCP5aPSqgqxAvAPJpQj19abRvFCntor"
        ],
        "fid": 0,
        "fulfillment": "cf:4:B6VAa7KAMD1v-pyvDx9RuBLb6l2Qs3vhucgXqzU_RbuRucOp6tNY8AoNMoC-HAOZl",
        "input": {
          "cid": 1,
          "txid": "d3f5e78f6d4346466178745f1c01cbcaf1c1dce1932a16cd653051b16ee29bac"
        }
      }
    ],
    "operation": "TRANSFER",
    "timestamp": "1464242340.227787"
  },
  "version": 1
}
```

At any given moment `testuser1` and `testuser2` can try to fulfill the `execute` and `abort` branch respectively. Whether the fulfillment will validate depends on the timeout condition.

We'll illustrate this by example.

In the case of `testuser1`, we create the `execute` fulfillment:

```
# Create a base template for execute fulfillment
tx_escrow_execute = b.create_transaction([testuser2_pub, testuser1_pub], testuser1_pub, {'txid': tx_escrow_execute_txid})

# Parse the Escrow cryptocondition
escrow_fulfillment = cc.Fulfillment.from_dict(
    tx_escrow['transaction']['conditions'][0]['condition']['details'])

subfulfillment_testuser1 = escrow_fulfillment.get_subcondition_from_vk(testuser1_pub)[0]
subfulfillment_testuser2 = escrow_fulfillment.get_subcondition_from_vk(testuser2_pub)[0]
subfulfillment_timeout = escrow_fulfillment.subconditions[0]['body'].subconditions[1]['body']
subfulfillment_timeout_inverted = escrow_fulfillment.subconditions[1]['body'].subconditions[1]['body']

# Get the fulfillment message to sign
tx_escrow_execute_fulfillment_message = \
    util.get_fulfillment_message(tx_escrow_execute,
                                tx_escrow_execute['transaction']['fulfillments'][0],
                                serialized=True)

# Clear the subconditions of the escrow fulfillment
escrow_fulfillment.subconditions = []

# Fulfill the execute branch
fulfillment_execute = cc.ThresholdSha256Fulfillment(threshold=2)
subfulfillment_testuser1.sign(tx_escrow_execute_fulfillment_message, crypto.SigningKey(testuser1_privkey))
fulfillment_execute.add_subfulfillment(subfulfillment_testuser1)
fulfillment_execute.add_subfulfillment(subfulfillment_timeout)
escrow_fulfillment.add_subfulfillment(fulfillment_execute)

# Do not fulfill the abort branch
condition_abort = cc.ThresholdSha256Fulfillment(threshold=2)
condition_abort.add_subfulfillment(subfulfillment_testuser2)
condition_abort.add_subfulfillment(subfulfillment_timeout_inverted)
escrow_fulfillment.add_subcondition(condition_abort.condition) # Adding only the condition here
```

```
# Update the execute transaction with the fulfillment
tx_escrow_execute['transaction']['fulfillments'][0]['fulfillment'] = escrow_fulfillment.serialize_uri()
```

In the case of testuser2, we create the abort fulfillment:

```
# Create a base template for execute fulfillment
tx_escrow_abort = b.create_transaction([testuser2_pub, testuser1_pub], testuser2_pub, {'txid': tx_escrow_execute['transaction']['txid']})

# Parse the threshold cryptocondition
escrow_fulfillment = cc.Fulfillment.from_dict(
    tx_escrow['transaction']['conditions'][0]['condition']['details'])

subfulfillment_testuser1 = escrow_fulfillment.get_subcondition_from_vk(testuser1_pub)[0]
subfulfillment_testuser2 = escrow_fulfillment.get_subcondition_from_vk(testuser2_pub)[0]
subfulfillment_timeout = escrow_fulfillment.subconditions[0]['body'].subconditions[1]['body']
subfulfillment_timeout_inverted = escrow_fulfillment.subconditions[1]['body'].subconditions[1]['body']

# Get the fulfillment message to sign
tx_escrow_abort_fulfillment_message = \
    util.get_fulfillment_message(tx_escrow_abort,
                                tx_escrow_abort['transaction']['fulfillments'][0],
                                serialized=True)

# Clear the subconditions of the escrow fulfillment
escrow_fulfillment.subconditions = []

# Do not fulfill the execute branch
condition_execute = cc.ThresholdSha256Fulfillment(threshold=2)
condition_execute.add_subfulfillment(subfulfillment_testuser1)
condition_execute.add_subfulfillment(subfulfillment_timeout)
escrow_fulfillment.add_subcondition(condition_execute.condition) # Adding only the condition here

# Fulfill the abort branch
fulfillment_abort = cc.ThresholdSha256Fulfillment(threshold=2)
subfulfillment_testuser2.sign(tx_escrow_abort_fulfillment_message, crypto.SigningKey(testuser2_priv))
fulfillment_abort.add_subfulfillment(subfulfillment_testuser2)
fulfillment_abort.add_subfulfillment(subfulfillment_timeout_inverted)
escrow_fulfillment.add_subfulfillment(fulfillment_abort)

# Update the abort transaction with the fulfillment
tx_escrow_abort['transaction']['fulfillments'][0]['fulfillment'] = escrow_fulfillment.serialize_uri()
```

The following demonstrates that the transaction validation switches once the timeout occurs:

```
for i in range(time_sleep - 4):
    valid_execute = b.is_valid_transaction(tx_escrow_execute) == tx_escrow_execute
    valid_abort = b.is_valid_transaction(tx_escrow_abort) == tx_escrow_abort

    seconds_to_timeout = int(float(time_expire) - float(util.timestamp()))
    print('tx_execute valid: {} - tx_abort valid {} ({}s to timeout)'.format(valid_execute, valid_abort, seconds_to_timeout))
    sleep(1)
```

If you execute in a timely fashion, you should see the following:

```
tx_execute valid: True - tx_abort valid False (3s to timeout)
tx_execute valid: True - tx_abort valid False (2s to timeout)
tx_execute valid: True - tx_abort valid False (1s to timeout)
tx_execute valid: True - tx_abort valid False (0s to timeout)
tx_execute valid: False - tx_abort valid True (0s to timeout)
```

```
tx_execute valid: False - tx_abort valid True (-1s to timeout)
tx_execute valid: False - tx_abort valid True (-2s to timeout)
tx_execute valid: False - tx_abort valid True (-3s to timeout)
```

Of course, when the `execute` transaction was accepted in-time by `bigchaindb`, then writing the `abort` transaction after expiry will yield a `Doublespend` error.

1.7.2 The HTTP Client-Server API

Note: The HTTP client-server API is currently quite rudimentary. For example, there is no ability to do complex queries using the HTTP API. We plan to add querying capabilities in the future. If you want to build a full-featured proof-of-concept, we suggest you use [the Python Server API](#) for now.

When you start `Bigchaindb` using `bigchaindb start`, an HTTP API is exposed at the address stored in the `BigchainDB` node configuration settings. The default is for that address to be:

`http://localhost:9984/api/v1/`

but that address can be changed by changing the “API endpoint” configuration setting (e.g. in a local config file). There’s more information about setting the API endpoint in [the section about BigchainDB Configuration Settings](#).

There are other configuration settings related to the web server (serving the HTTP API). In particular, the default is for the web server socket to bind to `localhost:9984` but that can be changed (e.g. to `0.0.0.0:9984`). For more details, see the “server” settings (“bind”, “workers” and “threads”) in [the section about BigchainDB Configuration Settings](#).

The HTTP API currently exposes two endpoints, one to get information about a specific transaction, and one to push a new transaction to the `BigchainDB` cluster.

GET /transactions/{tx_id}

Get the transaction with the ID `tx_id`.

This endpoint returns only a transaction from a `VALID` or `UNDECIDED` block on `bigchain`, if exists.

Parameters

- `tx_id` (*hex string*) – transaction ID

Example request:

```
GET /transactions/7ad5a4b83bc8c70c4fd7420ff3c60693ab8e6d0e3124378ca69ed5acd2578792 HTTP/1.1
Host: example.com
```

Example response:

```
HTTP/1.1 200 OK
Content-Type: application/json

{
  "id": "7ad5a4b83bc8c70c4fd7420ff3c60693ab8e6d0e3124378ca69ed5acd2578792",
  "transaction": {
    "conditions": [
      {
        "cid": 0,
        "condition": {
          "details": {
            "bitmask": 32,
            "public_key": "CwA8s2QYQBfNz4WvjEwmJi83zYr7JhxRhidx6uZ5KBVd",
            "signature": null,
            "type": "fulfillment",
            "type_id": 4
          }
        }
      }
    ]
  }
}
```

```

    },
    "uri": "cc:4:20:sVA_3p8gv18yRFNTomqm6MaavKewka6dGYcFAuPrRXQ:96"
  },
  "owners_after": [
    "CwA8s2QYQBfNz4WvjEwmJi83zYr7JhxRhidx6uZ5KBVd"
  ]
},
],
"data": {
  "payload": null,
  "uuid": "a9999d69-6cde-4b80-819d-ed57f6abe257"
},
"fulfillments": [
  {
    "owners_before": [
      "JEAkJqLbbgDRAtMm8YAjGp759Aq2qTn9eaEHUj2XePE"
    ],
    "fid": 0,
    "fulfillment": "cf:4:___Y_Um6H73iwPe6ejWXEw930SQhqVGjtAHTXilPp0P01vE_Cx6zs3GJV0l jhP",
    "input": {
      "cid": 0,
      "txid": "598ce4e9a29837a1c6fc337ee4a41b61c20ad25d01646754c825b1116abd8761"
    }
  }
],
"operation": "TRANSFER",
"timestamp": "1471423869",
"version": 1
}
}

```

Status Codes

- **200 OK** – A transaction with that ID was found.
- **404 Not Found** – A transaction with that ID was not found.

GET /transactions/{tx_id}/status

Get the status of a transaction with the ID `tx_id`.

This endpoint returns the status of a transaction if exists.

Possible values are `valid`, `invalid`, `undecided` or `backlog`.

Parameters

- **tx_id** (*hex string*) – transaction ID

Example request:

```
GET /transactions/7ad5a4b83bc8c70c4fd7420ff3c60693ab8e6d0e3124378ca69ed5acd2578792/status HTTP/1.1
Host: example.com
```

Example response:

```
HTTP/1.1 200 OK
Content-Type: application/json

{
  "status": "valid"
}
```

Status Codes

- **200 OK** – A transaction with that ID was found and the status is returned.
- **404 Not Found** – A transaction with that ID was not found.

POST /transactions/

Push a new transaction.

Example request:

```
POST /transactions/ HTTP/1.1
Host: example.com
Content-Type: application/json

{
  "id": "7ad5a4b83bc8c70c4fd7420ff3c60693ab8e6d0e3124378ca69ed5acd2578792",
  "transaction": {
    "conditions": [
      {
        "cid": 0,
        "condition": {
          "details": {
            "bitmask": 32,
            "public_key": "CwA8s2QYQBfNz4WvjEwmJi83zYr7JhXRhidx6uZ5KBVd",
            "signature": null,
            "type": "fulfillment",
            "type_id": 4
          },
          "uri": "cc:4:20:sVA_3p8gvl8yRFNTomqm6MaavKewka6dGYcFAuPrRXQ:96"
        },
        "owners_after": [
          "CwA8s2QYQBfNz4WvjEwmJi83zYr7JhXRhidx6uZ5KBVd"
        ]
      }
    ],
    "data": {
      "payload": null,
      "uuid": "a9999d69-6cde-4b80-819d-ed57f6abe257"
    },
    "fulfillments": [
      {
        "owners_before": [
          "JEAkJqLbbgDRAtMm8YAjGp759Aq2qTn9eaEHUj2XePE"
        ],
        "fid": 0,
        "fulfillment": "cf:4:___Y_Um6H73iwPe6ejWXEw930SQhqVGjtAHTXilPp0P01vE_Cx6zs3GJV0ljhP",
        "input": {
          "cid": 0,
          "txid": "598ce4e9a29837a1c6fc337ee4a41b61c20ad25d01646754c825b1116abd8761"
        }
      }
    ],
    "operation": "TRANSFER",
    "timestamp": "1471423869",
    "version": 1
  }
}
```



```
}
}
```

Example response:

```
HTTP/1.1 201 Created
Content-Type: application/json

{
  "assignee": "4XYfCbabAWVUCbjTmRTFEu2sc3dFEdkse4r6X498Bls8",
  "id": "7ad5a4b83bc8c70c4fd7420ff3c60693ab8e6d0e3124378ca69ed5acd2578792",
  "transaction": {
    "conditions": [
      {
        "cid": 0,
        "condition": {
          "details": {
            "bitmask": 32,
            "public_key": "CwA8s2QYQBfNz4WvjEwmJi83zYr7JhxRhidx6uZ5KBVd",
            "signature": null,
            "type": "fulfillment",
            "type_id": 4
          },
          "uri": "cc:4:20:sVA_3p8gvl8yRFNTomqm6MaavKewka6dGYcFAuPrRXQ:96"
        },
        "owners_after": [
          "CwA8s2QYQBfNz4WvjEwmJi83zYr7JhxRhidx6uZ5KBVd"
        ]
      }
    ],
    "data": {
      "payload": null,
      "uuid": "a9999d69-6cde-4b80-819d-ed57f6abe257"
    },
    "fulfillments": [
      {
        "owners_before": [
          "JEAkJqLbbgDRAtMm8YAjGp759Aq2qTn9eaEHUj2XePE"
        ],
        "fid": 0,
        "fulfillment": "cf:4:___Y_Um6H73iwPe6ejWXEw930SQhqVGjtAHTXilPp0P01vE_Cx6zs3GJV01jhP",
        "input": {
          "cid": 0,
          "txid": "598ce4e9a29837a1c6fc337ee4a41b61c20ad25d01646754c825b1116abd8761"
        }
      }
    ],
    "operation": "TRANSFER",
    "timestamp": "1471423869",
    "version": 1
  }
}
```

Status Codes

- 201 Created – A new transaction was created.
- 400 Bad Request – The transaction was invalid and not created.

Disclaimer

CREATE transactions are treated differently from TRANSFER assets. The reason is that a CREATE transaction needs to be signed by a federation node and not by the client.

The following python snippet in a client can be used to generate CREATE transactions before they can be pushed to the API server:

