Troubleshoot User Plane Node Reload Cases

Contents

Introduction

This document describes the process to identify different scenarios for User Plane (UP) reloads based on the symptoms to troubleshoot the issue.

Acronyms

RCM - Redundancy Configuration Manager

SSD - Show Support Details

UPF/UP - User Plane Function

VPP - Vector Packet processing

BFD - Bidirectional Forwarding Detection

Possible Scenarios Leading to UP Reloads

1. Configuration Push Issue between RCM and UPF
2. BFD, BGP, or Sx Flap
3. VPP Crash or Sessmgr Crashes

Required Logs for Symptom Identification

- RCM Logs (Controller/ConfigMgr)
- Syslogs
- Simple Network Management Protocol (SNMP) Traps
- RCM Command Outputs
- Show Support Details (SSDs) - If available, SSD before and after the reload

Analysis and Identification of Symptoms

Approach to Identify Symptoms of UP Reload Scenarios:

In a CUPS setup, UP reload scenarios frequently encounter challenges, demanding effective symptom identification and subsequent troubleshooting.

To initiate the process, examine the system uptime to pinpoint the exact time of the last UP restart. This information facilitates a focused analysis of RCM logs corresponding to the reload event.

Use this command to check for system uptime as follows:

```
******** show system uptime *******
Friday July 22 09:28:14 IST 2022
System uptime: 0D 0H 6M
```

> **Note**: Verify that RCM and UP timestamps are synchronized to the same timezone. If there is a discrepancy, make necessary correlations. For instance, if UP time is in IST and RCM time is in UTC, note that RCM time is consistently 5:30 hours behind UP time.

Verify if any crashes have occurred during the reload time. You can use this command to check for crash occurrences:

```
******** show crash list *******
Sunday January 23 12:12:14 IST 2022
=== ==================== ======== ========== =============== ======================
#           Time         Process  Card/CPU/         SW          HW_SER_NUM
                                     PID          VERSION       VPO / Crash Card
=== ==================== ======== ========== =============== ======================
1   2022-Jan-14+13:16:40 sessmgr  01/0/11287 21.25.5         NA
2   2022-Jan-19+20:51:01 sessmgr  01/0/16142 21.25.5         NA
3   2022-Jan-22+15:51:55 vpp      01/0/07307 21.25.5         NA
4   2022-Jan-22+15:52:08 sessmgr  01/0/27011 21.25.5         NA
5   2022-Jan-22+16:07:43 sessmgr  01/0/13528 21.25.5         NA
```

In this step, you need to examine whether any crashes have occurred, such as vpp/sessmgr crashes. If a vpp crash is detected, the UP reloads immediately due to the crash, prompting RCM to initiate a switchover to another UP.

If there is a consistent sequence of sessmgr crashes, it can potentially trigger a VPP crash, resulting in a reload of the UP.

Whenever you encounter such crashes, ensure to gather core files for vpp/sessmgr.

> **Note**: In the case of vpp, a minicore might be accessible instead of a full core file.

Action Plan: Once you obtain the core file or minicore, the next step is to perform core file debugging to pinpoint the root cause of the crash.

Analysis Based on UP syslogs/snmp

BFD Fluctuations or Failure in BFD Monitoring

The errors found in syslogs related to BFD monitoring failures are explained here.

These errors emerge when there is a BFD flap or packet loss between RCM and UP, especially in cases where ACI is involved in the connection between them.

Essentially, a timer is configured to monitor BFD packets. If, for any reason mentioned, this timer expires, it triggers a monitoring failure. This event prompts RCM to initiate a switchover.

```
Jan 22 15:51:55 <NODENAME> evlogd: [local-60sec55.823] [bfd 170500 error] [1/0/9345 <bfdlc:0> bfd_networ
Jan 22 15:51:55 <NODENAME> evlogd: [local-60sec55.856] [bfd 170500 error] [1/0/9345 <bfdlc:0> bfd_networ
Jan 22 15:51:55 <NODENAME> evlogd: [local-60sec55.859] [srp 84220 error] [1/0/10026 <vpnmgr:7> pnmgr_rcm
Jan 22 15:51:56 <NODENAME> evlogd: [local-60sec55.979] [srp 84220 error] [1/0/10026 <vpnmgr:7> pnmgr_rcm
```

To address this, it is important to conduct a comprehensive examination of the system and identify any potential issues that might have caused the BFD flap. If a problematic timestamp is pinpointed, coordinating with ACI is necessary to investigate whether there were any flaps or issues on their end corresponding to that timestamp.

