

Chassis Cluster Upgrade with Minimal Downtime (v1.0) SRX HighEnd (SRX5400, SRX5600, SRX5800)

Prerequisites:

- Prerequisites may be performed outside of MW with no impact to traffic
- Console connections setup to both chassis cluster nodes are necessary to allow unique config adjustments and due to device power off via 'halt' method used
- Download Junos software from [Juniper Download website](#)
- Backup current configuration
 - Local device file storage – 'show configuration | save /var/tmp/<name>
 - To attached USB – Use [KB12880](#) to mount USB then 'save configuration | save /var/tmp/usb'
- Upload Junos OS image to Device storage
 - e.g., /var/tmp/junos-install-srx5000-x86-64-18.4R2.7.tgz
- Verify upgrade image package to current configuration
 - >request system software validate <image location>
- Temporarily disable connected switch settings for mac moves/duplications such as 'mac-move-limit' and 'duplicate-mac-detection' due to possible duplicate location mac addresses during Step 12.

Upgrade Directions:

The below steps assume that Node0 is the primary for control plane (RG0) and data plane (RG1+) and configured with a higher priority than the secondary node.

As needed, please failover all RG groups to primary node
>request chassis cluster failover redundancy-group [x] node 0

Node0 Directions	Node1 Directions
<p>1. Disable all physical interfaces used for transit traffic on Node1 (secondary node). <i>Note: Alternatively, you may physically remove cables or 'shut' connected device interfaces.</i></p> <p>e.g., <code>set interfaces xe-12/0/0 disable</code> <code>set interfaces xe-12/3/0 disable</code></p> <p>2. Disable TCP SYN check and sequence check</p> <p><code>set security flow tcp-session no-syn-check</code> <code>set security flow tcp-session no-sequence-check</code></p> <p>3. Deactivate preempt for all RG1+</p> <p><code>deactivate chassis cluster redundancy-group 1 preempt</code></p> <p>4. Deactivate all interface-monitor and ip-monitoring</p> <p><code>deactivate chassis cluster redundancy-group 1 interface-monitor</code> <code>deactivate chassis cluster redundancy-group 1 ip-monitoring</code></p>	

<p>5. Commit the configuration</p> <p><code>commit</code></p> <p>6. Adjust control-ports control link and delete fab interfaces via configuration.</p> <p>Change the control ports to non-existing ports and delete fabric ports.</p> <ul style="list-style-type: none"> - Control ports need to be set in any SPC port on the device, which does not have a physical connection - Fabric ports can be set in any IOC slots (existing or not) on the device. A simple way is that change the fabric ports to undefined port numbers (e.g., port 40) on the same slot. <p><i>Note: Before configuration adjustments, make a note of current configured control and fabric port interfaces for later addition to configuration in step 19 & 21</i></p> <p>e.g.,</p> <pre>show chassis cluster control-ports show interfaces fab0 show interfaces fab1</pre> <pre>delete chassis cluster control-ports set chassis cluster control-ports fpc 10 port 0 (SPC port) set chassis cluster control-ports fpc 22 port 0 (SPC port) delete interface fab0 delete interface fab1</pre> <p><i>NOTE: If configured using dual control links, you need to also include the configuration change for the second control link.</i></p> <pre>set chassis cluster control-ports fpc 10 port 1 (SPC port) set chassis cluster control-ports fpc 22 port 1 (SPC port)</pre> <p>7. Commit the configuration</p> <p><code>commit and-quit</code></p> <p><i>NOTE: Upon commit completion, the following errors maybe generated due to control link down and may report commit failure for Node0. These are expected error messages. Under such errors, the "candidate" configuration is now in "active" configuration on Node1 and do not require another commit.</i></p> <pre>{primary:node0}[edit] root@srx5k# commit</pre> <pre>node0: configuration check succeeds error: error communicating with error: remote commit-configuration failed on node1 error: commit failed error: Connection to node1 has been broken error: remote unlock-configuration failed on node1</pre> <p><i>Under condition of commit-failure error, if you want to exit the configuration mode on Node0, you may execute "commit" again on Node0.</i></p> <p>e.g.,</p> <pre>{primary:node0}[edit] root@srx5k# exit The configuration has been changed but not committed Discard uncommitted changes? [yes,no] (yes) no <<< SHOULD be "no"</pre>	<p>6. Adjust control-ports control link and delete fab interfaces via configuration.</p> <p><i>NOTE: Upon commit completion, the following errors maybe generated due to control link down and may report commit failure for Node0. These are expected error messages. Under such errors, the "candidate" configuration is now in "active" configuration on Node1 and do not require another commit.</i></p> <p><i>NOTE: It is expected that Node1 will report transitioning to Primary for all RGs due to loss of Control and Fabric links but will have no impact on traffic due to disabling of transit interfaces in Step 1.</i></p>
---	--