```
from bigchaindb import util
tx = util.create_and_sign_tx(my_privkey, my_pubkey, my_pubkey, None, 'CREATE')
```

When POSTing tx to the API, the CREATE transaction will be signed by a federation node.

A TRANSFER transaction, that takes an existing input transaction to change ownership can be generated in multiple ways:

```
from bigchaindb import util, Bigchain
tx = util.create_and_sign_tx(my_privkey, my_pubkey, other_pubkey, input_tx, 'TRANSFER')
# or
b = Bigchain()
tx_unsigned = b.create_transaction(my_pubkey, other_pubkey, input_tx, 'TRANSFER')
tx = b.sign_transaction(tx_unsigned, my_privkey)
```

More information on generating transactions can be found in the Python server API examples

1.7.3 The Python Driver API by Example

The Python driver API is used by app developers to develop client apps which can communicate with one or more BigchainDB clusters. Under the hood, the Python driver API communicates with the BigchainDB cluster using the BigchainDB HTTP client-server API.

Here's an example of how to use the Python driver API. First, launch an interactive Python session, then:

```
In [1]: from bigchaindb.client import temp_client
In [2]: c1 = temp_client()
In [3]: c2 = temp_client()
In [4]: tx1 = c1.create()
In [5]: tx1
Out[5]:
{'assignee': '3NsvDXiiuf2BRPnqfRuBM9yHNjsH4L33gcZ4rh4GMY2J',
'id': '00f530d210c06671ab2de4330e3e2cf0d0b47b2826302ee25ceea9b2f47b097f',
'transaction': {'conditions': [{'cid': 0,
'condition': {'details': {'bitmask': 32,
'public_key': '9FGRd2jLxmwtRkwsWTPeQy1rZpg6ycuT7NwmCR4QV3',
'signature': None,
'type': 'fulfillment',
'type_id': 4},
'uri': 'cc:4:20:eoUROTxUArrpXGVBBrvrYqkcEGG81B_leliNvSvSddDg:96'},
'owners_after': ['9FGRd2jLxmwtRkwsWTPeQy1rZpg6ycuT7NwmCR4QV3']}]},
'data': {'payload': None, 'uuid': 'b4884e37-3c8e-4cc2-bfc8-68a05ed090ad'},
'fulfillments': [{'owners_before': ['3NsvDXiiuf2BRPnqfRuBM9yHNjsH4L33gcZ4rh4GMY2J'],
'fid': 0,
'fulfillment': 'cf:4:I1IkuhCSf_hGqJ-JKHTQIO1g4apbQuaZXNMEX4isyxd7azkJreyGKYaMLs6Xk9kxQClwz1nQiKM',
'input': None}],
'operation': 'CREATE',
'timestamp': '1466676327'},
'version': 1}

In [6]: c1.transfer(c2.public_key, {'txid': tx1['id'], 'cid': 0})
```

```

Out[6]:
{'assignee': '3NsvDXiiuf2BRPnqfRuBM9yHNjsH4L33gcZ4rh4GMY2J',
 'id': 'd6da7b42c1d82b6a14514bef5c919e7b21e77bc32d537993bc4e5c98d1885e1d',
 'transaction': {'conditions': [{'cid': 0,
   'condition': {'details': {'bitmask': 32,
     'public_key': '89tbMBospYsTNDgpqFS4RLszNsxuE4JEumNuY3WTAnT5',
     'signature': None,
     'type': 'fulfillment',
     'type_id': 4},
     'uri': 'cc:4:20:akjKWxLO2hbe6RVva_FsWNDJmnUKYjQ57HIhUQbwb2Q:96'},
     'owners_after': ['89tbMBospYsTNDgpqFS4RLszNsxuE4JEumNuY3WTAnT5']}]},
 'data': {'payload': None, 'uuid': 'a640a9d6-9384-4e9c-a130-e899ea6416aa'},
 'fulfillments': [{'owners_before': ['9FGRd2jLxmwtRkwsWTPeoqylrZpg6ycuT7NwmCR4QV3'],
   'fid': 0,
   'fulfillment': 'cf:4:eoUROTxUArrpXGVBrvrYqkcEGG8lB_leliNvSvSddDgVmY6O7YTER04mWjAVd6m0qOv5R44Cxp',
   'input': {'cid': 0,
     'txid': '00f530d210c06671ab2de4330e3e2cf0d0b47b2826302ee25ceea9b2f47b097f'}]}],
 'operation': 'TRANSFER',
 'timestamp': '1466676463'},
 'version': 1}

```

1.7.4 Example Apps

There are some example BigchainDB apps (i.e. apps which use BigchainDB) in the [GitHub repository](#) named `bigchaindb-examples`.

As noted there, the examples are for demonstration purposes only and should not be used as-is for production.

1.8 Clusters & Federations

1.8.1 Set Up a Federation

This section is about how to set up a BigchainDB *federation*, where each node is operated by a different operator. If you want to set up and run a testing cluster on AWS (where all nodes are operated by you), then see the section about that.

Initial Checklist

- Do you have a governance process for making federation-level decisions, such as how to admit new members?
- What will you store in creation transactions (data payload)? Is there a data schema?
- Will you use transfer transactions? Will they include a non-empty data payload?
- Who will be allowed to submit transactions? Who will be allowed to read or query transactions? How will you enforce the access rules?

Set Up the Initial Cluster

The federation must decide some things before setting up the initial cluster (initial set of BigchainDB nodes):

1. Who will operate a node in the initial cluster?
2. What will the replication factor be? (It must be 3 or more for [RethinkDB failover](#) to work.)

3. Which node will be responsible for sending the commands to configure the RethinkDB database?

Once those things have been decided, each node operator can begin setting up their BigchainDB (production) node.

Each node operator will eventually need two pieces of information from all other nodes in the federation:

1. Their RethinkDB hostname, e.g. `rdb.farm2.organization.org`
2. Their BigchainDB public key, e.g. `Eky3nkbxDTMgkmiJC8i5hKyVFIAQNmPP4a2G4JdDxJCK`

1.8.2 Backing Up & Restoring Data

There are several ways to backup and restore the data in a BigchainDB cluster.

RethinkDB's Replication as a form of Backup

RethinkDB already has internal replication: every document is stored on R different nodes, where R is the replication factor (set using `bigchaindb set-replicas R`). Those replicas can be thought of as “live backups” because if one node goes down, the cluster will continue to work and no data will be lost.

At this point, there should be someone saying, “But replication isn’t backup!”

It’s true. Replication alone isn’t enough, because something bad might happen *inside* the database, and that could affect the replicas. For example, what if someone logged in as a RethinkDB admin and did a “drop table”? We currently plan for each node to be protected by a next-generation firewall (or something similar) to prevent such things from getting very far. For example, see [issue #240](#).

Nevertheless, you should still consider having normal, “cold” backups, because bad things can still happen.

Live Replication of RethinkDB Data Files

Each BigchainDB node stores its subset of the RethinkDB data in one directory. You could set up the node’s file system so that directory lives on its own hard drive. Furthermore, you could make that hard drive part of a [RAID](#) array, so that a second hard drive would always have a copy of the original. If the original hard drive fails, then the second hard drive could take its place and the node would continue to function. Meanwhile, the original hard drive could be replaced.

That’s just one possible way of setting up the file system so as to provide extra reliability.

Another way to get similar reliability would be to mount the RethinkDB data directory on an [Amazon EBS](#) volume. Each Amazon EBS volume is, “automatically replicated within its Availability Zone to protect you from component failure, offering high availability and durability.”

See the section on setting up storage for RethinkDB for more details.

As with shard replication, live file-system replication protects against many failure modes, but it doesn’t protect against them all. You should still consider having normal, “cold” backups.

rethinkdb dump (to a File)

RethinkDB can create an archive of all data in the cluster (or all data in specified tables), as a compressed file. According to [the RethinkDB blog post when that functionality became available](#):

Since the backup process is using client drivers, it automatically takes advantage of the MVCC [multiversion concurrency control] functionality built into RethinkDB. It will use some cluster resources, but will not lock out any of the clients, so you can safely run it on a live cluster.

To back up all the data in a BigchainDB cluster, the RethinkDB admin user must run a command like the following on one of the nodes:

```
rethinkdb dump -e bigchain.bigchain -e bigchain.votes
```

That should write a file named `rethinkdb_dump_<date>_<time>.tar.gz`. The `-e` option is used to specify which tables should be exported. You probably don't need to export the backlog table, but you definitely need to export the `bigchain` and `votes` tables. `bigchain.votes` means the `votes` table in the RethinkDB database named `bigchain`. It's possible that your database has a different name: the database name is a BigchainDB configuration setting. The default name is `bigchain`. (Tip: you can see the values of all configuration settings using the `bigchaindb show-config` command.)

There's [more information about the `rethinkdb dump` command in the RethinkDB documentation](#). It also explains how to restore data to a cluster from an archive file.

Notes

- If the `rethinkdb dump` subcommand fails and the last line of the Traceback says “NameError: name ‘file’ is not defined”, then you need to update your RethinkDB Python driver; do a `pip install --upgrade rethinkdb`
- It might take a long time to backup data this way. The more data, the longer it will take.
- You need enough free disk space to store the backup file.
- If a document changes after the backup starts but before it ends, then the changed document may not be in the final backup. This shouldn't be a problem for BigchainDB, because blocks and votes can't change anyway.
- `rethinkdb dump` saves data and secondary indexes, but does *not* save cluster metadata. You will need to recreate your cluster setup yourself after you run `rethinkdb restore`.
- RethinkDB also has [subcommands to import/export](#) collections of JSON or CSV files. While one could use those for backup/restore, it wouldn't be very practical.

Client-Side Backup

In the future, it will be possible for clients to query for the blocks containing the transactions they care about, and for the votes on those blocks. They could save a local copy of those blocks and votes.

How could we be sure blocks and votes from a client are valid?

All blocks and votes are signed by federation nodes. Only federation nodes can produce valid signatures because only federation nodes have the necessary private keys. A client can't produce a valid signature for a block or vote.

Could we restore an entire BigchainDB database using client-saved blocks and votes?

Yes, in principle, but it would be difficult to know if you've recovered every block and vote. Votes link to the block they're voting on and to the previous block, so one could detect some missing blocks. It would be difficult to know if you've recovered all the votes.

Backup by Copying RethinkDB Data Files

It's *possible* to back up a BigchainDB database by creating a point-in-time copy of the RethinkDB data files (on all nodes, at roughly the same time). It's not a very practical approach to backup: the resulting set of files will be much larger (collectively) than what one would get using `rethinkdb dump`, and there are no guarantees on how consistent that data will be, especially for recently-written data.

If you're curious about what's involved, see the [MongoDB documentation about “Backup by Copying Underlying Data Files”](#). (Yes, that's documentation for MongoDB, but the principles are the same.)

See the last subsection of this page for a better way to use this idea.

Incremental or Continuous Backup

Incremental backup is where backup happens on a regular basis (e.g. daily), and each one only records the changes since the last backup.

Continuous backup might mean incremental backup on a very regular basis (e.g. every ten minutes), or it might mean backup of every database operation as it happens. The latter is also called transaction logging or continuous archiving.

At the time of writing, RethinkDB didn't have a built-in incremental or continuous backup capability, but the idea was raised in RethinkDB issues [#89](#) and [#5890](#). On July 5, 2016, Daniel Mewes (of RethinkDB) wrote the following comment on issue [#5890](#): “We would like to add this feature [continuous backup], but haven't started working on it yet.”

To get a sense of what continuous backup might look like for RethinkDB, one can look at the continuous backup options available for MongoDB. MongoDB, the company, offers continuous backup with [Ops Manager](#) (self-hosted) or [Cloud Manager](#) (fully managed). Features include:

- It “continuously maintains backups, so if your MongoDB deployment experiences a failure, the most recent backup is only moments behind...”
- It “offers point-in-time backups of replica sets and cluster-wide snapshots of sharded clusters. You can restore to precisely the moment you need, quickly and safely.”
- “You can rebuild entire running clusters, just from your backups.”
- It enables, “fast and seamless provisioning of new dev and test environments.”

The MongoDB documentation has more [details about how Ops Manager Backup works](#).

Considerations for BigchainDB:

- We'd like the cost of backup to be low. To get a sense of the cost, MongoDB Cloud Manager backup [costed \\$30 / GB / year prepaid](#). One thousand gigabytes backed up (i.e. about a terabyte) would cost 30 thousand US dollars per year. (That's just for the backup; there's also a cost per server per year.)
- We'd like the backup to be decentralized, with no single point of control or single point of failure. (Note: some file systems have a single point of failure. For example, HDFS has one Namenode.)
- We only care to back up blocks and votes, and once written, those never change. There are no updates or deletes, just new blocks and votes.

Combining RethinkDB Replication with Storage Snapshots

Although it's not advertised as such, RethinkDB's built-in replication feature is similar to continous backup, except the “backup” (i.e. the set of replica shards) is spread across all the nodes. One could take that idea a bit farther by creating a set of backup-only servers with one full backup:

- Give all the original BigchainDB nodes (RethinkDB nodes) the server tag `original`. This is the default if you used the RethinkDB config file suggested in the section titled [Configure RethinkDB Server](#).
- Set up a group of servers running RethinkDB only, and give them the server tag `backup`. The backup servers could be geographically separated from all the `original` nodes (or not; it's up to the federation).
- Clients shouldn't be able to read from or write to servers in the `backup` set.
- Send a RethinkDB reconfigure command to the RethinkDB cluster to make it so that the `original` set has the same number of replicas as before (or maybe one less), and the `backup` set has one replica. Also, make sure the `primary_replica_tag='original'` so that all primary shards live on the `original` nodes.

The [RethinkDB documentation on sharding and replication](#) has the details of how to set server tags and do RethinkDB reconfiguration.

Once you’ve set up a set of backup-only RethinkDB servers, you could make a point-in-time snapshot of their storage devices, as a form of backup.

You might want to disconnect the backup set from the original set first, and then wait for reads and writes in the backup set to stop. (The backup set should have only one copy of each shard, so there’s no opportunity for inconsistency between shards of the backup set.)

You will want to re-connect the backup set to the original set as soon as possible, so it’s able to catch up.

If something bad happens to the entire original BigchainDB cluster (including the backup set) and you need to restore it from a snapshot, you can, but before you make BigchainDB live, you should 1) delete all entries in the backlog table, 2) delete all blocks after the last voted-valid block, 3) delete all votes on the blocks deleted in part 2, and 4) rebuild the RethinkDB indexes.

NOTE: Sometimes snapshots are *incremental*. For example, [Amazon EBS snapshots](#) are incremental, meaning “only the blocks on the device that have changed after your most recent snapshot are saved. **This minimizes the time required to create the snapshot and saves on storage costs.**” [Emphasis added]

1.8.3 Deploy a Testing Cluster on AWS

This section explains a way to deploy a cluster of BigchainDB nodes on Amazon Web Services (AWS) for testing purposes.

Why?

Why would anyone want to deploy a centrally-controlled BigchainDB cluster? Isn’t BigchainDB supposed to be decentralized, where each node is controlled by a different person or organization?

Yes! These scripts are for deploying a testing cluster, not a production cluster.

How?

We use some Bash and Python scripts to launch several instances (virtual servers) on Amazon Elastic Compute Cloud (EC2). Then we use Fabric to install RethinkDB and BigchainDB on all those instances.

Python Setup

The instructions that follow have been tested on Ubuntu 14.04, but may also work on similar distros or operating systems.

Note: Our Python scripts for deploying to AWS use Python 2 because Fabric doesn’t work with Python 3.

Maybe create a Python 2 virtual environment and activate it. Then install the following Python packages (in that virtual environment):