BGP Flaps or BGP Monitoring Failure

BGP flaps or monitoring failures within UP can trigger a switchover initiated by RCM. These specific errors are characterized as described here.

```
Mar 21 09:10:37 <NODENAME> evlogd: [local-60sec37.482] [vpn 5572 info] [1/0/10038 <vpnmgr:7> pnmgr_rcm_b
Mar 21 09:10:37 <NODENAME> evlogd: [local-60sec37.482] [vpn 5572 info] [1/0/10038 <vpnmgr:7> pnmgr_rcm_b
Mar 21 09:10:37 <NODENAME> evlogd: [local-60sec37.482] [srp 84220 error] [1/0/10038 <vpnmgr:7> pnmgr_rcm
```

Possible factors contributing to BGP flaps and methods to identify them. SNMP traps could unveil errors that signal BGP flapping occurrences:

```
Wed Jan 18 10:30:03 2023 Internal trap notification 1289 (BGPPeerSessionIPv6Down)  vpn upf-in ipaddr abc
Wed Jan 18 10:30:09 2023 Internal trap notification 1288 (BGPPeerSessionIPv6Up)  vpn upf-in ipaddr abcd:
Wed Jan 18 10:30:19 2023 Internal trap notification 1289 (BGPPeerSessionIPv6Down)  vpn upf-in ipaddr abc
Wed Jan 18 10:30:03 2023 Internal trap notification 1289 (BGPPeerSessionIPv6Down)  vpn upf-in ipaddr abc
Wed Jan 18 10:30:09 2023 Internal trap notification 1288 (BGPPeerSessionIPv6Up)  vpn upf-in ipaddr abcd:
```

1. Initiate the process by identifying the context associated with the error that indicates BGP flaps, utilizing the context ID. With the context established, you can precisely determine the particular service involved and retrieve the corresponding IP details.

2. Employ the furnished information to collaboratively troubleshoot both the User Plane (UP) and ACI/Nexus components. This collaborative approach aims to precisely identify the root cause of the BGP flapping.
3. These instances of BGP flapping can lead to BGP monitoring failure, prompting the Redundancy Configuration Manager (RCM) to initiate a switchover.
4. BGP monitoring failure refers to the situation where the monitoring process for BGP encounters issues, resulting in an inability to ensure the continuous and expected functionality of the BGP connections.

In both RCM-based CUPS setups and ICSR-based CUPS setups, individual contexts are created within the UPs. For instance, in an RCM setup, the "rcm" context is established within the UP, while an ICSR setup involves the creation of the "srp" context. Here is a sample configuration for RCM-based CUPS:

<#root>

```
******** show rcm info *******
Thursday March 17 20:51:40 IST 2022
Redundancy Configuration Module:
-------------------------------------------------------------------------------
Context: rcm
Bind Address: <UPF IP binding with RCM controller>
Chassis State: Active
Session State: SockActive
Route-Modifier: 30
RCM Controller Address: <RCM controller IP>
RCM Controller Port: 9200
RCM Controller Connection State: Connected
Ready To Connect: Yes
Management IP Address: <UPF management IP>
Host ID: Active7
SSH IP Address: (Deactivated)
SSH IP Installation: Enabled

redundancy-configuration-module rcm
rcm controller-endpoint dest-ip-addr <Destination RCM controller IP> port 9200 upf-mgmt-ip-addr <UPF mar
bind address <UPF IP binding with RCM controller>
monitor bfd peer X.X.X.X
monitor bgp failure reload active
monitor bgp context GnS5S8-U X.X.X.X group 1
monitor bgp context GnS5S8-U X.X.X.X group 1
monitor bgp context GnS5S8-U abcd:defc:c:f::XXXX group 2
monitor bgp context GnS5S8-U defg:abcg:c:f::XXXX group 2
monitor bgp context SGi Z.Z.Z.Z group 3
monitor bgp context SGi G.G.G.G group 3
monitor bgp context SGi XXXX:YYYY:c:f::aaaa group 4
monitor bgp context SGi XXXX:YYYY:c:f::bbbb group 4
monitor bgp context Li XXXX:YYYY:c:f::cccc group 5
monitor bgp context Li XXXX:YYYY:c:f::dddd group 5
monitor sx context GnS5S8-U bind-address XXXX:YYYY:c:f::eeee peer-address XXXX:YYYY:c:f::ffff
#exit
```

Sample config for ICSR based CUPs without RCM

```
******** show srp info *******
Sunday April 23 04:39:49 JST 2023
Service Redundancy Protocol:
-------------------------------------------------------------------------------
Context: SRP
```

```
Local Address: <UP IP>
Chassis State: Active
Chassis Mode: Backup
Chassis Priority: 10
Local Tiebreaker: FA-02-1B-E8-C1-7E
Route-Modifier: 3

Peer Remote Address: <UP IP>
Peer State: Standby
Peer Mode: Primary
Peer Priority: 1
Peer Tiebreaker: FA-02-1B-13-31-D1
Peer Route-Modifier: 6
Last Hello Message received: Sun Apr 23 04:39:47 2023 (2 seconds ago)
Peer Configuration Validation: Complete
Last Peer Configuration Error: None
Last Peer Configuration Event: Sun Apr 23 04:21:10 2023 (1119 seconds ago)
Last Validate Switchover Status: None
Connection State: Connected

service-redundancy-protocol
monitor bfd context SRP <bfd peer IP> chassis-to-chassis
monitor bfd context SRP <bfd peer IP>  chassis-to-chassis
monitor bgp context SAEGW-U-1 <IP> group 1
monitor bgp context SAEGW-U-1 <IP> group 1
monitor bgp context SAEGW-U-1 <IP> group 2
monitor bgp context SAEGW-U-1 <IP> group 2
monitor bgp context SAEGW-U-1 <IP> group 3
monitor bgp context SAEGW-U-1 <IP> group 3
monitor bgp context SGI-1 <IP> group 4
monitor bgp context SGI-1 <IP> group 4
monitor system vpp delay-period 30
peer-ip-address <IP>
bind address <IP>
#exit
```

In both configurations, monitoring is implemented for BGP (similar to monitoring for BFD) within their respective contexts.

Each monitoring instance is assigned a unique group number, and distinct services are allocated separate group numbers. For instance, in the RCM context, "SGi" is associated with group number 3, "SGi IPv6" is linked to group number 4, and "Li" is connected with group number 5.

Using the provided configuration as a foundation, the RCM setup involves monitoring the specified BGP links within this context. Monitoring can encounter failure if any of these BGP links experience flapping or if there are difficulties in detecting the BGP link. In an ICSR setup, where RCM UP is absent, BGP link monitoring is conducted by SRP. This mechanism functions similarly to the explanation outlined in this point.

The primary objective is to oversee the links. When encountering these monitoring errors, the initial step is to ascertain the reasons behind the links not being monitored. Possible causes could include BGP flaps, configuration discrepancies in the IPs enlisted for monitoring versus the IPs specified in their respective contexts, or packet loss issues.

