Splunk® Search Head Clustering (SHC) Cheat Sheet (v2.0)

https://www.aplura.com/cheatsheets

Search Head Clustering

A Search Head Cluster is a group of Search Heads that work together to create high availability and horizontal scaling by sharing configurations, job schedules, and search artifacts

- Captain is the Member whom delegates any activity on the Members that is not an ad-hoc search through out the SHC effectively spreading the work load. This includes controlling replication and pushing knowledge bundles to Search Peers (Indexers.)
  - The Captain is a roll that rotates between Members using an election process. To win the election the Member must receive majority votes from all the other Members
  - The Captain election process starts when either the current Captain restarts, some of the Members are separated on the network, and/or the current Captain steps down due to missing a majority of the other Members
- Deployer is the instances that pushes apps, user data, and configurations to the members
- Members are the instances which searches and jobs are preformed on

Note: It is always recommended to use a load balancer with “sticky sessions” enabled to allow for continuity of which system a user is using

Note: Enterprise Security should only be used on a SHC if the number of concurrent users and/or concurrent searches exceeds the capacity of a maximum speced stand alone instance.

Deployment Steps

1 Create the Deployer

- To create a Deployer configure server.conf
- Restart the Member after this command

Conf
```
[shclustering]
pass4SymmKey = YurPwd
shcluster_label = YurName
```

2 Initialize the Members

- To Initialize the Search Heads run the follow command **after** you modify it with your chosen information on **each** Search Head Member

CLI
```
/splunk init shcluster-config -mgmt_uri https://
<LocalSearchheadAddress>:8089 -replication_port <yourPort> -
conf_deploy_fetch_url https://<DeployerAddressHere>:8089 -secret
shc1node -shcluster_label shcluster1
```

3 Boot-Strap the Members

- Boot-strapping initializes the Captain role for the SHC joining the Members together. **Only** run this command on 1 Member

CLI
```
/splunk bootstrap shcluster-captain -servers_list "https://spl-
srch01:8089,https://spl-srch02:8089,https://spl-srch03.edu:
8089"
```

- Restart the member after this command

4 Apply the First Search Head Cluster Bundle

- It does not matter which Member is used in the command
- Always add the “preserve-lookups true” flag

CLI
```
/splunk apply shcluster-bundle --answer-yes -target https://
<AnyMemberNameInSHC>:8089 -preserve-lookups true
```

TIPS

- The “boot-strap” command and the first “apply shcluster-bundle” command will take a short while to run
- Their will be a different version of the initialize command for each Member.
- The order of initializing Members does not matter
- Always build a SHC from a new install of Splunk

CLI Splunk CLI command

Conf Splunk .conf file configuration

APLURA
Many Solutions, One Goal.

v2.0

Management Tricks

- You have the option to remotely load the Deployer from another Splunk Server.

CLI
```
/splunk apply shcluster-bundle --answer-yes -target https://<Any1MemberNameInSHC>:8089 -preserve-lookups true -uri https://<deployeraddress>:8089
```

TIPS

- This link discusses how you can run most other Splunk commands remotely

Troubleshooting

- Replication errors: If a Seach Head failes to sync more than 20 consecutive times then it becomes an Error and requires a forced sync.

CLI
```
/splunk resync shcluster-replicated-config
```

- Replication errors: Make sure the Replication Factor is the same on all Search Head Cluster Members.
- Replication errors: Check the Members to see if an extremely large Lookup file exists or if the members are timing out while pushing/pulling configurations.
- If a SHC looses its majority members for a prolonged period of time and can not elect a Captain by itself. Manually set a temporary Captain as static. Revert the SHC to dynamic after the majority of Members are restored. **Do not** leave the SHC in a static Captain state.

KV store

Once the SHC is created the Members also cluster the KV store. A single instance assumes the role KV store Captain.

- The KV store Captain handles all write requests for the cluster's KV store. Any other KV store request is handled locally on the instance.
- All instances in the SHC sync from the KV store Captain.

To view the KV store status:

CLI
```
/splunk show kvstore-status
```

If the KV store becomes out of sync you can resync the KV store manually from the SHC Captain:

CLI
```
/splunk rsync kvstore
```

Add the flag below to the command above if you wish to use a different Member as the source for the sync instead of the SHC Captain

CLI
```
-source sourceId
```

Gotchas

- The setup guide on back is meant for a **new** SHC deployment only. A SHC **migration** involves more steps and is more complicated.
- The "Deployer" is **!=** the "Deployment Server"
- The "KV store replication" is **!=** the "SHC replication"
- A SHC has a limit of 5000 active and/or unexpired alerts
- Once the SHC is created, avoid modifying conf files directly on the members themselves. The SHC **will not** replicate changes made directly to any conf file.
- The Captain role is chosen by election and will not always be the same server. Use the Monitoring Console or Splunk command "show shcluster-status" to find the Captain
- The majority of the SHC Members must be online for dynamic Captain selection (Captain election) to occur.
- Make sure to have necessary network ports for the KV store, SHC replication, and management open before you initialize the SHC Members

Replication

What causes SHC Replication?

- Changes made through: Splunk Web, Splunk CLI commands, and REST API

What gets Replicated (by Default)?

alert_actions	manager	searchscripts
authentication	models	segmenters
authorize	multikv	tags
datamodels	nav	times
event_renderers	panels	transforms
eventtypes	passwd	transactiontypes
fields	passwords	ui-prefs
html	props	user-prefs
literals	quickstart	views
lookups	savedsearches	viewstates
macros	searchbnf	workflow_actions

lookup table, datamodel JSON, nav XML, meta files

- Change what is replicated in the SHC by modifying the replication whitelist in server.conf on **all** members

CLI Splunk CLI command

Conf Splunk .conf file configuration

APLURA
Many Solutions, One Goal.