```
pip install fabric fabtools requests boto3 awscli
```

What did you just install?

- “[Fabric](#) is a Python (2.5-2.7) library and command-line tool for streamlining the use of SSH for application deployment or systems administration tasks.”
- [fabtools](#) are “tools for writing awesome Fabric files”

- `requests` is a Python package/library for sending HTTP requests
- “`Boto` is the Amazon Web Services (AWS) SDK for Python, which allows Python developers to write software that makes use of Amazon services like S3 and EC2.” (`boto3` is the name of the latest Boto package.)
- The `aws-cli` package, which is an AWS Command Line Interface (CLI).

Basic AWS Setup

See the page about basic AWS Setup in the Appendices.

Get Enough Amazon Elastic IP Addresses

The AWS cluster deployment scripts use elastic IP addresses (although that may change in the future). By default, AWS accounts get five elastic IP addresses. If you want to deploy a cluster with more than five nodes, then you will need more than five elastic IP addresses; you may have to apply for those; see [the AWS documentation on elastic IP addresses](#).

Create an Amazon EC2 Security Group

Go to the AWS EC2 Console and select “Security Groups” in the left sidebar. Click the “Create Security Group” button. Name it `bigchaindb`. The description probably doesn’t matter; you can also put `bigchaindb` for that.

Add these rules for Inbound traffic:

- Type = All TCP, Protocol = TCP, Port Range = 0-65535, Source = 0.0.0.0/0
- Type = SSH, Protocol = SSH, Port Range = 22, Source = 0.0.0.0/0
- Type = All UDP, Protocol = UDP, Port Range = 0-65535, Source = 0.0.0.0/0
- Type = All ICMP, Protocol = ICMP, Port Range = 0-65535, Source = 0.0.0.0/0

Note: These rules are extremely lax! They’re meant to make testing easy. For example, Source = 0.0.0.0/0 is [CIDR notation](#) for “allow this traffic to come from *any* IP address.”

Deploy a BigchainDB Monitor

This step is optional.

One way to monitor a BigchainDB cluster is to use the monitoring setup described in the Monitoring section of this documentation. If you want to do that, then you may want to deploy the monitoring server first, so you can tell your BigchainDB nodes where to send their monitoring data.

You can deploy a monitoring server on AWS. To do that, go to the AWS EC2 Console and launch an instance:

1. Choose an AMI: select Ubuntu Server 14.04 LTS.
2. Choose an Instance Type: a `t2.micro` will suffice.
3. Configure Instance Details: you can accept the defaults, but feel free to change them.
4. Add Storage: A “Root” volume type should already be included. You *could* store monitoring data there (e.g. in a folder named `/influxdb-data`) but we will attach another volume and store the monitoring data there instead. Select “Add New Volume” and an EBS volume type.
5. Tag Instance: give your instance a memorable name.
6. Configure Security Group: choose your `bigchaindb` security group.

7. Review and launch your instance.

When it asks, choose an existing key pair: the one you created earlier (named `bigchaindb`).

Give your instance some time to launch and become able to accept SSH connections. You can see its current status in the AWS EC2 Console (in the “Instances” section). SSH into your instance using something like:

```
cd deploy-cluster-aws
ssh -i pem/bigchaindb.pem ubuntu@ec2-52-58-157-229.eu-central-1.compute.amazonaws.com
```

where `ec2-52-58-157-229.eu-central-1.compute.amazonaws.com` should be replaced by your new instance’s EC2 hostname. (To get that, go to the AWS EC2 Console, select Instances, click on your newly-launched instance, and copy its “Public DNS” name.)

Next, create a file system on the attached volume, make a directory named `/influxdb-data`, and set the attached volume’s mount point to be `/influxdb-data`. For detailed instructions on how to do that, see the AWS documentation for [Making an Amazon EBS Volume Available for Use](#).

Then install Docker and Docker Compose:

```
# in a Python 2.5-2.7 virtual environment where fabric, boto3, etc. are installed
fab --fabfile=fabfile-monitor.py --hosts=<EC2 hostname> install_docker
```

After Docker is installed, we can run the monitor with:

```
fab --fabfile=fabfile-monitor.py --hosts=<EC2 hostname> run_monitor
```

For more information about monitoring (e.g. how to view the Grafana dashboard in your web browser), see the Monitoring section of this documentation.

To configure a BigchainDB node to send monitoring data to the monitoring server, change the `statsd` host in the configuration of the BigchainDB node. The section on [Configuring a BigchainDB Node](#) explains how you can do that. (For example, you can change the `statsd` host in `$HOME/.bigchaindb`.)

Deploy a BigchainDB Cluster

Step 1

Suppose N is the number of nodes you want in your BigchainDB cluster. If you already have a set of N BigchainDB configuration files in the `deploy-cluster-aws/confiles` directory, then you can jump to the next step. To create such a set, you can do something like:

```
# in a Python 3 virtual environment where bigchaindb is installed
cd bigchaindb
cd deploy-cluster-aws
./make_confiles.sh confiles 3
```

That will create three (3) *default* BigchainDB configuration files in the `deploy-cluster-aws/confiles` directory (which will be created if it doesn’t already exist). The three files will be named `bcd_b_conf0`, `bcd_b_conf1`, and `bcd_b_conf2`.

You can look inside those files if you’re curious. For example, the default keyring is an empty list. Later, the deployment script automatically changes the keyring of each node to be a list of the public keys of all other nodes. Other changes are also made. That is, the configuration files generated in this step are *not* what will be sent to the deployed nodes; they’re just a starting point.

Step 2

Step 2 is to make an AWS deployment configuration file, if necessary. There's an example AWS configuration file named `example_deploy_conf.py`. It has many comments explaining each setting. The settings in that file are (or should be):

```
NUM_NODES=3
BRANCH="master"
WHAT_TO_DEPLOY="servers"
SSH_KEY_NAME="not-set-yet"
USE_KEYPAIRS_FILE=False
IMAGE_ID="ami-accff2b1"
INSTANCE_TYPE="m3.2xlarge"
USING_EBS=False
EBS_VOLUME_SIZE=30
EBS_OPTIMIZED=False
```

Make a copy of that file and call it whatever you like (e.g. `cp example_deploy_conf.py my_deploy_conf.py`). You can leave most of the settings at their default values, but you must change the value of `SSH_KEY_NAME` to the name of your private SSH key. You can do that with a text editor. Set `SSH_KEY_NAME` to the name you used for `<key-name>` when you generated an RSA key pair for SSH (in basic AWS setup).

If you want your nodes to have a predictable set of pre-generated keypairs, then you should 1) set `USE_KEYPAIRS_FILE=True` in the AWS deployment configuration file, and 2) provide a `keypairs.py` file containing enough keypairs for all of your nodes. You can generate a `keypairs.py` file using the `write_keypairs_file.py` script. For example:

```
# in a Python 3 virtual environment where bigchaindb is installed
cd bigchaindb
cd deploy-cluster-aws
python3 write_keypairs_file.py 100
```

The above command generates a `keypairs.py` file with 100 keypairs. You can generate more keypairs than you need, so you can use the same list over and over again, for different numbers of servers. The deployment scripts will only use the first `NUM_NODES` keypairs.

Step 3

Step 3 is to launch the nodes ("instances") on AWS, to install all the necessary software on them, configure the software, run the software, and more. Here's how you'd do that:

```
# in a Python 2.5-2.7 virtual environment where fabric, boto3, etc. are installed
cd bigchaindb
cd deploy-cluster-aws
./awsdeploy.sh my_deploy_conf.py
# Only if you want to set the replication factor to 3
fab set_replicas:3
# Only if you want to start BigchainDB on all the nodes:
fab start_bigchaindb
```

`awsdeploy.sh` is a Bash script which calls some Python and Fabric scripts. If you're curious what it does, [the source code](#) has many explanatory comments.

It should take a few minutes for the deployment to finish. If you run into problems, see the section on **Known Deployment Issues** below.

The EC2 Console has a section where you can see all the instances you have running on EC2. You can `ssh` into a running instance using a command like:

```
ssh -i pem/bigchaindb.pem ubuntu@ec2-52-29-197-211.eu-central-1.compute.amazonaws.com
```

except you'd replace the `ec2-52-29-197-211.eu-central-1.compute.amazonaws.com` with the public DNS name of the instance you want to `ssh` into. You can get that from the EC2 Console: just click on an instance and look in its details pane at the bottom of the screen. Some commands you might try:

```
ip addr show
sudo service rethinkdb status
bigchaindb --help
bigchaindb show-config
```

You can also check out the RethinkDB web interface at port 8080 on any of the instances; just go to your web browser and visit a web address like `http://ec2-52-29-197-211.eu-central-1.compute.amazonaws.com:8080/`.

Server Monitoring with New Relic

New Relic is a business that provides several monitoring services. One of those services, called Server Monitoring, can be used to monitor things like CPU usage and Network I/O on BigchainDB instances. To do that:

1. Sign up for a New Relic account
2. Get your New Relic license key
3. Put that key in an environment variable named `NEWRELIC_KEY`. For example, you might add a line like the following to your `~/.bashrc` file (if you use Bash): `export NEWRELIC_KEY=<insert your key here>`
4. Once you've deployed a BigchainDB cluster on AWS as above, you can install a New Relic system monitor (agent) on all the instances using:

```
# in a Python 2.5-2.7 virtual environment where fabric, boto3, etc. are installed
fab install_newrelic
```

Once the New Relic system monitor (agent) is installed on the instances, it will start sending server stats to New Relic on a regular basis. It may take a few minutes for data to show up in your New Relic dashboard (under New Relic Servers).

Shutting Down a Cluster

There are fees associated with running instances on EC2, so if you're not using them, you should terminate them. You can do that using the AWS EC2 Console.

The same is true of your allocated elastic IP addresses. There's a small fee to keep them allocated if they're not associated with a running instance. You can release them using the AWS EC2 Console, or by using a handy little script named `release_eips.py`. For example:

```
$ python release_eips.py
You have 2 allocated elastic IPs which are not associated with instances
0: Releasing 52.58.110.110
(It has Domain = vpc.)
1: Releasing 52.58.107.211
(It has Domain = vpc.)
```

Known Deployment Issues

NetworkError

If you tested with a high sequence it might be possible that you run into an error message like this:

```
NetworkError: Host key for ec2-xx-xx-xx-xx.eu-central-1.compute.amazonaws.com
did not match pre-existing key! Server's key was changed recently, or possible
man-in-the-middle attack.
```

If so, just clean up your `known_hosts` file and start again. For example, you might copy your current `known_hosts` file to `old_known_hosts` like so:

```
mv ~/.ssh/known_hosts ~/.ssh/old_known_hosts
```

Then terminate your instances and try deploying again with a different tag.

Failure of `sudo apt-get update`

The first thing that's done on all the instances, once they're running, is basically `sudo apt-get update`. Sometimes that fails. If so, just terminate your instances and try deploying again with a different tag. (These problems seem to be time-bounded, so maybe wait a couple of hours before retrying.)

Failure when Installing Base Software

If you get an error with installing the base software on the instances, then just terminate your instances and try deploying again with a different tag.

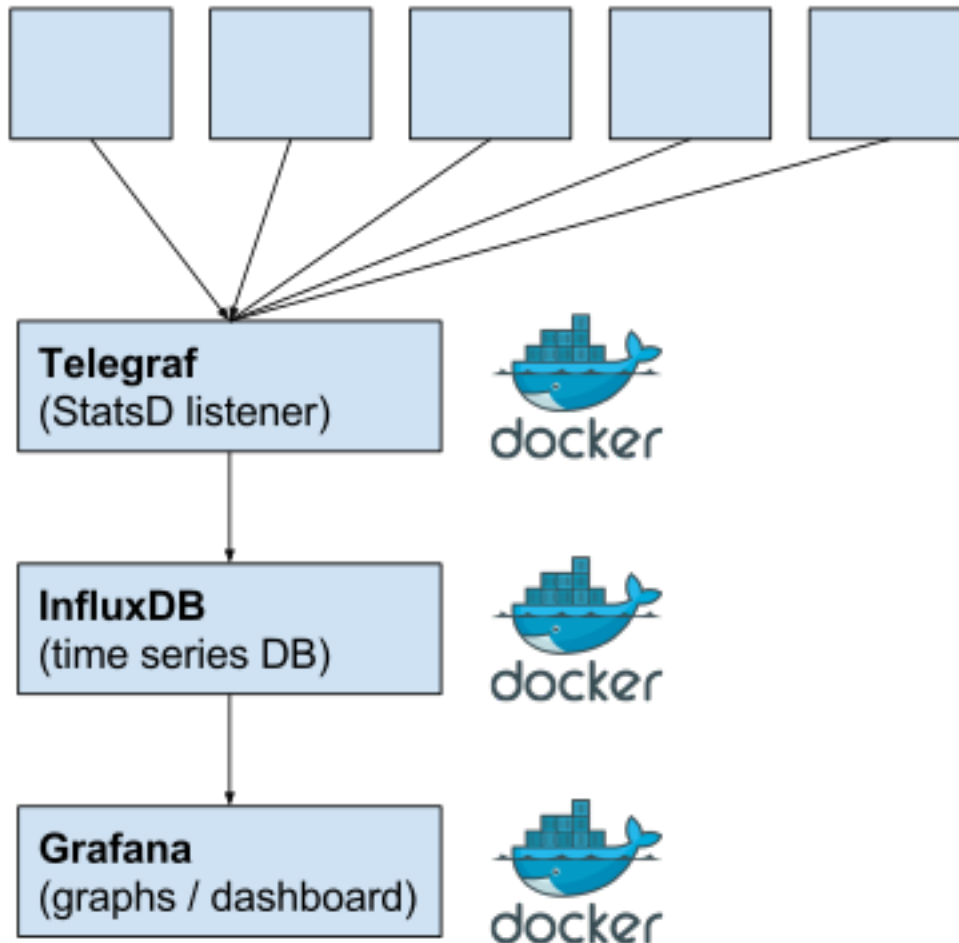
1.8.4 Cluster Monitoring

BigchainDB uses [StatsD](#) for cluster monitoring. We require some additional infrastructure to take full advantage of its functionality:

- an agent to listen for metrics: [Telegraf](#),
- a time-series database: [InfluxDB](#), and
- a frontend to display analytics: [Grafana](#).

We put each of those inside its own Docker container. The whole system is illustrated below.

Servers (nodes) running BigchainDB Server
(includes the statsd Python package to send metrics using StatsD protocol over UDP)



For ease of use, we've created a Docker *Compose file* (named `docker-compose-monitor.yml`) to define the monitoring system setup. To use it, just go to the top `bigchaindb` directory and run:

```
$ docker-compose -f docker-compose-monitor.yml build
$ docker-compose -f docker-compose-monitor.yml up
```

It is also possible to mount a host directory as a data volume for InfluxDB by setting the `INFLUXDB_DATA` environment variable:

```
$ INFLUXDB_DATA=/data docker-compose -f docker-compose-monitor.yml up
```

You can view the Grafana dashboard in your web browser at:

http://localhost:3000/dashboard/script/bigchaindb_dashboard.js

(You may want to replace `localhost` with another hostname in that URL, e.g. the hostname of a remote monitoring server.)

The login and password are `admin` by default. If BigchainDB is running and processing transactions, you should see analytics—if not, start BigchainDB and load some test transactions:

```
$ bigchaindb load
```

then refresh the page after a few seconds.

If you're not interested in monitoring, don't worry: BigchainDB will function just fine without any monitoring setup.

Feel free to modify the [custom Grafana dashboard](#) to your liking!

1.8.5 Documentation to Come

- Adding a node (including resharding etc.)
- Removing a node
- Upgrading BigchainDB or components
- Logging
- Node monitoring & crash recovery
- Node Security
 - (Private) key management
 - RethinkDB security
- Cluster monitoring
- Internal watchdogs
- External watchdogs

1.9 Topic Guides

Topic guides give background and explain concepts at a high level.

1.9.1 How BigchainDB is Decentralized

Decentralization means that no one owns or controls everything, and there is no single point of failure.

Ideally, each node in a BigchainDB cluster is owned and controlled by a different person or organization. Even if the cluster lives within one organization, it's still preferable to have each node controlled by a different person or subdivision.

We use the phrase “BigchainDB federation” (or just “federation”) to refer to the set of people and/or organizations who run the nodes of a BigchainDB cluster. A federation requires some form of governance to make decisions such as membership and policies. The exact details of the governance process are determined by each federation, but it can be very decentralized (e.g. purely vote-based, where each node gets a vote, and there are no special leadership roles).

The actual data is decentralized in that it doesn't all get stored in one place. Each federation node stores the primary of one shard and replicas of some other shards. (A shard is a subset of the total set of documents.) Sharding and replication are handled by RethinkDB.

A federation can increase its decentralization (and its resilience) by increasing its jurisdictional diversity, geographic diversity, and other kinds of diversity. This idea is expanded upon in the section on node diversity.

There's no node that has a long-term special position in the federation. All nodes run the same software and perform the same duties.

RethinkDB has an "admin" user which can't be deleted and which can make big changes to the database, such as dropping a table. Right now, that's a big security vulnerability, but we have plans to mitigate it by:

1. Locking down the admin user as much as possible.
2. Having all nodes inspect RethinkDB admin-type requests before acting on them. Requests can be checked against an evolving whitelist of allowed actions (voted on by federation nodes).

It's worth noting that the RethinkDB admin user can't transfer assets, even today. The only way to create a valid transfer transaction is to fulfill the current (crypto) conditions on the asset, and the admin user can't do that because the admin user doesn't have the necessary private keys (or preimages, in the case of hashlock conditions). They're not stored in the database.

1.9.2 Kinds of Node Diversity

Steps should be taken to make it difficult for any one actor or event to control or damage "enough" of the nodes. ("Enough" is usually a quorum.) There are many kinds of diversity to consider, listed below. It may be quite difficult to have high diversity of all kinds.

1. **Jurisdictional diversity.** The nodes should be controlled by entities within multiple legal jurisdictions, so that it becomes difficult to use legal means to compel enough of them to do something.
2. **Geographic diversity.** The servers should be physically located at multiple geographic locations, so that it becomes difficult for a natural disaster (such as a flood or earthquake) to damage enough of them to cause problems.
3. **Hosting diversity.** The servers should be hosted by multiple hosting providers (e.g. Amazon Web Services, Microsoft Azure, Digital Ocean, Rackspace), so that it becomes difficult for one hosting provider to influence enough of the nodes.
4. **Operating system diversity.** The servers should use a variety of operating systems, so that a security bug in one OS can't be used to exploit enough of the nodes.
5. **Diversity in general.** In general, membership diversity (of all kinds) confers many advantages on a federation. For example, it provides the federation with a source of various ideas for addressing challenges.

Note: If all the nodes are running the same code, i.e. the same implementation of BigchainDB, then a bug in that code could be used to compromise all of the nodes. Ideally, there would be several different, well-maintained implementations of BigchainDB Server (e.g. one in Python, one in Go, etc.), so that a federation could also have a diversity of server implementations.

1.9.3 How BigchainDB is Immutable / Tamper-Resistant

The blockchain community often describes blockchains as "immutable." If we interpret that word literally, it means that blockchain data is unchangeable or permanent, which is absurd. The data *can* be changed. For example, a plague might drive humanity extinct; the data would then get corrupted over time due to water damage, thermal noise, and the general increase of entropy. In the case of Bitcoin, nothing so drastic is required: a 51% attack will suffice.

It's true that blockchain data is more difficult to change than usual: it's more tamper-resistant than a typical file system or database. Therefore, in the context of blockchains, we interpret the word "immutable" to mean tamper-resistant. (Linguists would say that the word "immutable" is a *term of art* in the blockchain community.)

BigchainDB achieves strong tamper-resistance in the following ways:

1. **Replication.** All data is sharded and shards are replicated in several (different) places. The replication factor can be set by the federation. The higher the replication factor, the more difficult it becomes to change or delete all replicas.
2. **Internal watchdogs.** All nodes monitor all changes and if some unallowed change happens, then appropriate action is taken. For example, if a valid block is deleted, then it is put back.
3. **External watchdogs.** Federations may opt to have trusted third-parties to monitor and audit their data, looking for irregularities. For federations with publicly-readable data, the public can act as an auditor.
4. **Cryptographic signatures** are used throughout BigchainDB as a way to check if messages (transactions, blocks and votes) have been tampered with enroute, and as a way to verify who signed the messages. Each block is signed by the node that created it. Each vote is signed by the node that cast it. A creation transaction is signed by the node that created it, although there are plans to improve that by adding signatures from the sending client and multiple nodes; see [Issue #347](#). Transfer transactions can contain multiple fulfillments (one per asset transferred). Each fulfillment will typically contain one or more signatures from the owners (i.e. the owners before the transfer). Hashlock fulfillments are an exception; there's an open issue ([#339](#)) to address that.
5. **Full or partial backups** of the database may be recorded from time to time, possibly on magnetic tape storage, other blockchains, printouts, etc.
6. **Strong security.** Node owners can adopt and enforce strong security policies.
7. **Node diversity.** Diversity makes it so that no one thing (e.g. natural disaster or operating system bug) can compromise enough of the nodes. See the section on the kinds of node diversity.

Some of these things come “for free” as part of the BigchainDB software, and others require some extra effort from the federation and node owners.

1.9.4 BigchainDB and Byzantine Fault Tolerance

We have Byzantine fault tolerance (BFT) in our roadmap, as a switch that people can turn on. We anticipate that turning it on will cause a severe dropoff in performance (to gain some extra security). See [Issue #293](#).

Among the big, industry-used distributed databases in production today (e.g. DynamoDB, Bigtable, MongoDB, Cassandra, Elasticsearch), none of them are BFT. Indeed, almost all wide-area distributed systems in production are not BFT, including military, banking, healthcare, and other security-sensitive systems.

There are many more practical things that nodes can do to increase security (e.g. firewalls, key management, access controls).

From a [recent essay by Ken Birman](#) (of Cornell):

Oh, and with respect to the BFT point: Jim [Gray] felt that real systems fail by crashing [54]. Others have since done studies reinforcing this view, or finding that even crash-failure solutions can sometimes defend against application corruption. One interesting study, reported during a SOSP WIPS session by Ben Reed (one of the co-developers of Zookeeper), found that at Yahoo, Zookeeper itself had never experienced Byzantine faults in a one-year period that they studied closely.

[54] Jim Gray. Why Do Computers Stop and What Can Be Done About It? SOSP, 1985.

Ben Reed never published those results, but Birman wrote more about them in the book *Guide to Reliable Distributed Systems: Building High-Assurance Applications*. From page 358 of that book:

But the cloud community, led by Ben Reed and Flavio Junqueira at Yahoo, sees things differently (these are the two inventor's [sic] of Yahoo's ZooKeeper service). **They have described informal studies of how applications and machines at Yahoo failed, concluding that the frequency of Byzantine failures was extremely small relative to the frequency of crash failures** [emphasis added]. Sometimes they did see data corruption, but then they often saw it occur in a correlated way that impacted many replicas all at once. And very often they saw failures occur in the client layer, then propagate into the service.

BFT techniques tend to be used only within a service, not in the client layer that talks to that service, hence offer no protection against malfunctioning clients. **All of this, Reed and Junqueira conclude, lead to the realization that BFT just does not match the real needs of a cloud computing company like Yahoo, even if the data being managed by a service really is of very high importance** [emphasis added]. Unfortunately, they have not published this study; it was reported at an “outrageous opinions” session at the ACM Symposium on Operating Systems Principles, in 2009.

The practical use of the Byzantine protocol raises another concern: The timing assumptions built into the model [i.e. synchronous or partially-synchronous nodes] are not realizable in most computing environments...

1.9.5 How BigchainDB is Good for Asset Registrations & Transfers

BigchainDB can store data of any kind (within reason), but it’s designed to be particularly good for storing asset registrations and transfers:

- The fundamental thing that one submits to a BigchainDB federation to be checked and stored (if valid) is a *transaction*, and there are two kinds: creation transactions and transfer transactions.
- A creation transaction can be used to register any kind of indivisible asset, along with arbitrary metadata.
- An asset can have zero, one, or several owners.
- The owners of an asset can specify (crypto-)conditions which must be satisfied by anyone wishing transfer the asset to new owners. For example, a condition might be that at least 3 of the 5 current owners must cryptographically sign a transfer transaction.
- BigchainDB verifies that the conditions have been satisfied as part of checking the validity of transfer transactions. (Moreover, anyone can check that they were satisfied.)
- BigchainDB prevents double-spending of an asset.
- Validated transactions are strongly tamper-resistant; see the section about immutability / tamper-resistance.

You can read more about the details of BigchainDB transactions in the section on Transaction, Block and Vote Models (data models).

BigchainDB Integration with Other Blockchains

BigchainDB works with the [Interledger protocol](#), enabling the transfer of assets between BigchainDB and other blockchains, ledgers, and payment systems.

We’re actively exploring ways that BigchainDB can be used with other blockchains and platforms.

1.9.6 BigchainDB and Smart Contracts

BigchainDB isn’t intended for running smart contracts. That said, it can do many of the things that smart contracts are used to do. For example, the owners of an asset can impose conditions that must be fulfilled by anyone wishing to transfer the asset to new owners; see the section on conditions. BigchainDB also supports a form of escrow.

1.9.7 The Transaction, Block and Vote Models

BigchainDB stores all its records in JSON documents.

The three main kinds of records are transactions, blocks and votes. *Transactions* are used to register, issue, create or transfer things (e.g. assets). Multiple transactions are combined with some other metadata to form *blocks*. Nodes vote on blocks. This section is a reference on the details of transactions, blocks and votes.

Below we often refer to cryptographic hashes, keys and signatures. The details of those are covered in the section on cryptography.

Some Words of Caution

BigchainDB is still in the early stages of development. The data models described below may change substantially before BigchainDB reaches a production-ready state (i.e. version 1.0 and higher).

Transaction Concepts

Transactions are the most basic kind of record stored by BigchainDB. There are two kinds: creation transactions and transfer transactions.

A *creation transaction* can be used to register, issue, create or otherwise initiate the history of a single thing (or asset) in BigchainDB. For example, one might register an identity or a creative work. The things are often called “assets” but they might not be literal assets. A creation transaction also establishes the initial owner or owners of the asset. Only a federation node can create a valid creation transaction (but it’s usually made based on a message from a client).

Currently, BigchainDB only supports indivisible assets. You can’t split an asset apart into multiple assets, nor can you combine several assets together into one. [Issue #129](#) is an enhancement proposal to support divisible assets.

A *transfer transaction* can transfer one or more assets to new owners.

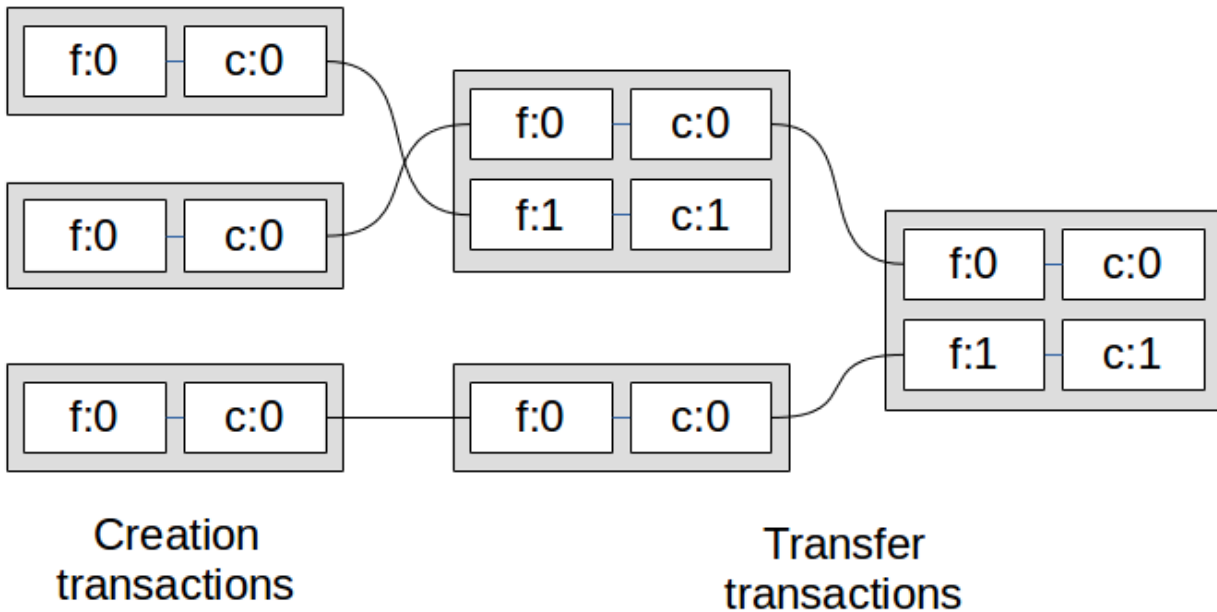
BigchainDB works with the [Interledger Protocol \(ILP\)](#), a protocol for transferring assets between different ledgers, blockchains or payment systems.

The owner(s) of an asset can specify conditions (ILP crypto-conditions) which others must fulfill (satisfy) in order to become the new owner(s) of the asset. For example, a crypto-condition might require a signature from the owner, or from m-of-n owners (a threshold condition, e.g. 3-of-4).

When someone creates a transfer transaction with the goal of changing an asset’s owners, they must fulfill the asset’s current crypto-conditions (i.e. in a fulfillment), and they must provide new conditions (including the list of new owners).

Every create transaction contains exactly one fulfillment-condition pair. A transfer transaction can contain multiple fulfillment-condition pairs: one per asset transferred. Every fulfillment in a transfer transaction (input) must correspond to a condition (output) in a previous transaction. The diagram below illustrates some of these concepts: transactions are represented by light grey boxes, fulfillments have a label like $f : 0$, and conditions have a label like $c : 0$.

Tracking the stories of three assets



To determine the current owner(s) of an asset, find the most recent valid transaction involving it, and then look at the list of owners in the *conditions* (not the fulfillments).

Transaction Validation

When a node is asked to check the validity of a transaction, it must do several things; the main things are:

- schema validation,
- double-spending checks (for transfer transactions),
- hash validation (i.e. is the calculated transaction hash equal to its id?), and
- validation of all fulfillments, including validation of cryptographic signatures if they're among the conditions.

The full details of transaction validation can be found in the code for `validate_transaction()` in the `BaseConsensusRules` class of `consensus.py` (unless other validation rules are being used by a federation, in which case those should be consulted instead).

Mutable and Immutable Assets

Assets can be mutable (changeable) or immutable. To change a mutable asset, you must create a valid transfer transaction with a payload specifying how it changed (or will change). The data structure (schema) of the change depends on the asset class. If you're inventing a new asset class, you can make up how to describe changes. For a mutable asset in an existing asset class, you should find out how changes are specified for that asset class. That's not something determined by BigchainDB.

The Transaction Model