Sx Flaps

Similarly, as explained for BGP flaps, monitoring for Sx flaps between CP and UP is implemented. If an Sx

flap is detected, RCM initiates a switchover accordingly.

```
Errors for Sx flap which can be seen from snmp traps

Thu Apr 28 15:22:55 2022 Internal trap notification 1382 (SxPathFailure) Context Name:gwctx, Service Nam
```

Analysis Based on RCM Logs and Command Outputs

RCM Controller Logs:

<#root>

Monitoring failure for BFD

```
{"log":"2022/11/12 13:33:31.138  [ERROR] [red.go:2144] [rcm_ctrl.control.main] [handleUpfActiveToDownAct
```

Monitoring failure for BGP

```
{"log":"2022/11/12 15:34:27.644  [ERROR] [red.go:2144] [rcm_ctrl.control.main] [handleUpfActiveToDownAct
```

Monitoring failure for Sx

```
{"log":"2022/11/12 15:34:46.763  [ERROR] [red.go:2144] [rcm_ctrl.control.main] [handleUpfActiveToDownAct
```

RCM command outputs:

```
rcm show-status
(to check RCM in Master or Backup state)

rcm show-statistics configmgr
(to check number of UPs connected to this configmgr and current stat of about which are the active UPs a

rcm show-statistics controller
(to check number of UPs connected to this controller and current stat of about which are the active UPs

rcm show-statistics switchover
rcm show-statistics switchover-verbose
(to check which UP got switchovered to which UP and at what time and with what reason)
```

Examples of the command output:

```
root@Nodename:
[unknown] ram# ram show-status
message :
{"statusâ€�: â€œMASTER"}


[unknown] rcm# rcm show-statistics switchover
message :
{
  "stats_history": [
    {
      "status": "Success",
      "started": "Mar 21 03:40:37.480",
      "ended": "Mar 21 03:40:41.659",
      "switchoverreason": "BGP Failure",
      "source_endpoint": "X.X.X.X",
      "destination_endpoint": "Y.Y.Y.Y"
    }
  ],

  "num_switchover": 1

}
```

It is important to obtain the controller logs and carefully review them for any switchover scenarios, as previously discussed. This analysis aims to ensure that the switchover process was executed seamlessly and without any issues.

<#root>

```
{"log":"2022/05/10 00:30:48.553  [INFO] [events.go:87] [rcm_ctrl_ep.events.bfdmgr] eventsDbSetCallBack:


--------------------Indication of active UP bfd went down



{"log":"2022/05/10 00:30:48.553  [DEBUG] [control.go:2920] [rcm_ctrl.control.main] [stateMachine]: Recei
{"log":"\n","stream":"stdout","time":"2022-05-10T00:30:48.553661415Z"}
{"log":"2022/05/10 00:30:48.553  [INFO] [red.go:2353] [rcm_ctrl.control.main] [upfHandlUpfAction]: State
{"log":"2022/05/10 00:30:48.553  [ERROR] [red.go:2103] [rcm_ctrl.control.main] [handleUpfActiveToDownAct
{"log":"2022/11/12 13:33:27.759  [ERROR] [red.go:2144] [rcm_ctrl.control.main] [handleUpfActiveToDownAct

---------- Indication of BFD/BGD timer expired and there is a monitoring failure



{"log":"2022/05/10 00:30:48.553  [WARN] [red.go:2256] [rcm_ctrl.control.main] [handleUpfActiveToDownActi
----------

Indication of switchover initiated by RCM



{"log":"2022/05/10 00:32:03.555  [DEBUG] [control.go:3533] [rcm_ctrl.control.main] [snmpThread]: SNMP tr
{"log":"2022/05/10 00:32:03.603  [DEBUG] [control.go:1885] [rcm_ctrl.control.main] [handleUpfStateMsg]:
{"log":"2022/05/10 00:32:03.603  [DEBUG] [control.go:2048] [rcm_ctrl.control.main] [handleUpfStateMsg]:
----------

Indication of switchover completed and other UP became Active
```

```
{"log":"2022/05/10 00:32:03.646  [INFO] [control.go:1054] [rcm_ctrl.control.main] [handleUpfActiveAckMsg
--------------
```

