PERFORMING A FULL GRACEFUL RESTART OF A RADIANTONE FID CLUSTER

SUMMARY

This document details how to perform a full restart of a Radiantone FID cluster.

OVERVIEW

Radiant Logic strongly recommends performing a complete **all-nodes restart strategy** when performing maintenance on a cluster. The all nodes strategy is the safest, most reliable way to ensure cluster stability during maintenance**.** *It is a Radiant Logic best practice*. The all-nodes strategy creates a temporary complete cluster outage. If this temporary outage is not an option, the rolling restart strategy is the recommended approach; the guide for which is linked below.

Performing Rolling Restarts in a RadiantOne FID Cluster.

A full restart on a RadiantOne FID cluster is performed by completely shutting down each node in the cluster one by one, until all the nodes are shut down and restarting the nodes gradually, until all the nodes are back up again. This requires a specific order to be followed to ensure stability and data consistency.

1. Shut down the follower nodes first.
2. Shut down the leader node last.
3. Perform required maintenance.
4. Start the previous leader node first.
5. Start the follower nodes.

The full procedure is detailed in the next sections of this guide. The guide is formatted to be an easy to follow reference, with each major procedure broken down into grouped, detailed steps. Light gray indented text provides additional context for why steps are performed and what to expect. The overall process is nearly identical in Windows and Linux. But they are treated separately for clarity and to account for differing commands. Please refer to the appropriate section based on your operating system.

Radiant Logic recommends reading this guide in its entirety before using it as only a reference. When following this guide, successfully complete each step before moving to the next. If you have questions about this guide or its procedure, please contact Radiant Logic Support by emailing support@radiantlogic.com or creating a support case through the Support portal.

DETAILED FULL RESTART PROCEDURE ON WINDOWS

A. Shut down a follower node.

Note: All steps in this section are performed on a single follower node.

1. Select a follower node.

2. Record all currently running services and components.

> This is best done through the RadiantOne Main Control Panel. Take care to check if the server is running services or components that provide additional functionality such as Global Sync topologies, or capture connectors. These components will need to be started again when bringing the server back online.

3. Shutdown all RadiantOne services on the follower by running:
`%RLI_HOME%\bin\advanced\stop_servers.bat`

> This script attempts to gracefully stop each RadiantOne service's in the proper order. It may take a few minutes for the script to complete, so be patient.

4. Kill any lingering Java processes associated with RadiantOne components.

> This can be done using the Task Manager by finding, then ending all java.exe processes. Once all Java processes are stopped, the node will be offline and ready for maintenance.

B. Shut down the remaining follower nodes.

Repeat steps in the previous section (“A. Shut down a follower node”) for the remaining follower nodes, one at a time.

C. Shut down the leader node.

Repeat steps in the previous section (“A. Shut down a follower node”) for the leader node.

D. Perform all necessary maintenance on the server.
   Do not proceed until all maintenance is complete

E. Start the Control Panel service on all nodes.
   Starting the Control Panel also starts ZooKeeper. ZooKeeper needs to be running for the node to properly rejoin the cluster. You can confirm ZooKeeper is running by checking the node’s status indicators on the Main Control Panel. It will appear green (● ZK) if online.

   Also, the Control Panel sometimes takes a few minutes to start, so be patient. You'll know it's online when you can sign in to the Main Control Panel directly on the follower.

F. Start the RadiantOne FID service on the leader node.
   1. Start the FID service on the node that was the previous leader (which was stopped last).

   Starting the RadiantOne FID service will start the directory (LDAP/LDAPS) and administrative (HTTP/HTTPS) components. Wait for the component status indicators to turn green on the Main Control Panel before continuing.

Ensure that the yellow triangle (▲) identifying the current leader shows up on the Main Control Panel.

2. Start any remaining services and components that normally run on this specific node (e.g., Global Sync topologies, connectors, etc).

G. Start the RadiantOne FID service on the follower nodes.
   1. Start the FID service on one of the follower nodes.

   2. Wait for it to join the cluster as a follower and all the component status indicators to turn green on the Main Control Panel.

   3. Start any remaining services and components that normally run on this specific node (e.g., Global Sync topologies, connectors, etc).

   4. Repeat for all the nodes one by one, until all the nodes are up and running.

   5. Return to the Main Control Panel to confirm all cluster nodes and services are now online.

DETAILED FULL RESTART PROCEDURE ON LINUX

A. Shut down a follower node.

Note: All steps in this section are performed on a single follower node.

1. Select a follower node.

2. Record all currently running daemons and components.

> This is best done through the RadiantOne Main Control Panel. Take care to check if the server is running services or components that provide additional functionality such as Global Sync topologies, or capture connectors. These components will need to be started again when bringing the server back online.

3. Shutdown all RadiantOne services on the follower by running:
`%RLI_HOME%\bin\advanced\stop_servers.sh`

> This script attempts to gracefully stop each RadiantOne service's in the proper order. It may take a few minutes for the script to complete, so be patient.

4. Kill any lingering Java processes associated with RadiantOne components.

> This can be done using the command `PS -EF | GREP JAVA` to identify running Java processes, then using the `KILL` command to stop them one by one. Once all Java processes are stopped, the node will be offline and ready for maintenance.

B. Shut down the remaining follower nodes.

Repeat steps in the previous section ("A. Shut down a follower node") for the remaining follower nodes, one at a time.

C. Shut down the leader node.

Repeat steps in the previous section (“A. Shut down a follower node”) for the leader node.

D. Perform all necessary maintenance on the server.

Do not proceed until all maintenance is complete

E. Start the Control Panel daemon on all nodes.

Starting the Control Panel also starts ZooKeeper. ZooKeeper needs to be running for the node to properly rejoin the cluster. You can confirm ZooKeeper is running by checking the node’s status indicators on the Main Control Panel. It will appear green (● ZK) if online.

Also, the Control Panel sometimes takes a few minutes to start, so be patient. You'll know it's online when you can sign in to the Main Control Panel directly on the follower.

F. Start the RadiantOne FID daemon on the leader node.

1. Start the FID daemon on the node that was the previous leader (which was stopped last).

Starting the RadiantOne FID daemon will start the directory (LDAP/LDAPS) and administrative (HTTP/HTTPS) components. Wait for the component status indicators to turn green on the Main Control Panel before continuing.

Ensure that the yellow triangle (▲) identifying the current leader shows up on the Main Control Panel.

2. Start any remaining services and components that normally run on this specific node (e.g., Global Sync topologies, connectors, etc).

G. Start the RadiantOne FID daemon on the follower nodes.

1. Start the FID daemon on one of the follower nodes.

2. Wait for it to join the cluster as a follower and all the component status indicators to turn green on the Main Control Panel.

3. Start any remaining services and components that normally run on this specific node (e.g., Global Sync topologies, connectors, etc.).

4. Repeat for all the nodes one by one, until all the nodes are up and running.

5. Return to the Main Control Panel to confirm all cluster nodes and services are now online.