Chassis Cluster Upgrade with Minimal Downtime
SRX Mid-Range (SRX1500, SRX4k)
Junos 21.1 and higher

Prerequisites:

- Minimum Junos release required for the procedure in this document:
  - SRX1500: 20.4R1
  - SRX4100/4200/4600: 21.1R1
  - If the currently installed Junos release is lower, please use the instructions available for earlier Junos releases.
- This procedure applies to all HA cabling variations. It can be used with single or dual control links or single or dual fabric links. Fabric links may also be using the dedicated HA Fabric ports or revenue ports.
- The following prerequisites may be performed outside of MW with no impact to traffic
- Console connections setup to both chassis cluster nodes are necessary to keep full visibility on the progress while rebooting and to keep control in case anything goes wrong during boot.
- 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-srxentedge-x86-64-23.2R1.13.tgz
- 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 redundancy-groups (RGs) to primary node
>request chassis cluster failover redundancy-group [x] node 0

Node0 Directions	Node1 Directions
1. Disable all physical interfaces used for transit traffic on Node1 (secondary node) e.g., set interfaces ge-7/0/4 disable set interfaces ge-7/0/5 disable *NOTE: make sure to use the correct interface name ge/xe/et as is used in your device configuration.* 2. Disable TCP SYN check and sequence check set security flow tcp-session no-syn-check set security flow tcp-session no-sequence-check	

3. Deactivate preempt for all RG1+

deactivate chassis cluster redundancy-group 1 preempt

4. Deactivate all interface-monitor and ip-monitoring

deactivate chassis cluster redundancy-group 1 interface-monitor
deactivate chassis cluster redundancy-group 1 ip-monitoring

5. Commit the configuration

commit

6. Disable control link and fab interfaces in the configuration.

set chassis cluster control-interface node 0 disable
set chassis cluster control-interface node 1 disable

set interfaces fab0 disable
set interfaces fab1 disable

7. Commit the configuration

commit and-quit

NOTE: error messages are expected and not an issue.

 # commit and-quit
...
error: error communicating with node0:
error: remote unlock-configuration failed on node1
error:
Connection to node1 has been broken
commit complete

NOTE: there is no need to commit on node1 at step 7.

NOTE: After step 7, the nodes are isolated from each other. Node1 is fully isolated, also from the network.

8. Before starting Node1 upgrade, make sure that the "active" configuration reflects the changes made on step 1-6 and that Node0 reports Node1 as lost.

e.g.,
{primary:node0}
root@srx1500> show configuration | display set | match disable

8. Before starting Node1 upgrade, make sure that the "active" configuration reflects the changes made on step 1-6 and that Node1 reports Node0 as lost.

e.g.,
{primary:node1}
root@srx1500> show configuration | display set | match disable

NOTE: In case the step 6 configuration is not committed successfully on Node1, configure and commit it on Node1 at this step.

```
{primary:node0}
root@srx1500> show chassis cluster status
…
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       None
node1  0        lost            n/a     n/a      n/a
```

NOTE: On SRX4600 you may see a "HW" Monitor-failure for RG1+ and a Major Alarm due to "FPC 0 Major Errors" while the control port is disabled. This is not traffic affecting and can be disregarded.

```
{primary:node1}
root@srx1500> show chassis cluster status
…
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       None
```

NOTE: The node1 status may be either "primary", "ineligible" or "disabled".

Start Node1 upgrade

9. Upgrade Junos OS on the Node1 and reboot

request system software add <install-package> no-validate reboot

NOTE: If for any reason you wish to stop the procedure after the reboot at this step and rollback, this can be done by performing a rollback of the Junos and then skip to step 16, with the instructions for node0 and node1 swapped.

10. After Node1 completes the boot process, verify the following before moving to the next step:
- Updated Junos OS
- All FPCs and PICs are online (may take upto 15 minutes)
- Node1 should be in primary state for all RGs, and reporting Node0 as 'lost'
- No major alarms being displayed

NOTE: On SRX4600 you may see a "HW" Monitor-failure for RG1+ and a Major Alarm due to "FPC 0 Major Errors" while the control port is disabled. This is not traffic affecting and can be disregarded.

show version
show chassis hardware
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:node1}[edit]
root@srx1500# run show chassis cluster status
…
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
```

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

```
e.g., (test commit)
   {primary:node0}[edit]
   root@srx1500# set interfaces reth0 description TEST

   {primary:node0}[edit]
   root@srx1500# commit
   node0:
   commit complete

   {primary:node0}[edit]
   root@srx1500# rollback 1
   load complete

   {primary:node0}[edit]
   root@srx1500# commit
   node0:
   commit complete
```

e.g.,
set interfaces ge-0/0/4 disable
set interfaces ge-0/0/5 disable

delete interfaces ge-7/0/4 disable
delete interfaces ge-7/0/5 disable
commit check

NOTE: make sure to use the correct interface name ge/xe/et as is used in your device configuration.

NOTE: Enable all physical interfaces of Node1 (slot 7) that were disabled in step 1. Disable the same interfaces of Node0 (slot 0).

NOTE: If there are any commit conflicts, they need to be resolved before moving to the next step.

12. Commit the configuration simultaneously on both nodes. This will cause all of the traffic to failover to the Node1.

commit

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

