This is a concise guide which should get you started with using BioCluster. If you have any questions about issues related to BioCluster that are not covered here, please use the BioCluster user's forum or the biocluster-users mailing list.

Installing BioCluster

Refer to the installation instructions on the download page.

Starting And Stopping BioCluster

BioCluster comes with an init script to control it, located at `/etc/init.d/atelis-biopeer`. You can use it as with any *init.d* script. To start BioCluster, issue `/etc/init.d/atelis-biopeer start` as root.

BioCluster includes an additional optional service, `/etc/init.d/atelis-biopeer-creditupdater`, used to update the credit of users after calls and to send CDRs to a central storage server you can define in the network settings. If you use credits or defined a storage server for CDRs, you should start this service too.

These services log their activities under `/var/log`. The BioCluster log is at `/var/log/peers.log`, and the credit updater log at `/var/log/creditupdater.log`.

Changing The Network Secret

Each peer in the BioCluster uses a shared secret to connect to the cluster. The network secret is the only configuration property that needs to be set separately on each machine in the cluster, since it's what separates different clusters from each other on the same network. To view the current secret, issue `biopeer-secret -p`. To change the secret, use `biopeer-secret -c` *newsecret* with your new secret. After changing the network secret, the BioCluster must be restarted.

The BioCluster Command Line Interface

The command line interface is a simple way to use the BioCluster. It can be used to connect to a peer in the cluster, query it for data and send updated data to it. BioCluster supports many kinds of data that are shared in the cluster, including extensions, trunks, AGI scripts, IVRs, voicemail messages and general network settings.

Connecting to the cluster through the command line interface can be done from any machine in the cluster, from a machine on the same network that's not part of the cluster, or from a remote machine that knows the address of some machine in the cluster. To run the command line interface, run `biopeer-cli`. If you're running the command line interface from a machine that's not part of the cluster, you'll have to know the network secret of the cluster you're connecting to. To set the secret to use, issue `secret mysecret` after starting the command line interface. Usually you'd then connect to the cluster by using `connect`, or if you're not on the same network as the cluster, `connect hostname.example.com` with the hostname or IP address of a machine in the cluster.

The command line interface includes help for all its commands. Issue `help` to see the list of commands available. Use `help` with the name of a command for help on a specific command, e.g. `help connect`. The command line interface can also run in batch mode, and has some command line arguments that can be passed to it. Issue `biopeer-cli --help` from the shell to see available command line arguments.

Querying Using The Command Line Interface

All data it the cluster is identified with a URI which starts with `bio:///`. For example, SIP extension 1000 has the URI `bio:///sip/1000`. The configuration property interconnect-tech has the URI bio:///configuration/interconnect-tech...
Querying for a data units is done using the command `query` with the URI or list of URIs to query. For example:
`query bio:///route/route1` - query for the priority route named *route1*.
`query bio:///dundi/` - query for all DUNDi peers in the network.
Data received as a result of queries is placed in the local cache of the command line interface, which can be viewed using the `db` command.

Adding And Updating Data

In order to add or update data in the cluster, use the command `put`. See `help put` for a list of different data families you can add, as each data family has different arguments that are passed to `put`. To see instructions on adding a data unit of a specific family, use `help put` *`family`* with the name of the data family, e.g. `help put configuration`.
Examples for `put`:
`put sip 1000 1000 local dynamic friend all ulaw;alaw;gsm` - Add a SIP peer *1000*, with secret *1000*, local context, dynamic host, type "friend", "all" as disallow field and "ulaw;alaw;gsm" as allow field.
`put configuration network-timeout INTEGER 4000` - Set the *network-timeout* configuration property to 4000.
After adding data to add or update using `put`, the data needs to be sent to the cluster using `update`. To see a list of all the data families your BioCluster supports, with the data currently in your local cache, issue `db`.

More Command Line Interface Features

There are many more features in the BioCluster command line interface you can discover. Use `help` to learn more or ask in the user's forum or the biocluster-users mailing list.