Exit aborted

```
{primary:node0}[edit]
root@srx5k# commit and-quit
node0:
commit complete
Exiting configuration mode

{primary:node0}[edit]
root@srx5k# exit
```

NOTE: Before starting Node1 upgrade, make sure the "active" configuration reflects the changes made on step 6 and Node0 reports Node1 as lost.

e.g.,

```
{primary:node0}
root@srx5k> show configuration | display set | match "control-ports | fab[01]"
set chassis cluster control-ports fpc 10 port 0
set chassis cluster control-ports fpc 22 port 0
```

```
{primary:node0}
root@srx5k> show chassis cluster status
...
Cluster ID: 5
Node Priority Status Preempt Manual Monitor-failures

Redundancy group: 0 , Failover count: 1
node0 200 primary no no None
node1 0 lost n/a n/a n/a

Redundancy group: 1 , Failover count: 1
node0 200 primary no no None
node1 0 lost n/a n/a n/a
```

NOTE: Before starting Node1 upgrade, make sure the "active" configuration reflects the changes made on step 6 and Node1 reports Node0 as lost.

e.g.,

```
{primary:node1}
root@srx5k> show configuration | display set | match
"control-ports | fab[01]"
```

```
set chassis cluster control-ports fpc 10 port 0
set chassis cluster control-ports fpc 22 port 0
```

```
{primary:node1}
root@srx5k> show chassis cluster status
...
Cluster ID: 5
Node Priority Status Preempt Manual Monitor-failures

Redundancy group: 0 , Failover count: 1
node0 0 lost n/a n/a n/a
node1 100 primary no no None

Redundancy group: 1 , Failover count: 1
node0 0 lost n/a n/a n/a
node1 100 primary no no None
```

Start Node1 upgrade

8. Upgrade Junos OS on the Node1

```
request system software add no-copy <install-package>
```

Note: If upgrade was verified previously as part of prerequisite steps, 'no-validate' may be used to speed up install process.

9. Reboot

```
request system reboot
```

10. After Node1 completes boot process, verify the following before moving to next step:

- Updated Junos OS
- All FPCs and PICs are online (may take upto 15 minutes depending on the type and number of FPCs)
- Node1 should be in primary state for all RGs, and reporting Node0 as 'lost'
- No major alarms being displayed

```
show version
show chassis fpc pic-status
show chassis cluster status
show chassis alarms
show system alarms
```