```
{
  "id": "<hash of transaction, excluding signatures (see explanation)>",
  "transaction": {
    "version": "<version number of the transaction model>",
    "fulfillments": ["<list of fulfillments>"],
    "conditions": ["<list of conditions>"],
    "operation": "<string>",
    "timestamp": "<timestamp from client>",
    "data": {
      "uuid": "<uuid4>",
      "payload": "<any JSON document>"
    }
  }
}
```

Here's some explanation of the contents of a transaction:

- **id**: The hash of everything inside the serialized transaction body (i.e. fulfillments, conditions, operation, timestamp and data; see below), with one wrinkle: for each fulfillment in fulfillments, fulfillment is set to null. The id is also the database primary key.
- **transaction**:
 - **version**: Version number of the transaction model, so that software can support different transaction models.
 - **fulfillments**: List of fulfillments. Each *fulfillment* contains a pointer to an unspent asset and a *crypto fulfillment* that satisfies a spending condition set on the unspent asset. A *fulfillment* is usually a signature proving the ownership of the asset. See [Conditions and Fulfillments](#) below.
 - **conditions**: List of conditions. Each *condition* is a *crypto-condition* that needs to be fulfilled by a transfer transaction in order to transfer ownership to new owners. See [Conditions and Fulfillments](#) below.
 - **operation**: String representation of the operation being performed (currently either “CREATE” or “TRANSFER”). It determines how the transaction should be validated.
 - **timestamp**: The Unix time when the transaction was created. It's provided by the client. See the section on timestamps.
 - **data**:
 - * **uuid**: UUID version 4 (random) converted to a string of hex digits in standard form.
 - * **payload**: Can be any JSON document. It may be empty in the case of a transfer transaction.

Later, when we get to the models for the block and the vote, we'll see that both include a signature (from the node which created it). You may wonder why transactions don't have signatures... The answer is that they do! They're just hidden inside the *fulfillment* string of each fulfillment. A creation transaction is signed by the node that created it. A transfer transaction is signed by whoever currently controls or owns it.

What gets signed? For each fulfillment in the transaction, the “fulfillment message” that gets signed includes the operation, timestamp, data, version, id, corresponding condition, and the fulfillment itself, except with its fulfillment string set to null. The computed signature goes into creating the fulfillment string of the fulfillment.

One other note: Currently, transactions contain only the public keys of asset-owners (i.e. who own an asset or who owned an asset in the past), inside the conditions and fulfillments. A transaction does *not* contain the public key of the client (computer) which generated and sent it to a BigchainDB node. In fact, there's no need for a client to *have* a public/private keypair. In the future, each client may also have a keypair, and it may have to sign each sent transaction (using its private key); see [Issue #347 on GitHub](#). In practice, a person might think of their keypair as being both their

“ownership-keypair” and their “client-keypair,” but there is a difference, just like there’s a difference between Joe and Joe’s computer.

Conditions and Fulfillments

To create a transaction that transfers an asset to new owners, one must fulfill the asset’s current conditions (crypto-conditions). The most basic kinds of conditions are:

- **A hashlock condition:** One can fulfill a hashlock condition by providing the correct “preimage” (similar to a password or secret phrase)
- **A simple signature condition:** One can fulfill a simple signature condition by a providing a valid cryptographic signature (i.e. corresponding to the public key of an owner, usually)
- **A timeout condition:** Anyone can fulfill a timeout condition before the condition’s expiry time. After the expiry time, nobody can fulfill the condition. Another way to say this is that a timeout condition’s fulfillment is valid (TRUE) before the expiry time and invalid (FALSE) after the expiry time. Note: at the time of writing, timeout conditions are BigchainDB-specific (i.e. not part of the Interledger specs).

A more complex condition can be composed by using n of the above conditions as inputs to an m -of- n threshold condition (a logic gate which outputs TRUE iff m or more inputs are TRUE). If there are n inputs to a threshold condition:

- 1-of- n is the same as a logical OR of all the inputs
- n -of- n is the same as a logical AND of all the inputs

For example, one could create a condition requiring that m (of n) owners provide signatures before their asset can be transferred to new owners.

One can also put different weights on the inputs to threshold condition, along with a threshold that the weighted-sum-of-inputs must pass for the output to be TRUE. Weights could be used, for example, to express the number of shares that someone owns in an asset.

The (single) output of a threshold condition can be used as one of the inputs of other threshold conditions. This means that one can combine threshold conditions to build complex logical expressions, e.g. $(x \text{ OR } y) \text{ AND } (u \text{ OR } v)$.

Aside: In BigchainDB, the output of an m -of- n threshold condition can be inverted on the way out, so an output that would have been TRUE would get changed to FALSE (and vice versa). This enables the creation of NOT, NOR and NAND gates. At the time of writing, this “inverted threshold condition” is BigchainDB-specific (i.e. not part of the Interledger specs). It should only be used in combination with a timeout condition.

When one creates a condition, one can calculate its fulfillment length (e.g. 96). The more complex the condition, the larger its fulfillment length will be. A BigchainDB federation can put an upper limit on the allowed fulfillment length, as a way of capping the complexity of conditions (and the computing time required to validate them).

If someone tries to make a condition where the output of a threshold condition feeds into the input of another “earlier” threshold condition (i.e. in a closed logical circuit), then their computer will take forever to calculate the (infinite) “condition URI”, at least in theory. In practice, their computer will run out of memory or their client software will timeout after a while.

Aside: In what follows, the list of `owners_after` (in a condition) is always who owned the asset at the time the transaction completed, but before the next transaction started. The list of `owners_before` (in a fulfillment) is always equal to the list of `owners_after` in that asset’s previous transaction.

Conditions

One New Owner If there is only one *new owner*, the condition will be a simple signature condition (i.e. only one signature is required).

```
{
  "cid": "<condition index>",
  "condition": {
    "details": {
      "bitmask": "<base16 int>",
      "public_key": "<new owner public key>",
      "signature": null,
      "type": "fulfillment",
      "type_id": "<base16 int>"
    },
    "uri": "<string>"
  },
  "owners_after": ["<new owner public key>"]
}
```

- **Condition header:**

- cid: Condition index so that we can reference this output as an input to another transaction. It also matches the input fid, making this the condition to fulfill in order to spend the asset used as input with fid.
- owners_after: A list containing one item: the public key of the new owner.

- **Condition body:**

- bitmask: A set of bits representing the features required by the condition type.
- public_key: The new owner's public key.
- type_id: The fulfillment type ID; see the [ILP spec](#).
- uri: Binary representation of the condition using only URL-safe characters.

Multiple New Owners If there are multiple *new owners*, they can create a ThresholdCondition requiring a signature from each of them in order to spend the asset. For example:

```
{
  "cid": "<condition index>",
  "condition": {
    "details": {
      "bitmask": 41,
      "subfulfillments": [
        {
          "bitmask": 32,
          "public_key": "<new owner 1 public key>",
          "signature": null,
          "type": "fulfillment",
          "type_id": 4,
          "weight": 1
        },
        {
          "bitmask": 32,
          "public_key": "<new owner 2 public key>",
          "signature": null,
          "type": "fulfillment",
          "type_id": 4,
          "weight": 1
        }
      ]
    },
    "threshold": 2,
  }
}
```

```

    "type": "fulfillment",
    "type_id": 2
  },
  "uri": "cc:2:29:ytNK3X6-bZsbF-nCGDTuopUIMi1HCyCkyPewm6oLI3o:206",
  "owners_after": [
    "owner 1 public key>",
    "owner 2 public key>"
  ]
}

```

- `subfulfillments`: a list of fulfillments
 - `weight`: integer weight for each subfulfillment's contribution to the threshold
- `threshold`: threshold to reach for the subfulfillments to reach a valid fulfillment

The weights and threshold could be adjusted. For example, if the threshold was changed to 1 above, then only one of the new owners would have to provide a signature to spend the asset.

Fulfillments

One Current Owner If there is only one *current owner*, the fulfillment will be a simple signature fulfillment (i.e. containing just one signature).

