Convert Stand-Alone AtScale To Clustered AtScale

This document describes conversion of Stand-Alone AtScale to Clustered AtScale. These instructions are for CentOS/RedHat operating systems. You may have to modify these instructions for your supported OS.

Before You Begin

This document refers to the host where AtScale is currently installed in a single-instance deployment as the current AtScale node.

- **Important**: You must have a directory service configured (Google G Suite Directory, LDAP, or Microsoft Active Directory) that AtScale can connect to for managing users and security on your cluster. AtScale clusters cannot use AtScale's embedded directory service. AtScale does not provide a migration utility for moving users and groups from its embedded directory service to a supported external one.
- **Important**: You must have an external load balancer configured. See Configure an External Load Balancer for instructions.
- Unless otherwise specified, commands are to be run as the `atscale` user. If you installed AtScale to run as a different user than `atscale`, run the following commands as that user.
- Review Architecture of an AtScale Cluster.

Preparing The Machines

Machines required

Running Clustered AtScale requires a customer-provided external load balancer and three AtScale machines: one that will run AtScale Coordinator (a stripped-down version of AtScale), and two that will run the full AtScale application in the cluster.

These AtScale machines will be referred to as follows:

- Coordinator Host - runs the AtScale Coordinator
- Master Host - runs full AtScale (only designated "Master" during installation but not at run-time)
- Standby Host - runs full AtScale (only designated "Standby" during installation but not at run-time)

Upgrade a Stand-Alone installation before converting to a clustered installation

Before executing this procedure, the stand-alone installation must be upgraded to the same version as the downloaded Clustered AtScale binary. Complete these steps before continuing: Upgrading Stand-Alone AtScale.

Backup AtScale:

On the stand-alone host, backup the AtScale application before the conversion:

1. Stop AtScale:

   ```
   /opt/atscale/bin/atscale_stop
   ```

2. After AtScale stops, copy to backup:

   ```
   cp -Rp /opt/atscale ~/atscale_standalone_backup
   ```

3. Start AtScale:

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

Converting To Clustered AtScale

1. Verify that AtScale is up and running in stand-alone mode. The existing host will be referred to as the "Master" host in these instructions.
2. If you wish to run the system with an account other than the default `atscale` then assign the desired user name to the ATSCALE_USER environment variable:

   ```
   export ATSCALE_USER=myatscaleuser
   ```

3. Copy the installation package to the new "Standby" host. From the "Standby" host, as root, install the AtScale installer package, but do not run the `confgurator.sh` script. If you are not root, but have sudo privileges install the rpm with the `sudo -E` option to preserve the ATSCALE_USER environment variable in the subshell.

   ```
   rpm -i installer_package_filename.rpm
   ```

   or if using Debian packages:

   ```
   dpkg --install installer_package_filename.deb
   ```

4. Copy the installation package to the new "Coordinator" host. From the "Coordinator" host, as root, install the AtScale installer package, but do not run the `confgurator.sh` script. If you are not root, but have sudo privileges install the rpm with the `sudo -E` option to preserve the ATSCALE_USER environment variable in the subshell.

```
rpm -i installer_package_filename.rpm
```

or if using Debian packages:

```
dpkg --install installer_package_filename.deb
```

5. On the currently running stand-alone "Master" host, edit `/opt/atscale/conf/atscale.yaml`:

    1. Change `loadbalancer_dns_name` to the dns of the load balancer
    2. Under "hosts", add a config block for the new "Standby" host with two services, `atscale` and `coordinator`. Verify that the `override.coordinator.id` property has a unique value. For an example see the `hosts -> name -> atscale-02` section in Figure 1.
    3. Under "hosts", add a config block for the new "Coordinator" host with only one service, `coordinator`. Verify that the `override.coordinator.id` property has a unique value. For an example see the `hosts -> name -> atscale-03` section in Figure 1.

    Figure 1: Example of a basic atscale.yaml file.

```
installation_location: "/opt/atscale"
service_account: "atscale"
loadbalancer_dns_name: "atscale-lb"

tls:
    enabled: false
    certificate: "/opt/atscale/conf/server.cert"
    key: "/opt/atscale/conf/server.key"

kerberos:
    enabled: false
    keytab: "/opt/atscale/conf/atscale.keytab"
    principal: "atscale/atscale-01@REALM"

hosts:
    - name: atscale-01
      services:
          - atscale
          - coordinator
      override:
          coordinator:
            id: 12
    - name: atscale-02
      services:
          - atscale
          - coordinator
      override:
          coordinator:
            id: 13
    - name: atscale-03
      services:
          - coordinator
      override:
          coordinator:
            id: 14
```

4. **Advanced Configuration:** If the machine's host name, as returned by the `hostname` command, is not publicly routable then you must make several changes to `atscale.yaml`. This situation is common when running in cloud environments like AWS.