<p>11. Before failing over to Node1, it is best to verify the configuration change will occur successfully by testing a commit first then, - disable all physical interfaces for transit traffic on Node0 - enable all physical interfaces for transit traffic on Node1</p> <p>e.g.,</p> <pre>{primary:node0}[edit] root@srx5k# set interfaces reth0 description TEST {primary:node0}[edit] root@srx5k# commit node0: commit complete {primary:node0}[edit] root@srx5k# rollback 1 load complete {primary:node0}[edit] root@srx5k# commit node0: commit complete</pre> <p>e.g.,</p> <pre>set interfaces xe-0/0/0 disable set interfaces xe-0/3/0 disable delete interfaces xe-12/0/0 disable delete interfaces xe-12/3/0 disable commit check</pre> <p><i>NOTE: Enable all physical interfaces of Node1 that were disabled on step 1.</i></p> <p><i>NOTE: If there are any commit conflicts, they need to be resolved before moving to the next step.</i></p> <p><i>NOTE: Alternatively, you may physically remove cables or 'shut' connected device interfaces for Node0.</i></p> <p>12. Commit the configuration simultaneously on both nodes. This will cause all of the traffic to failover to Node1</p> <pre>commit</pre>	<p><i>NOTE: Priorities of RG1+ will report priority 0 as part of normal behavior.</i></p> <pre>{primary:node1}[edit] root@srx5k# run show chassis cluster status ... Cluster ID: 6 Node Priority Status Preempt Manual Monitor-failures Redundancy group: 0 , Failover count: 1 node0 0 Lost n/a n/a n/a node1 100 primary no no None Redundancy group: 1 , Failover count: 1 node0 0 Lost n/a n/a n/a node1 0 primary no no CS</pre> <p>11. Before failing over to Node1, it is best to verify the configuration change will occur successfully by testing a commit first then, - disable all physical interfaces for transit traffic on Node0 - enable all physical interfaces for transit traffic on Node1</p> <p>e.g.,</p> <pre>{primary:node1}[edit] root@srx5k# set interfaces reth0 description TEST {primary:node1}[edit] root@srx5k# commit node1: commit complete {primary:node1}[edit] root@srx5k# rollback 1 load complete {primary:node1}[edit] root@srx5k# commit node1: commit complete</pre> <p>e.g.,</p> <pre>set interfaces xe-0/0/0 disable set interfaces xe-0/3/0 disable delete interfaces xe-12/0/0 disable delete interfaces xe-12/3/0 disable commit check</pre> <p><i>NOTE: Enable all physical interfaces of Node1 that were disabled on step 1.</i></p> <p><i>NOTE: If there are any commit conflicts, they need to be resolved before moving to the next step.</i></p> <p><i>NOTE: Alternatively, you may physically add cables or 'un-shut' connected device interfaces for Node1.</i></p> <p>12. Commit the configuration simultaneously on both nodes. This will cause all of the traffic to failover to the Node1</p> <pre>commit</pre>
<p><i>NOTE: The total amount of downtime will vary depending on switching/routing environment. (e.g., dynamic routing, STP, MSTP, RSTP, VSTP, edge, PortFast, and etc).</i></p>	
	<p>13. Verify traffic is passing through Node1</p> <pre>show security flow session summary monitor interface traffic</pre>

Start Node0 upgrade

14. Upgrade Junos OS on the Node0

```
request system software add no-copy no-validate <install-package>
```

15. Reboot

```
request system reboot
```

16. After Node0 completes boot process verify the following, before moving to next step:

- Updated Junos OS
- All FPCs and PICs are online (may take upto 15 minutes depending on the type and number of FPCs)
- Node0 should be in primary state for all RGs, and reporting Node1 as 'lost'
- No major alarms being displayed

```
show version
show chassis fpc pic-status
show chassis cluster status
show chassis alarms
show system alarms
```

NOTE: Priorities of RG1+ will report priority 0 as part of normal behavior.

```
{primary:node0}
root@srx5k> show chassis cluster status
...
Cluster ID: 6
Node Priority Status Preempt Manual Monitor-failures

Redundancy group: 0 , Failover count: 1
node0 200 primary no no None
node1 0 Lost n/a n/a n/a

Redundancy group: 1 , Failover count: 1
node0 0 primary no no CS
node1 0 Lost n/a n/a n/a
```

17. Before re-configuring control and re-configuring fab interfaces, enable interface-monitor which was disabled in step 4.

e.g.,

```
activate chassis cluster redundancy-group 1 interface-monitor
commit check
```

18. Commit the configuration on both nodes

```
commit
```

19. Reverse the changes made to control-ports and fab interfaces on Node0 that were previously adjusted in step 6. (You will configure the below on Node1 at step 21)

17. Before re-configuring control and re-configuring fab interfaces, enable interface-monitor which was disabled in step 4.

e.g.,

```
activate chassis cluster redundancy-group 1 interface-monitor
commit check
```

18. Commit the configuration on both nodes

```
commit
```

```
delete chassis cluster control-ports
set chassis cluster control-ports fpc 11 port 0
set chassis cluster control-ports fpc 23 port 0
set interfaces fab0 fabric-options member-interfaces xe-1/3/0
set interfaces fab1 fabric-options member-interfaces xe-13/3/0
commit check
commit and-quit
```