```

{
  "owners_before": ["<public key of the owner before the transaction happened>"],
  "fid": 0,
  "fulfillment": "cf:4:RxFzIE679tFBk8zwEgizhmTuciAylvTUwy6EL6ehddHFJOHk5F4IjwQ1xLu2oQK9iyRCZJdfWAe",
  "input": {
    "cid": 0,
    "txid": "11b3e7d893cc5fdcf1a1706809c7def290a3b10b0bef6525d10b024649c42d3"
  }
}

```

- `fid`: Fulfillment index. It matches a `cid` in the conditions with a new *crypto-condition* that the new owner needs to fulfill to spend this asset.
- `owners_before`: A list of public keys of the owners before the transaction; in this case it has just one public key.
- `fulfillment`: A crypto-conditions URI that encodes the cryptographic fulfillments like signatures and others, see [crypto-conditions](#).
- `input`: Pointer to the asset and condition of a previous transaction
 - `cid`: Condition index
 - `txid`: Transaction id

The Block Model

```

{
  "id": "<hash of block>",
  "block": {
    "timestamp": "<block-creation timestamp>",
    "transactions": ["<list of transactions>"],
    "node_pubkey": "<public key of the node creating the block>",
    "voters": ["<list of federation nodes public keys>"]
  }
}

```

```
} ,
  "signature": "<signature of block>",
}
```

- **id**: The hash of the serialized block (i.e. the timestamp, transactions, node_pubkey, and voters). This is also a database primary key; that's how we ensure that all blocks are unique.
- **block**:
 - **timestamp**: The Unix time when the block was created. It's provided by the node that created the block. See the section on timestamps.
 - **transactions**: A list of the transactions included in the block.
 - **node_pubkey**: The public key of the node that create the block.
 - **voters**: A list of public keys of federation nodes. Since the size of the federation may change over time, this will tell us how many nodes existed in the federation when the block was created, so that at a later point in time we can check that the block received the correct number of votes.
- **signature**: Signature of the block by the node that created the block. (To create the signature, the node serializes the block contents and signs that with its private key.)

The Vote Model

Each node must generate a vote for each block, to be appended the `votes` table. A vote has the following structure:

```
{
  "id": "<RethinkDB-generated ID for the vote>",
  "node_pubkey": "<the public key of the voting node>",
  "vote": {
    "voting_for_block": "<id of the block the node is voting for>",
    "previous_block": "<id of the block previous to this one>",
    "is_block_valid": "<true|false>",
    "invalid_reason": "<None|DOUBLE_SPEND|TRANSACTIONS_HASH_MISMATCH|NODES_PUBKEYS_MISMATCH>",
    "timestamp": "<Unix time when the vote was generated, provided by the voting node>"
  },
  "signature": "<signature of vote>",
}
```

Note: The `invalid_reason` was not being used as of v0.1.3 and may be dropped in a future version of BigchainDB.

1.9.8 Timestamps in BigchainDB

Each transaction, block and vote has an associated timestamp. Interpreting those timestamps is tricky, hence the need for this section.

Timestamp Sources & Accuracy

A transaction's timestamp is provided by the client which created and submitted the transaction to a BigchainDB node. A block's timestamp is provided by the BigchainDB node which created the block. A vote's timestamp is provided by the BigchainDB node which created the vote.

When a BigchainDB client or node needs a timestamp, it calls a BigchainDB utility function named `timestamp()`. There's a detailed explanation of how that function works below, but the short version is that it gets the **Unix time** from its system clock, rounded to the nearest second.

We can't say anything about the accuracy of the system clock on clients. Timestamps from clients are still potentially useful, however, in a statistical sense. We say more about that below.

We advise BigchainDB nodes to run special software (an "NTP daemon") to keep their system clock in sync with standard time servers. (NTP stands for [Network Time Protocol](#).)

Converting Timestamps to UTC

To convert a BigchainDB timestamp (a Unix time) to UTC, you need to know how the node providing the timestamp was set up. That's because different setups will report a different "Unix time" value around leap seconds! There's a [nice Red Hat Developer Blog post about the various setup options](#). If you want more details, see [David Mills' pages about leap seconds, NTP, etc.](#) (David Mills designed NTP.)

We advise BigchainDB nodes to run an NTP daemon with particular settings, so that their timestamps are consistent.

If a timestamp comes from a node that's set up as we advise, it can be converted to UTC as follows:

1. Use a standard "Unix time to UTC" converter to get a UTC timestamp.
2. Is the UTC timestamp a leap second, or the second before/after a leap second? There's [a list of all the leap seconds on Wikipedia](#).
3. If no, then you are done.
4. If yes, then it might not be possible to convert it to a single UTC timestamp. Even if it can't be converted to a single UTC timestamp, it *can* be converted to a list of two possible UTC timestamps. Showing how to do that is beyond the scope of this documentation. In all likelihood, you will never have to worry about leap seconds because they are very rare. (There were only 26 between 1972 and the end of 2015.)

Calculating Elapsed Time Between Two Timestamps

There's another gotcha with (Unix time) timestamps: you can't calculate the real-world elapsed time between two timestamps (correctly) by subtracting the smaller timestamp from the larger one. The result won't include any of the leap seconds that occurred between the two timestamps. You could look up how many leap seconds happened between the two timestamps and add that to the result. There are many library functions for working with timestamps; those are beyond the scope of this documentation.

Avoid Doing Transactions Around Leap Seconds

Because of the ambiguity and confusion that arises with Unix time around leap seconds, we advise users to avoid creating transactions around leap seconds.

Interpreting Sets of Timestamps

You can look at many timestamps to get a statistical sense of when something happened. For example, a transaction in a decided-valid block has many associated timestamps:

- its own timestamp
- the timestamps of the other transactions in the block; there could be as many as 999 of those
- the timestamp of the block
- the timestamps of all the votes on the block

Those timestamps come from many sources, so you can look at all of them to get some statistical sense of when the transaction “actually happened.” The timestamp of the block should always be after the timestamp of the transaction, and the timestamp of the votes should always be after the timestamp of the block.

How BigchainDB Uses Timestamps

BigchainDB *doesn't* use timestamps to determine the order of transactions or blocks. In particular, the order of blocks is determined by RethinkDB's changefeed on the bigchain table.

BigchainDB does use timestamps for some things. It uses them to determine if a transaction has been waiting in the backlog for too long (i.e. because the node assigned to it hasn't handled it yet). It also uses timestamps to determine the status of timeout conditions (used by escrow).

Including Trusted Timestamps

If you want to create a transaction payload with a trusted timestamp, you can.

One way to do that would be to send a payload to a trusted timestamping service. They will send back a timestamp, a signature, and their public key. They should also explain how you can verify the signature. You can then include the original payload, the timestamp, the signature, and the service's public key in your transaction. That way, anyone with the verification instructions can verify that the original payload was signed by the trusted timestamping service.

How the `timestamp()` Function Works

BigchainDB has a utility function named `timestamp()` which amounts to:

```
timestamp() = str(round(time.time()))
```

In other words, it calls the `time()` function in Python's `time` module, **rounds** that to the nearest integer, and converts the result to a string.

It rounds the output of `time.time()` to the nearest second because, according to [the Python documentation for `time.time\(\)`](#), “...not all systems provide time with a better precision than 1 second.”

How does `time.time()` work? If you look in the C source code, it calls `floattime()` and `floattime()` calls `clock_gettime()`, if it's available.

```
ret = clock_gettime(CLOCK_REALTIME, &tp);
```

With `CLOCK_REALTIME` as the first argument, it returns the “Unix time.” (“Unix time” is in quotes because its value around leap seconds depends on how the system is set up; see above.)

Why Not Use UTC, TAI or Some Other Time that Has Unambiguous Timestamps for Leap Seconds?

It would be nice to use UTC or TAI timestamps, but unfortunately there's no commonly-available, standard way to get always-accurate UTC or TAI timestamps from the operating system on typical computers today (i.e. accurate around leap seconds).

There *are* commonly-available, standard ways to get the “Unix time,” such as `clock_gettime()` function available in C. That's what we use (indirectly via Python). (“Unix time” is in quotes because its value around leap seconds depends on how the system is set up; see above.)

The Unix-time-based timestamps we use are only ambiguous circa leap seconds, and those are very rare. Even for those timestamps, the extra uncertainty is only one second, and that's not bad considering that we only report timestamps to a precision of one second in the first place. All other timestamps can be converted to UTC with no ambiguity.

1.10 Release Notes

You can find a list of all BigchainDB releases and release notes on GitHub at:

<https://github.com/bigchaindb/bigchaindb/releases>

The `CHANGELOG.md` file contains much the same information, but it also has notes about what to expect in the *next* release.

We also have a [roadmap document](#) in `ROADMAP.md`.

1.11 Appendices

1.11.1 How to Install OS-Level Dependencies

BigchainDB Server has some OS-level dependencies that must be installed.

On Ubuntu 14.04 and 16.04, we found that the following was enough:

```
sudo apt-get update
sudo apt-get install g++ python3-dev
```

On Fedora 23 and 24, we found that the following was enough:

```
sudo dnf update
sudo dnf install gcc-c++ redhat-rpm-config python3-devel
```

(If you're using a version of Fedora before version 22, you may have to use `yum` instead of `dnf`.)

1.11.2 How to Install the Latest pip and setuptools

You can check the version of `pip` you're using (in your current `virtualenv`) by doing:

```
pip -V
```

If it says that `pip` isn't installed, or it says `pip` is associated with a Python version less than 3.4, then you must install a `pip` version associated with Python 3.4+. In the following instructions, we call it `pip3` but you may be able to use `pip` if that refers to the same thing. See [the pip installation instructions](#).

On Ubuntu 14.04, we found that this works:

```
sudo apt-get install python3-pip
```

That should install a Python 3 version of `pip` named `pip3`. If that didn't work, then another way to get `pip3` is to do `sudo apt-get install python3-setuptools` followed by `sudo easy_install3 pip`.

You can upgrade `pip` (`pip3`) and `setuptools` to the latest versions using:

```
pip3 install --upgrade pip setuptools
```

1.11.3 Run BigchainDB with Docker

NOT for Production Use

For those who like using Docker and wish to experiment with BigchainDB in non-production environments, we currently maintain a Docker image and a `Dockerfile` that can be used to build an image for `bigchaindb`.

Pull and Run the Image from Docker Hub

Assuming you have Docker installed, you would proceed as follows.

In a terminal shell, pull the latest version of the BigchainDB Docker image using:

```
docker pull bigchaindb/bigchaindb
```

then do a one-time configuration step to create the config file; we will use the `-y` option to accept all the default values. The configuration file will be stored in a file on your host machine at `~/bigchaindb_docker/.bigchaindb`:

```
docker run --rm -v "$HOME/bigchaindb_docker:/data" -ti \
  bigchaindb/bigchaindb -y configure
Generating keypair
Configuration written to /data/.bigchaindb
Ready to go!
```

Let's analyze that command:

- `docker run` tells Docker to run some image
- `--rm` remove the container once we are done
- `-v "$HOME/bigchaindb_docker:/data"` map the host directory `$HOME/bigchaindb_docker` to the container directory `/data`; this allows us to have the data persisted on the host machine, you can read more in the [official Docker documentation](#)
- `-t` allocate a pseudo-TTY
- `-i` keep STDIN open even if not attached
- `'bigchaindb/bigchaindb'` the image to use
- `-y configure` execute the `configure` sub-command (of the `bigchaindb` command) inside the container, with the `-y` option to automatically use all the default config values

After configuring the system, you can run BigchainDB with the following command:

```
docker run -v "$HOME/bigchaindb_docker:/data" -d \
  --name bigchaindb \
  -p "58080:8080" -p "59984:9984" \
  bigchaindb/bigchaindb start
```

The command is slightly different from the previous one, the differences are:

- `-d` run the container in the background
- `--name bigchaindb` give a nice name to the container so it's easier to refer to it later
- `-p "58080:8080"` map the host port 58080 to the container port 8080 (the RethinkDB admin interface)
- `-p "59984:9984"` map the host port 59984 to the container port 9984 (the BigchainDB API server)
- `start` start the BigchainDB service

Another way to publish the ports exposed by the container is to use the `-P` (or `--publish-all`) option. This will publish all exposed ports to random ports. You can always run `docker ps` to check the random mapping.

You can also access the RethinkDB dashboard at <http://localhost:58080/>

If that doesn't work, then replace `localhost` with the IP or hostname of the machine running the Docker engine. If you are running `docker-machine` (e.g. on Mac OS X) this will be the IP of the Docker machine (`docker-machine ip machine_name`).

Load Testing with Docker

Now that we have BigchainDB running in the Docker container named `bigchaindb`, we can start another BigchainDB container to generate a load test for it.

First, make sure the container named `bigchaindb` is still running. You can check that using:

```
docker ps
```

You should see a container named `bigchaindb` in the list.

You can load test the BigchainDB running in that container by running the `bigchaindb load` command in a second container:

```
docker run --rm -v "$HOME/bigchaindb_docker:/data" -ti \
  --link bigchaindb \
  bigchaindb/bigchaindb load
```

Note the `--link` option to link to the first container (named `bigchaindb`).

Aside: The `bigchaindb load` command has several options (e.g. `-m`). You can read more about it in the documentation about the BigchainDB command line interface.

If you look at the RethinkDB dashboard (in your web browser), you should see the effects of the load test. You can also see some effects in the Docker logs using:

```
docker logs -f bigchaindb
```

Building Your Own Image

Assuming you have Docker installed, you would proceed as follows.

In a terminal shell:

```
git clone git@github.com:bigchaindb/bigchaindb.git
```

Build the Docker image:

```
docker build --tag local-bigchaindb .
```

Now you can use your own image to run BigchainDB containers.

1.11.4 JSON Serialization

We needed to clearly define how to serialize a JSON object to calculate the hash.

The serialization should produce the same byte output independently of the architecture running the software. If there are differences in the serialization, hash validations will fail although the transaction is correct.

For example, consider the following two methods of serializing `{ 'a' : 1 }`:

```
# Use a serializer provided by RethinkDB
a = r.expr({'a': 1}).to_json().run(b.connection)
u'{"a":1}'

# Use the serializer in Python's json module
b = json.dumps({'a': 1})
'{"a": 1}'
```

```
a == b
False
```

The results are not the same. We want a serialization and deserialization so that the following is always true:

```
deserialize(serialize(data)) == data
True
```

Since BigchainDB performs a lot of serialization we decided to use `python-rapidjson` which is a python wrapper for `rapidjson` a fast and fully RFC compliant JSON parser.