```
e.g., (test commit)
   {primary:node1}[edit]
   root@srx1500# set interfaces reth0 description TEST

   {primary:node1}[edit]
   root@srx1500# commit
   node1:
   commit complete

   {primary:node1}[edit]
   root@srx1500# rollback 1
   load complete

   {primary:node1}[edit]
   root@srx1500# commit
   node1:
   commit complete
```

e.g.,
set interfaces ge-0/0/4 disable
set interfaces ge-0/0/5 disable

delete interfaces ge-7/0/4 disable
delete interfaces ge-7/0/5 disable
commit check

NOTE: make sure to use the correct interface name ge/xe/et as is used in your device configuration.

NOTE: Enable all physical interfaces of Node1(slot 7) that were disabled in step 1. Disable the same interfaces of Node0 (slot 0).

NOTE: If there are any commit conflicts, they need to be resolved before moving to the next step.

12. Commit the configuration simultaneously on both nodes. This will cause all of the traffic to failover to the Node1

commit

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).

13. Verify traffic is passing through Node1

show security flow session summary
monitor interface traffic

Start Node0 upgrade

14. Upgrade Junos OS on the Node0

request system software add <install-package> no-validate reboot

15. 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

NOTE: On SRX4600 you may see a "HW" Monitor-failure for RG1+ and a Major Alarm due to "FPC 0 Major Errors" while the control port is disabled. This is not traffic affecting and can be disregarded.

show version
show chassis hardware
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@srx1500> show chassis cluster status
…
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
```

16. Before enabling control link, enable interface-monitor which was disabled in step 4.

e.g.,
activate chassis cluster redundancy-group 1 interface-monitor
commit check

17. Commit the configuration on both nodes

commit

18. Re-enable the control link interfaces on Node0 only

delete chassis cluster control-interface node 0 disable
delete chassis cluster control-interface node 1 disable

commit and-quit

16. Before enabling control link and fab interfaces, enable interface-monitor which was disabled in step 4.

e.g.,
activate chassis cluster redundancy-group 1 interface-monitor
commit check

17. Commit the configuration on both nodes

commit

19. Use the following operational CLI command to disable the control link on Node0 again until the next reboot.

request chassis cluster control-interface node 0 disable

NOTE: The control link will stay down at this point, but will come up after the next reboot of Node0. This allows the Node0 to join the cluster correctly after the reboot.

The fabric link will stay down at this point, since it is still disabled in the Node0 configuration. After the next reboot, when Node0 joins the cluster, the configuration will be synched from Node1 to Node0, which will enable the fabric link as well.

22. Reboot Node0.

request system reboot

```
{primary:node0}
root@srx1500> request system reboot
Reboot the system ? [yes,no] (no) yes
```

NOTE: Do NOT use "`request system power-off`", which would require someone to be on-site to enable the power again.
NOTE: SRX1500/4100/4200/4600 platforms do not support the command "`request system halt`"

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
- Both nodes showing same version

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
show version

20. Re-enable the control link interfaces on Node1 only

delete chassis cluster control-interface node 0 disable
delete chassis cluster control-interface node 1 disable

commit

21. Re-enable the fabric link interfaces on Node1 only

delete interfaces fab0 disable
delete interfaces fab1 disable

commit and-quit

23. When Node0 completes bootup, verify the following:

Chassis cluster statistics should reflect increasing counts on control and fabric links

show chassis cluster statistics

NOTE: If for some reason the configuration has synched from node0 to node1 instead of the other way around, leading to "set interfaces fabX disable" still being present in both nodes at this point, then configure the following again to bring up the fabric link:

```
delete interfaces fab0 disable
delete interfaces fab1 disable

commit and-quit
```

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.

e.g.,

```
delete interfaces ge-0/0/4 disable
delete interfaces ge-0/0/5 disable
delete security flow tcp-session no-syn-check
delete security flow tcp-session no-sequence-check
commit check
commit
```

25. Activate "preempt" and ip-monitoring if they were configured before for RG1+

```
activate chassis cluster redundancy-group 1 preempt
activate chassis cluster redundancy-group 1 ip-monitoring
commit and-quit
```

26. Verify chassis cluster priorities have returned to normal configured values

```
show chassis cluster status
```

27. Optional: Failover RG groups to Node0
(in case "preempt" is not configured, or is used with higher priority on node1)

```
request chassis cluster failover redundancy-group 1 node 0
request chassis cluster failover redundancy-group 0 node 0
request chassis cluster failover reset redundancy-group 1
request chassis cluster failover reset redundancy-group 0
```