20. Make Node0 in halt status by “request system halt”

```
{primary:node0}
root@srx5k> request system halt
warning: This command will not halt the other routing-engine.
If planning to switch off power, use the both-routing-engines
option.
Halt the system ? [yes,no] (no) yes

*** FINAL System shutdown message from root@srx5k ***

System going down IMMEDIATELY

Shutdown NOW!
[pid 2193]

{primary:node0}
root@srx5k> failed to set the server tnp addresswaiting (max
60 seconds) for system process `vnlru_mem' to stop...done
Waiting (max 60 seconds) for system process `vnlru' to
stop...done
Waiting (max 60 seconds) for system process `bufdaemon' to
stop...done
Waiting (max 60 seconds) for system process `syncer' to
stop...
Syncing disks, vnodes remaining...3 3 1 1 1 1 1 1 0 0 0 0 0 0
done

syncing disks... All buffers synced.
Uptime: 1h25m0s
recorded reboot as normal shutdown

The operating system has halted.
Please press any key to reboot.
```

NOTE: DO NOT press any key before step 21 is completed.

NOTE: Make sure DO NOT commit until node0 is in halt status in step 20.

NOTE: Make sure node1 is primary for all RGs (show chassis cluster status).

21. When node0 console prints out “The operating system has halted.”, Reverse the changes made to control-ports and fab interfaces on Node1 that were previously adjusted in step 6. (Same configuration adjusted on Node 19 for Node 0)

```
delete chassis cluster control-ports
set chassis cluster control-ports fpc 11 port 0
set chassis cluster control-ports fpc 23 port 0
set interfaces fab0 fabric-options member-interfaces xe-1/3/0
set interfaces fab1 fabric-options member-interfaces xe-13/3/0
commit check
commit and-quit
```

22. Press any key to reboot Node0

<p>23. When Node0 completes bootup, verify the following: All FPCs and PICs are online (may take upto 15 minutes depending on the type and number of FPCs) Chassis cluster status should reflect nodes as Primary/Secondary Chassis cluster statistics should reflect increasing counts on control and fabric links</p> <pre> show chassis fpc pic-status (verify all slots and pics are "Online") show security flow session summary (verify both nodes reporting similar session counts) show chassis cluster status show chassis cluster statistics </pre> <p><i>NOTE: Alternatively, physically add cables or 'un-shut' connected device interfaces for Node0 after commit of flow settings.</i></p> <p>24. Enable all physical interfaces for transit traffic on Node0, which was disabled in step 11 and enable TCP syn-check/ sequence-check which were disabled in step 2.</p> <p>e.g.,</p> <pre> delete interfaces xe-0/0/0 disable delete interfaces xe-0/3/0 disable delete security flow tcp-session no-syn-check delete security flow tcp-session no-sequence-check commit check commit </pre> <p><i>NOTE: Alternatively, physically add cables or 'un-shut' connected device interfaces for Node0 after commit of flow settings</i></p> <p>25. Enable "preempt" and ip-monitoring if they were configured before for RG1+</p> <pre> activate chassis cluster redundancy-group 1 preempt activate chassis cluster redundancy-group 1 ip-monitoring commit and-quit </pre> <p>26. Verify chassis cluster priorities have returned to normal values</p> <pre> show chassis cluster status </pre> <p>27. Optional: Failover RG groups to Node0 (in case "preempt" is not configured, or is used with higher priority on Node1)</p> <pre> request chassis cluster failover redundancy-group 0 node 0 request chassis cluster failover redundancy-group 1 node 0 request chassis cluster failover reset redundancy-group 0 request chassis cluster failover reset redundancy-group 1 </pre>	<p>23. When Node0 completes bootup, verify the following: All FPCs and PICs are online (may take upto 15 minutes depending on the type and number of FPCs) Chassis cluster status should reflect nodes as Primary/Secondary Chassis cluster statistics should reflect increasing counts on control and fabric links</p> <pre> show chassis fpc pic-status (verify all slots and pics are "Online") show security flow session summary (verify both nodes reporting similar session counts) show chassis cluster status show chassis cluster statistics </pre>
---	--