```
import rapidjson

rapidjson.dumps(data, skipkeys=False,
                 ensure_ascii=False,
                 sort_keys=True)
```

- `skipkeys`: With `skipkeys` `False` if the provided keys are not a string the serialization will fail. This way we enforce all keys to be strings
- `ensure_ascii`: The RFC recommends `utf-8` for maximum interoperability. By setting `ensure_ascii` to `False` we allow unicode characters and `python-rapidjson` forces the encoding to `utf-8`.
- `sort_keys`: Sorted output by keys.

Every time we need to perform some operation on the data like calculating the hash or signing/verifying the transaction, we need to use the previous criteria to serialize the data and then use the `byte` representation of the serialized data (if we treat the data as bytes we eliminate possible encoding errors e.g. unicode characters). For example:

```
# calculate the hash of a transaction
# the transaction is a dictionary
tx_serialized = bytes(serialize(tx))
tx_hash = hashlib.sha3_256(tx_serialized).hexdigest()

# signing a transaction
tx_serialized = bytes(serialize(tx))
signature = sk.sign(tx_serialized)

# verify signature
tx_serialized = bytes(serialize(tx))
vk.verify(signature, tx_serialized)
```

1.11.5 Cryptography

The section documents the cryptographic algorithms and Python implementations that we use.

Before hashing or computing the signature of a JSON document, we serialize it as described in the section on JSON serialization.

Hashes

We compute hashes using the SHA3-256 algorithm and `pysha3` as the Python implementation. We store the hex-encoded hash in the database. For example:

```
import hashlib
# monkey patch hashlib with sha3 functions
import sha3
```

```
data = "message"
tx_hash = hashlib.sha3_256(data).hexdigest()
```

Signature Algorithm and Keys

BigchainDB uses the [Ed25519](#) public-key signature system for generating its public/private key pairs (also called verifying/signing keys). Ed25519 is an instance of the [Edwards-curve Digital Signature Algorithm \(EdDSA\)](#). As of April 2016, EdDSA was in “Internet-Draft” status with the IETF but was already widely used.

BigchainDB uses the [ed25519](#) Python package, overloaded by the [cryptoconditions](#) library.

All keys are represented with the base58 encoding by default.

1.11.6 The Bigchain class

The Bigchain class is the top-level Python API for BigchainDB. If you want to create and initialize a BigchainDB database, you create a Bigchain instance (object). Then you can use its various methods to create transactions, write transactions (to the object/database), read transactions, etc.

```
class bigchaindb.Bigchain(host=None, port=None, dbname=None, public_key=None, private_key=None, keyring=[], consensus_plugin=None)
```

Bigchain API

Create, read, sign, write transactions to the database

```
__init__(host=None, port=None, dbname=None, public_key=None, private_key=None, keyring=[], consensus_plugin=None)
```

Initialize the Bigchain instance

A Bigchain instance has several configuration parameters (e.g. `host`). If a parameter value is passed as an argument to the Bigchain `__init__` method, then that is the value it will have. Otherwise, the parameter value will come from an environment variable. If that environment variable isn’t set, then the value will come from the local configuration file. And if that variable isn’t in the local configuration file, then the parameter will have its default value (defined in `bigchaindb.__init__`).

Parameters

- **host** (*str*) – hostname where RethinkDB is running.
- **port** (*int*) – port in which RethinkDB is running (usually 28015).
- **dbname** (*str*) – the name of the database to connect to (usually `bigchain`).
- **public_key** (*str*) – the base58 encoded public key for the ED25519 curve.
- **private_key** (*str*) – the base58 encoded private key for the ED25519 curve.
- **keyring** (*list[str]*) – list of base58 encoded public keys of the federation nodes.

```
create_transaction(*args, **kwargs)
```

Create a new transaction

Refer to the documentation of your consensus plugin.

Returns newly constructed transaction.

Return type `dict`

```
sign_transaction(transaction, *args, **kwargs)
```

Sign a transaction

Refer to the documentation of your consensus plugin.

Returns transaction with any signatures applied.

Return type `dict`

validate_fulfillments (*signed_transaction*, *args, **kwargs)

Validate the fulfillment(s) of a transaction.

Refer to the documentation of your consensus plugin.

Returns

True if the transaction's required fulfillments are present and correct, False otherwise.

Return type `bool`

write_transaction (*signed_transaction*, durability='soft')

Write the transaction to bigchain.

When first writing a transaction to the bigchain the transaction will be kept in a backlog until it has been validated by the nodes of the federation.

Parameters **signed_transaction** (*dict*) – transaction with the *signature* included.

Returns database response

Return type `dict`

get_transaction (*txid*, include_status=False)

Retrieve a transaction with *txid* from bigchain.

Queries the bigchain for a transaction, if it's in a valid or invalid block.

Parameters

- **txid** (*str*) – transaction id of the transaction to query
- **include_status** (*bool*) – also return the status of the transaction the return value is then a tuple: (tx, status)

Returns A dict with the transaction details if the transaction was found. Will add the transaction status to payload ('valid', 'undecided', or 'backlog'). If no transaction with that *txid* was found it returns *None*

get_status (*txid*)

Retrieve the status of a transaction with *txid* from bigchain.

Parameters **txid** (*str*) – transaction id of the transaction to query

Returns transaction status ('valid', 'undecided', or 'backlog'). If no transaction with that *txid* was found it returns *None*

Return type (string)

search_block_election_on_index (*value*, *index*)

Retrieve block election information given a secondary index and value

Parameters

- **value** – a value to search (e.g. transaction id string, payload hash string)
- **index** (*str*) – name of a secondary index, e.g. 'transaction_id'

Returns A list of blocks with with only election information

get_blocks_status_containing_tx (*txid*)

Retrieve block ids and statuses related to a transaction

Transactions may occur in multiple blocks, but no more than one valid block.

Parameters `txid` (*str*) – transaction id of the transaction to query

Returns A dict of blocks containing the transaction, e.g. `{block_id_1: 'valid', block_id_2: 'invalid' ...}`, or `None`

get_tx_by_payload_uuid (*payload_uuid*)

Retrieves transactions related to a digital asset.

When creating a transaction one of the optional arguments is the *payload*. The payload is a generic dict that contains information about the digital asset.

To make it easy to query the bigchain for that digital asset we create a UUID for the payload and store it with the transaction. This makes it easy for developers to keep track of their digital assets in bigchain.

Parameters `payload_uuid` (*str*) – the UUID for this particular payload.

Returns A list of transactions containing that payload. If no transaction exists with that payload it returns an empty list `[]`

get_spent (*tx_input*)

Check if a *txid* was already used as an input.

A transaction can be used as an input for another transaction. Bigchain needs to make sure that a given *txid* is only used once.

Parameters `tx_input` (*dict*) – Input of a transaction in the form `{'txid': 'transaction id', 'cid': 'condition id'}`

Returns The transaction that used the *txid* as an input if it exists else it returns *None*

get_owned_ids (*owner*)

Retrieve a list of *txids* that can be used as inputs.

Parameters `owner` (*str*) – base58 encoded public key.

Returns list of *txids* currently owned by *owner*

Return type `list`

validate_transaction (*transaction*)

Validate a transaction.

Parameters `transaction` (*dict*) – transaction to validate.

Returns The transaction if the transaction is valid else it raises an exception describing the reason why the transaction is invalid.

is_valid_transaction (*transaction*)

Check whether a transaction is valid or invalid.

Similar to *validate_transaction* but never raises an exception. It returns *False* if the transaction is invalid.

Parameters `transaction` (*dict*) – transaction to check.

Returns *transaction* if the transaction is valid, *False* otherwise

create_block (*validated_transactions*)

Creates a block given a list of *validated_transactions*.

Note that this method does not validate the transactions. Transactions should be validated before calling *create_block*.

Parameters `validated_transactions` (*list*) – list of validated transactions.

Returns created block.

Return type `dict`

validate_block (*block*)

Validate a block.

Parameters **block** (*dict*) – block to validate.

Returns The block if the block is valid else it raises an exception describing the reason why the block is invalid.

has_previous_vote (*block*)

Check for previous votes from this node

Parameters **block** (*dict*) – block to check.

Returns `True` if this block already has a valid vote from this node, `False` otherwise.

Return type `bool`

Raises `ImproperVoteError` – If there is already a vote, but the vote is invalid.

is_valid_block (*block*)

Check whether a block is valid or invalid.

Similar to `validate_block` but does not raise an exception if the block is invalid.

Parameters **block** (*dict*) – block to check.

Returns `True` if the block is valid, `False` otherwise.

Return type `bool`

write_block (*block*, *durability='soft'*)

Write a block to bigchain.

Parameters **block** (*dict*) – block to write to bigchain.

prepare_genesis_block ()

Prepare a genesis block.

create_genesis_block ()

Create the genesis block

Block created when bigchain is first initialized. This method is not atomic, there might be concurrency problems if multiple instances try to write the genesis block when the BigchainDB Federation is started, but it's a highly unlikely scenario.

vote (*block_id*, *previous_block_id*, *decision*, *invalid_reason=None*)

Cast your vote on the block given the `previous_block_hash` and the `decision` (valid/invalid) return the block to the updated in the database.

Parameters

- **block_id** (*str*) – The id of the block to vote.
- **previous_block_id** (*str*) – The id of the previous block.
- **decision** (*bool*) – Whether the block is valid or invalid.
- **invalid_reason** (*Optional[str]*) – Reason the block is invalid

write_vote (*vote*)

Write the vote to the database.

get_last_voted_block ()

Returns the last block that this node voted on.

get_unvoted_blocks ()

Return all the blocks that has not been voted by this node.

block_election_status (*block*)

Tally the votes on a block, and return the status: valid, invalid, or undecided.

1.11.7 Basic AWS Setup

Before you can deploy anything on AWS, you must do a few things.

Get an AWS Account

If you don't already have an AWS account, you can [sign up for one for free at aws.amazon.com](https://aws.amazon.com).

Install the AWS Command-Line Interface

To install the AWS Command-Line Interface (CLI), just do:

```
pip install awscli
```

Create an AWS Access Key

The next thing you'll need is an AWS access key. If you don't have one, you can create one using the [instructions in the AWS documentation](#). You should get an access key ID (e.g. AKIAIOSFODNN7EXAMPLE) and a secret access key (e.g. wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY).

You should also pick a default AWS region name (e.g. eu-central-1). That's where your cluster will run. The AWS documentation has [a list of them](#).

Once you've got your AWS access key, and you've picked a default AWS region name, go to a terminal session and enter:

```
aws configure
```

and answer the four questions. For example:

```
AWS Access Key ID [None]: AKIAIOSFODNN7EXAMPLE
AWS Secret Access Key [None]: wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
Default region name [None]: eu-central-1
Default output format [None]: [Press Enter]
```

This writes two files: `~/.aws/credentials` and `~/.aws/config`. AWS tools and packages look for those files.

Generate an RSA Key Pair for SSH

Eventually, you'll have one or more instances (virtual machines) running on AWS and you'll want to SSH to them. To do that, you need a public/private key pair. The public key will be sent to AWS, and you can tell AWS to put it in any instances you provision there. You'll keep the private key on your local workstation.

First you need to make up a key name. Some ideas:

- bcdb-troy-1
- bigchaindb-7
- bcdb-jupiter

If you already have key pairs on AWS (Amazon EC2), you have to pick a name that's not already being used. Below, replace every instance of `<key-name>` with your actual key name. To generate a public/private RSA key pair with that name:

```
ssh-keygen -t rsa -C "<key-name>" -f ~/.ssh/<key-name>
```

It will ask you for a passphrase. You can use whatever passphrase you like, but don't lose it. Two keys (files) will be created in `~/.ssh/`:

1. `~/.ssh/<key-name>.pub` is the public key
2. `~/.ssh/<key-name>` is the private key

To send the public key to AWS, use the AWS Command-Line Interface:

```
aws ec2 import-key-pair \
--key-name "<key-name>" \
--public-key-material file:///~/.ssh/<key-name>.pub
```

If you're curious why there's a `file://` in front of the path to the public key, see issue [aws/aws-cli#41](#) on GitHub.

If you want to verify that your key pair was imported by AWS, go to the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>, select the region you gave above when you did `aws configure` (e.g. eu-central-1), click on **Key Pairs** in the left sidebar, and check that `<key-name>` is listed.

1.11.8 BigchainDB Consensus Plugins

BigchainDB has a pluggable block/transaction validation architecture. The default consensus rules can be extended or replaced entirely.

Installing a plugin

Plugins can be installed via pip!

