

Starting Clustered AtScale

Here you can find how to start a Clustered AtScale environment. In short, you should first start the Coordinator node, then the Leader node, and finally the Replica node.

Prerequisites

You can access the Linux command line shell session at the AtScale system, and have account that allows you to manage the system.

Procedure

1. Establish a shell session to each AtScale node: Engine 1, Engine 2, and Coordinator node.
2. Switch user to the atscale (required for the remaining steps):

```
sudo su - atscale
```

3. Start the AtScale coordinator node by running the `atscale_start` command:

```
/opt/atscale/bin/atscale_start
```

4. Wait for all services to start.

To verify all services are started, you can execute the `atscale_service_control status` command. The result should look like in the example below:

```
/opt/atscale/bin/atscale_service_control status
agent                RUNNING  pid 23703, uptime 0:00:11
coordinator          RUNNING  pid 23704, uptime 0:00:11
egress               RUNNING  pid 23705, uptime 0:00:11
ingress              RUNNING  pid 23707, uptime 0:00:11
service_registry     RUNNING  pid 23702, uptime 0:00:11
```

5. Start the AtScale Engine node that was previously the AtScale Database PostgreSQL leader node using the `atscale_start` command:

```
/opt/atscale/bin/atscale_start
```

6. Wait for all services to start.

To verify all services are started you can execute the `atscale_service_control status`` command as shown above. Note there will be additional services running on engine nodes versus the coordinator node.

7. Start the AtScale Engine node that was previously the AtScale Database PostgreSQL Replica node using the `atscale_start` command:

```
/opt/atScale/bin/atScale_start
```

8. Wait for all services to start.

Again, you can use the `atScale_service_control status` command to verify all services are started.

The AtScale Cluster should now be started.

More Information

[Stopping Clustered AtScale](#)