Traffic routed towards other Active UP

```
{"log":"2022/05/10 00:32:53.861  [INFO] [red.go:859] [rcm_ctrl.control.main] [handleUpfSetStandby]: Assi
{"log":"2022/05/10 00:32:53.861  [INFO] [red.go:1681] [rcm_ctrl.control.main] [sendStateToUpf]:send stat
{"log":"2022/05/10 00:32:53.890  [INFO] [red.go:1176] [rcm_ctrl.control.main] [handleUpfNotifyMgrs]: Rec
```

---------- Switchovered UP became Standby

Configuration Push Issues between RCM and UPF

- During a switchover from one UP to another UP initiated by RCM, the necessary configuration is pushed by RCM. To ensure this configuration is successfully applied, RCM sets a timer to complete the process.
- Once the configuration is pushed and stored in the path of the UP, the UP executes the configuration within the specified time frame defined by RCM.
- Once the UP completes the execution of the configuration, it sends a signal to RCM. This signal is indicated by an event log entry in the syslogs, confirming the successful completion of the configuration push.

```
Nov 13 12:01:09 <NODENAME> evlogd: [local-60sec9.041] [cli 30000 debug] [1/0/10935 <cli:1010935> clipars
Nov 13 12:01:09 <NODENAME> evlogd: [local-60sec9.041] [cli 30000 debug] [1/0/10935 <cli:1010935> clipars
```

- When the event logs show up in syslogs, it indicates that the configuration process is complete. However, if this event is absent after post-configuration execution, it implies that a configuration issue might be causing a delay in completion. In such cases, a timer can expire, leading to UP reloads. RCM initiates the reload when it does not receive a response from UP. This cycle continues until the configuration is successfully completed.
- As a temporary solution, you can manually execute a specific command to halt the reload process temporarily. Nonetheless, it is crucial to pinpoint the root cause of the configuration problem to resolve it permanently.

```
rcm-config-push-complete end-of-config
```

Factors Leading to Configuration Problems

Problematic Configuration CLIs

Identify troublesome CLIs within the pushed configuration file, which can be determined from RCM ConfigMgr logs.

SFTP Issues

SFTP-related problems can occur when RCM tries to send configuration but faces difficulties in establishing a connection with the UP. These challenges could stem from password complications or other factors impacting SFTP operations.

Reviewing ConfigMgr logs allows for monitoring of SFTP status and identifying configuration errors. Here is a sample representation of typical error instances.

SFTP logs in RCM ConfigMgr logs appear as:

```
{"log":"2022/11/12 23:53:09.066 rcm-configmgr [DEBUG] [sshclient.go:395] [rcm_grpc_ep.msg-process.Int] I
{"log":"2022/11/12 23:53:09.066 rcm-configmgr [DEBUG] [sftpClient.go:26] [rcm_grpc_ep.grpc.Int] Connetin
{"log":"2022/11/12 23:53:09.203 rcm-configmgr [DEBUG] [sftpClient.go:58] [rcm_grpc_ep.grpc.Int] Successf
{"log":"2022/11/12 23:53:09.211 rcm-configmgr [DEBUG] [sftpClient.go:66] [rcm_grpc_ep.grpc.Int] Total by
```

Password expiration during SFTP observed in UP syslogs:

```
2022-May-16+17:45:02.834 [cli 30005 info] [1/0/14263 <cli:1014263> _commands_cli.c:1474] [software inter
2022-May-16+17:45:02.834 [cli 30024 error] [1/0/14263 <cli:1014263> cli.c:1657] [software internal syste
2022-May-16+17:45:02.834 [cli 30087 info] [1/0/14263 <cli:1014263> cli.c:1352] [software internal system
2022-May-16+17:45:02.594 [cli 30004 info] [1/0/14263 <cli:1014263> cli_sess.c:164] [software internal sy
2022-May-16+17:45:02.537 [cli 30028 debug] [1/0/9816 <vpnmgr:1> luser_auth.c:1598] [context: local, cont
```

If SFTP problems arise from passwords, consider generating a new password or extending the password expiration period.

If password issues are ruled out, examine the number of concurrent SFTP sessions, as an excessive number of sessions can lead to SFTP disruptions.