```
$ pip install bigchaindb-plugin-demo
```

Or using `setuptools`:

```
$ cd bigchaindb-plugin-demo/
$ python setup.py install # (or develop)
```

To activate your plugin, you can either set the `consensus_plugin` field in your config file (usually `~/.bigchaindb`) or by setting the `BIGCHAIN_CONSENSUS_PLUGIN` environment variable to the name of your plugin (see the section on *Packaging a plugin* for more about plugin names).

Plugin API

BigchainDB's current plugin API exposes five functions in an `AbstractConsensusRules` class:

```
validate_transaction(bigchain, transaction)
validate_block(bigchain, block)
create_transaction(*args, **kwargs)
sign_transaction(transaction, *args, **kwargs)
validate_fulfillments(transaction)
```

Together, these functions are sufficient for most customizations. For example:

- Replace the crypto-system with one your hardware can accelerate

- Re-implement an existing protocol
- Delegate validation to another application
- etc...

Extending BigchainDB behavior

A default installation of BigchainDB will use the rules in the `BaseConsensusRules` class. If you only need to modify this behavior slightly, you can inherit from that class and call `super()` in any methods you change, so long as the return values remain the same.

Here's a quick example of a plugin that adds nonsense rules:

```
from bigchaindb.consensus import BaseConsensusRules

class SillyConsensusRules(BaseConsensusRules):

    @staticmethod
    def validate_transaction(bigchain, transaction):
        transaction = super().validate_transaction(bigchain, transaction)
        # I only like transactions whose timestamps are even.
        if transaction['transaction']['timestamp'] % 2 != 0:
            raise StandardError("Odd... very odd indeed.")
        return transaction

    @staticmethod
    def validate_block(bigchain, block):
        block = super().validate_block(bigchain, block)
        # I don't trust Alice, I think she's shady.
        if block['block']['node_pubkey'] == '<ALICE_PUBKEY>':
            raise StandardError("Alice is shady, everybody ignore her blocks!")
        return block
```

Packaging a plugin

BigchainDB uses `setuptools`' `entry_points` to provide the plugin functionality. Any custom plugin needs to add this section to the `setup()` call in their `setup.py`:

```
entry_points={
    'bigchaindb.consensus': [
        'PLUGIN_NAME=package.module:ConsensusRulesClass'
    ]
},
```

1.11.9 Notes for Firewall Setup

This is a page of notes on the ports used by BigchainDB nodes and the traffic they should expect, to help with firewall setup (or security group setup on AWS). This page is *not* a firewall tutorial or step-by-step guide.

Port 22

Port 22 is the default SSH port (TCP) so you'll at least want to make it possible to SSH in from your remote machine(s).

Port 53

Port 53 is the default DNS port (UDP). It may be used, for example, by some package managers when look up the IP address associated with certain package sources.

Port 80

Port 80 is the default HTTP port (TCP). It's used by some package managers to get packages. It's *not* used by the RethinkDB web interface (see Port 8080 below) or the BigchainDB client-server HTTP API (Port 9984).

Port 123

Port 123 is the default NTP port (UDP). You should be running an NTP daemon on production BigchainDB nodes. NTP daemons must be able to send requests to external NTP servers and accept the responses.

Port 161

Port 161 is the default SNMP port (usually UDP, sometimes TCP). SNMP is used, for example, by some server monitoring systems.

Port 443

Port 443 is the default HTTPS port (TCP). You may need to open it up for outbound requests (and inbound responses) temporarily because some RethinkDB installation instructions use `wget` over HTTPS to get the RethinkDB GPG key. Package managers might also get some packages using HTTPS.

Port 8125

If you set up a cluster-monitoring server, then StatsD will send UDP packets to Telegraf (on the monitoring server) via port 8125.

Port 8080

Port 8080 is the default port used by RethinkDB for its administrative web (HTTP) interface (TCP). While you *can*, you shouldn't allow traffic arbitrary external sources. You can still use the RethinkDB web interface by binding it to `localhost` and then accessing it via a SOCKS proxy or reverse proxy; see “Binding the web interface port” on the [RethinkDB page about securing your cluster](#).

Port 9984

Port 9984 is the default port for the BigchainDB client-server HTTP API (TCP), which is served by Gunicorn HTTP Server. It's *possible* allow port 9984 to accept inbound traffic from anyone, but we recommend against doing that. Instead, set up a reverse proxy server (e.g. using Nginx) and only allow traffic from there. Information about how to do that can be found in the [Gunicorn documentation](#). (They call it a proxy.)

If Gunicorn and the reverse proxy are running on the same server, then you'll have to tell Gunicorn to listen on some port other than 9984 (so that the reverse proxy can listen on port 9984). You can do that by setting `server.bind` to `'localhost:PORT'` in the BigchainDB Configuration Settings, where PORT is whatever port you chose (e.g. 9983).

You may want to have Gunicorn and the reverse proxy running on different servers, so that both can listen on port 9984. That would also help isolate the effects of a denial-of-service attack.

Port 28015

Port 28015 is the default port used by RethinkDB client driver connections (TCP). If your BigchainDB node is just one server, then Port 28015 only needs to listen on localhost, because all the client drivers will be running on localhost. Port 28015 doesn't need to accept inbound traffic from the outside world.

Port 29015

Port 29015 is the default port for RethinkDB intracluster connections (TCP). It should only accept incoming traffic from other RethinkDB servers in the cluster (a list of IP addresses that you should be able to find out).

Other Ports

On Linux, you can use commands such as `netstat -tunlp` or `lsof -i` to get a sense of currently open/listening ports and connections, and the associated processes.

Cluster-Monitoring Server

If you set up a cluster-monitoring server (running Telegraf, InfluxDB & Grafana), Telegraf will listen on port 8125 for UDP packets from StatsD, and the Grafana web dashboard will use port 3000. (Those are the default ports.)

1.11.10 Notes on NTP Daemon Setup

There are several NTP daemons available, including:

- The reference NTP daemon (`ntpd`) from ntp.org; see [their support website](#)
- [chrony](#)
- [OpenNTPD](#)
- Maybe [NTPsec](#), once it's production-ready
- Maybe [Ntimed](#), once it's production-ready
- [More](#)

We suggest you run your NTP daemon in a mode which will tell your OS kernel to handle leap seconds in a particular way: the default NTP way, so that system clock adjustments are localized and not spread out across the minutes, hours, or days surrounding leap seconds (e.g. “slewing” or “smearing”). There's a [nice Red Hat Developer Blog post](#) about the various options.

Use the default mode with `ntpd` and `chronyd`. For another NTP daemon, consult its documentation.

It's tricky to make an NTP daemon setup secure. Always install the latest version and read the documentation about how to configure and run it securely. See the notes on firewall setup.

Amazon Linux Instances

If your BigchainDB node is running on an Amazon Linux instance (i.e. a Linux instance packaged by Amazon, not Canonical, Red Hat, or someone else), then an NTP daemon should already be installed and configured. See the EC2 documentation on [Setting the Time for Your Linux Instance](#).

That said, you should check *which* NTP daemon is installed. Is it recent? Is it configured securely?

Ubuntu's ntp Package

The [Ubuntu 14.04 \(Trusty Tahr\)](#) package `ntp` is based on the reference implementation of an NTP daemon (i.e. `ntpd`).

The following commands will uninstall the `ntp` and `ntpdate` packages, install the latest `ntp` package (which *might not be based on the latest ntpd code*), and start the NTP daemon (a local NTP server). (`ntpdate` is not reinstalled because it's [deprecated](#) and you shouldn't use it.)

```
sudo apt-get --purge remove ntp ntpdate
sudo apt-get autoremove
sudo apt-get update
sudo apt-get install ntp
# That should start the NTP daemon too, but just to be sure:
sudo service ntp restart
```

You can check if `ntpd` is running using `sudo ntpq -p`.

You may want to use different NTP time servers. You can change them by editing the NTP config file `/etc/ntp.conf`.

Note: A server running an NTP daemon can be used by others for DRDoS amplification attacks. The above installation procedure should install a default NTP configuration file `/etc/ntp.conf` with the lines:

```
restrict -4 default kod notrap nomodify nopeer noquery
restrict -6 default kod notrap nomodify nopeer noquery
```

Those lines should prevent the NTP daemon from being used in an attack. (The first line is for IPv4, the second for IPv6.)

There are additional things you can do to make NTP more secure. See the [NTP Support Website](#) for more details.

1.11.11 Example RethinkDB Storage Setups

Example Amazon EC2 Setups

We have some scripts for deploying a *test* BigchainDB cluster on AWS. Those scripts include command sequences to set up storage for RethinkDB. In particular, look in the file `/deploy-cluster-aws/fabfile.py`, under `def prep_rethinkdb_storage(USING_EBS)`. Note that there are two cases:

1. **Using EBS (Amazon Elastic Block Store).** This is always an option, and for some instance types (“EBS-only”), it’s the only option.
2. **Using an “instance store” volume provided with an Amazon EC2 instance.** Note that our scripts only use one of the (possibly many) volumes in the instance store.

There’s some explanation of the steps in the [Amazon EC2 documentation](#) about making an Amazon EBS volume available for use.

You shouldn't use an EC2 "instance store" to store RethinkDB data for a production node, because it's not replicated and it's only intended for temporary, ephemeral data. If the associated instance crashes, is stopped, or is terminated, the data in the instance store is lost forever. Amazon EBS storage is replicated, has incremental snapshots, and is low-latency.

Example Using Amazon EFS

TODO

Other Examples?

TODO

Maybe RAID, ZFS, ... (over EBS volumes, i.e. a DIY Amazon EFS)

1.11.12 Licenses

Information about how the BigchainDB code and documentation are licensed can be found in the [LICENSES.md](#) file (in the root directory of the repository).

1.11.13 Installing BigchainDB on LXC containers using LXD

You can visit this link to install LXD (instructions here): [LXD Install](#)

(assumption is that you are using Ubuntu 14.04 for host/container)

Let us create an LXC container (via LXD) with the following command:

```
lxc launch ubuntu:14.04 bigchaindb
```

(ubuntu:14.04 - this is the remote server the command fetches the image from) (bigchaindb - is the name of the container)

Below is the `install.sh` script you will need to install BigchainDB within your container.

Here is my `install.sh`:

```
#!/bin/bash
set -ex
export DEBIAN_FRONTEND=noninteractive
apt-get install -y wget
source /etc/lsb-release && echo "deb http://download.rethinkdb.com/apt $DISTRIB_CODENAME main" | sudo
wget -qO- https://download.rethinkdb.com/apt/pubkey.gpg | sudo apt-key add -
apt-get update
apt-get install -y rethinkdb python3-pip
pip3 install --upgrade pip wheel setuptools
pip install ptpython bigchaindb
```

Copy/Paste the above `install.sh` into the directory/path you are going to execute your LXD commands from (ie. the host).

Make sure your container is running by typing:

```
lxc list
```

Now, from the host (and the correct directory) where you saved `install.sh`, run this command:

```
cat install.sh | lxc exec bigchaindb /bin/bash
```

If you followed the commands correctly, you will have successfully created an LXC container (using LXD) that can get you up and running with BigchainDB in <5 minutes (depending on how long it takes to download all the packages).

From this point onwards, you can follow the [Python Example](#) .

/transactions

GET /transactions/{tx_id}, [42](#)
GET /transactions/{tx_id}/status, [43](#)
POST /transactions/, [44](#)

Symbols

`__init__()` (bigchaindb.core.Bigchain method), 75

B

Bigchain (class in bigchaindb), 75

`block_election_status()` (bigchaindb.Bigchain method), 78

C

`create_block()` (bigchaindb.Bigchain method), 77

`create_genesis_block()` (bigchaindb.Bigchain method), 78

`create_transaction()` (bigchaindb.Bigchain method), 75

G

`get_blocks_status_containing_tx()` (bigchaindb.Bigchain method), 76

`get_last_voted_block()` (bigchaindb.Bigchain method), 78

`get_owned_ids()` (bigchaindb.Bigchain method), 77

`get_spent()` (bigchaindb.Bigchain method), 77

`get_status()` (bigchaindb.Bigchain method), 76

`get_transaction()` (bigchaindb.Bigchain method), 76

`get_tx_by_payload_uuid()` (bigchaindb.Bigchain method), 77

`get_unvoted_blocks()` (bigchaindb.Bigchain method), 78

H

`has_previous_vote()` (bigchaindb.Bigchain method), 78

I

`is_valid_block()` (bigchaindb.Bigchain method), 78

`is_valid_transaction()` (bigchaindb.Bigchain method), 77

P

`prepare_genesis_block()` (bigchaindb.Bigchain method), 78

S

`search_block_election_on_index()` (bigchaindb.Bigchain method), 76

`sign_transaction()` (bigchaindb.Bigchain method), 75

V

`validate_block()` (bigchaindb.Bigchain method), 77

`validate_fulfillments()` (bigchaindb.Bigchain method), 76

`validate_transaction()` (bigchaindb.Bigchain method), 77

`vote()` (bigchaindb.Bigchain method), 78

W

`write_block()` (bigchaindb.Bigchain method), 78

`write_transaction()` (bigchaindb.Bigchain method), 76

`write_vote()` (bigchaindb.Bigchain method), 78