1. As the ATSCALE_USER user (`atscale` by default), open `/opt/atscale/conf/atscale.yaml`
2. **The following properties must be set to a routable host name. See Figure 2 for an example.**

    1. `loadbalancer_dns_name` (publicly-routable FQDN of the external load balancer)
    2. Each `hosts -> name` property
    3. `kerberos -> principal` (Only needed if setting up Kerberos)

3. Add an alias line under the hosts.name property. This alias must be the same as the host's machine name returned by the `hostname` command. See Figure 2 for an example.

    Figure 2: Example of atscale.yaml configured with routable and non-routable host names.
    In this example the host's names are set to routable domain names whereas their alias values are set to the value returned by the `hostname` command. The load balancer uses a publicly routable FQDN.

```
installation_location: "/opt/atscale"
service_account: "atscale"
loadbalancer_dns_name: "atscale-service.acme.com"

tls:
  enabled: false
  certificate: "/opt/atscale/conf/server.cert"
  key: "/opt/atscale/conf/server.key"

kerberos:
  enabled: false
  keytab: "/opt/atscale/conf/atscale.keytab"
  principal: "atscale/atscale-service.acme.com@REALM"

hosts:
  - name: atscale-01.acme.com
    alias: ip-172-18-2-36.us-west-1.compute.internal
    services:
      - atscale
      - coordinator
    override:
      coordinator:
        id: 12
  - name: atscale-02.acme.com
    alias: ip-172-18-2-37.us-west-2.compute.internal
    services:
        - atscale
        - coordinator
    override:
        coordinator:
          id: 13
  - name: atscale-03.acme.com
    alias: ip-172-18-2-38.us-west-3.compute.internal
    services:
        - coordinator
    override:
        coordinator:
          id: 14
```

6. If the host has multiple private IP addresses, you must set the IP address for the AtScale Service Registry by overriding the service_registry `bind_addr` property.

   Figure 3: Example of atscale.yaml hosts section configured to bind the AtScale Service Registry Service to a specific private IP address.
   In this example, the host contains multiple private IP addresses and the AtScale Service Registry is configured to use a specific address.

   ```
   hosts:
   - name: atscale-01.acme.com
     alias: ip-172-18-2-36.us-west-1.compute.internal
     override:
        service_registry:
           bind_addr: 1.2.3.4
   ```

   Alternatively, set the `ATSCALE_BIND_ADDRESS` environment variable. For each host that has more than one private IP address, set `ATSCALE_BIND_ADDRESS` to the desired private IP. Once completed, run `atscale_stop` (if AtScale is running) followed by `atscale_start`, OR run `configurator.sh` in `--apply` mode.

   ```
   export ATSCALE_BIND_ADDRESS=1.2.3.4
   ```

7. Copy the modified `/opt/atscale/conf/atscale.yaml` from the "Master" node to `/opt/atscale/conf/atscale.yaml` on both the "Standby" and "Coordinator" nodes.

8. Activate the "Coordinator" host:

   1. On the "Coordinator" host, as the ATSCALE_USER user (`atscale` by default), run the installed AtScale's `configurator.sh` script in `--activate` mode. If running AtScale under a different account, alter the switch user command to switch to the desired user (e.g. `su - myatscaleuser`):

      ```
      su - atscale
      cd /opt/atscale/versions/<package_version>
      ./bin/configurator.sh --activate
      ```

   2. You will be asked to confirm that you want to apply this configuration. Answer `y` to start the activation process, which will configure and run AtScale.

      Note: If an error occurs, you must perform a full reinstallation. You cannot simply rerun the installer. See Uninstall AtScale.

9. Apply configuration to the "Master" host:

1. On the "Master" host, as the ATSCALE_USER user (`atscale` by default), run the installed AtScale's `configurator.sh` script in `--apply` mode. If running AtScale under a different account, alter the switch user command to switch to the desired user (e.g. `su - myatscaleuser`):

   ```
   su - atscale
   cd /opt/atscale/versions/<package_version>
   ./bin/configurator.sh --apply
   ```

2. You will be asked to confirm that you want to apply this configuration. Answer `y` to start the config apply process, which will configure and run AtScale.

10. Activate the "Standby" host:

    1. On the "Standby" host, as the ATSCALE_USER user (`atscale` by default), run the installed AtScale's `configurator.sh` script in `--activate` mode. If running AtScale under a different account, alter the switch user command to switch to the desired user (e.g. `su - myatscaleuser`):

       ```
       su - atscale
       cd /opt/atscale/versions/<package_version>
       ./bin/configurator.sh --activate
       ```

    2. You will be asked to confirm that you want to apply this configuration. Answer `y` to start the activation process, which will configure and run AtScale.

> **Note:** If an error occurs, you must perform a full reinstallation. You cannot simply rerun the installer. See Uninstall AtScale.

What To Do Next

After you have installed AtScale on all of the nodes, doing the following:

- Confirm that Business Intelligence users can log in to Design Center and run queries against the Engine.

  > **Note:** If you have trouble logging in to AtScale as a non-super user after upgrading to Clustered AtScale, confirm that you have configured an external directory service. The AtScale embedded directory is not supported for Clustered AtScale installations. Log in to Design Center as an admin user to correct the configuration.

- If you wish to convert an AtScale cluster to a stand-alone installation, contact AtScale Customer Support.