Troubleshoot RPD DOCSIS Throughput Performance Issues

Contents

Introduction

Prerequisites

- Requirements
- Components Used

Background Information

Troubleshoot

- Late MAP Messages
  - Potential Cause 1. CIN Delay, Latency, Jitter
  - Potential Cause 2. Software Bug
- Upstream Latency
- Out of Order Layer 2 Tunneling Protocol (L2TP) Packets
  - Potential Cause 1. Load balance
  - Potential Cause 2. Packet Drops
- PTP Loss or Unlocks Periodically
- Congested CIN
  - Potential Cause 1. QoS
  - Potential Cause 2. Delayed Best Effort Traffic

Related Information

Introduction

This document describes how to troubleshoot Cisco Remote PHY Device (RPD) performance problems.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- RPD
- Cisco converged Broadband Router (cBR)-8
- Data Over Cable Service Interface Specification (DOCSIS)

Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

The scenario considered in this article involves an RPD provisioned by the Cisco cBR-8 as Converged Cable Access Platform (CCAP). Precision Time Protocol (PTP) is used to synchronize an external primary clock with the cBR-8 and RPD which act as secondary. For more information on how the PTP design in this environment, you can refer to PTP Design Recommendations For R-PHY Networks.

This is not a comprehensive list of steps in order to troubleshoot performance issues with RPD, although it is a good start in order to isolate the problem.

Troubleshoot

If you observe a performance degradation with RPD deployment, and you wish to perform an initial troubleshoot, it is not be clear as to where to begin.

This section describes some common problems that can be the possible cause for the RPDs performance issues.

Late MAP Messages

A late upstream bandwidth allocation map (MAP) message condition, occurs when a modem receives a MAP message at a point in time, after the time slots described in the message have already occurred. The modem is unable to use this MAP message, so is unable to send any traffic on the assigned grants.

A few late MAPs can cause reduced upstream traffic rates, as well as reduced downstream TCP traffic rates as upstream ACKs are delayed. If there are enough late MAPs, modems are unable to perform station maintenance and go offline.

Another symptom can be packet drops when you do a **ping docsis <MAC_ADDR>** from the cBR-8 to a modem connected to an RPD.

With Remote PHY (R-PHY), the cBR-8 sends MAP messages to the modems in a Downstream External PHY Interface (DEPI) tunnel and to the RPD in an Upstream External PHY Interface (UEPI) tunnel. These messages have higher Quality of Service (QoS) mark with a Differentiated Services Code Point (DSCP) value of 46 (express forwarding - EF).

If a MAP message destined for the RPD gets dropped in the CIN, the RPD is not able to use those minislots and counts them as “unmapped”. If the MAP message arrives late at the RPD, it initially counts the minislots as unmapped and then after it receives the late MAP, it increments the late minislots count.

Early MAPs are also possible but usually only happen when the PTP clock is off in either the cBR-8 or the RPD.

Overlap MAPs can happen when MAP messages come out of sequence but withjust 2 ms frequency, this is not usually a problem. The actual number of minislots in a MAP message is based on the minislot configuration for each upstream channel. If an upstream uses two ticks per minislot (popular for 6.4 MHz SC-QAM), a 2 ms MAP has 160 minislots.

In order to check if on the RPD you receive late MAP messages, perform these commands to access the RPD console. Then, run the command **show upstream map counter <rf port> <channel>** multiple times and check if the counter “Discarded minislots (late maps)” increases:

```
cbr8# ssh <RPD_IP_ADDR> -l admin
R-PHY>enable
R-PHY#

show upstream map counter 0 0

Map Processor Counters
==============================================
Mapped minislots                  :     553309
Discarded minislots (chan disable):          0
Discarded minislots (overlap maps):          0
Discarded minislots (early maps)  :          0
Discarded minislots (late maps)   :          0 <= check if the counter increases
Unmapped minislots                :          0
Last mapped minislot              :   21900956
```



Note: Every time you run the command **show upstream map counter**, all the counters are reset to 0 but Last mapped minislot



Tip: From RPD version 6.x, you can run the **show tech-support** command, which collects multiple occurrences of **show upstream map counter** and other commands, therefore useful for counters comparison.

If you run RPD software version 5.x or lower, you can run the **show tech** command with the use of the script available here: Capture show tech on rpd or limited command on both RPD, supervisor.

The linked page contains instructions on how to install the script and usage examples, at the end of which you can find the file **Script-Readme.tar** available for download. This file contains the **sh_tech_rpd.tcl** script and the **sh_tech_rpd-README.txt** file with the instructions and usage examples.

The script has an option (-c) in order to collect an additional set of commands listed in a text file, both commands to be issued on the RPD itself and on the cBR-8 supervisor are accepted (all procedures explained in the link previously mentioned and the readme file).

This feature makes the use of this script, interestingly, also in RPD versions that include the **show tech-support** command.

Potential Cause 1. CIN Delay, Latency, Jitter

The Converged Interconnect Network (CIN) that links the CCAP core and RPDs can introduce delays that must be accounted for in the MAP advance timer. If there is a change in the CIN, like for example another router was added, you possibly have introduced a higher delay.

The MAP advance timer is used by the CCAP to prevent late MAP messages. This timer is based in microseconds(µs), and can be statically configured per cable interface by the operator, or dynamically calculated by the cBR-8.

The dynamic value, is the sum of the downstream time interleaf (680 µs with SC-QAM with 256-QAM), modem MAP processing delay (600 µs), CCAP internal network delay (300 µs), MAP advance safety value (1000 µs by default), and maximum modem time offset (based on most distant modem).

With R-PHY, the CCAP internal delay is now replaced by a network delay, which defaults to 500 µs. In consideration of the CIN design, this value can be larger than the default parameter and can change over

time.

The MAP advance values for an upstream can be displayed with this command:

<#root>

```
cbr8#
```

show controllers upstream-Cable 2/0/5 us-channel 0 cdm-ump

```
<output omitted>
nom_map_adv_usecs 2899, max_map_adv_usecs 4080 mtn_map_adv 8080
map_adv_alg 1 dyn_map_adv_safety 1000 max_plant_delay 1800
cm_map_proc 600 intlv_delay 680 network_delay 500 max_tmoff 119
<output omitted>
```

MAPadvance = map_adv_safety (1000) + cm_map_proc (600) + intlv_delay (680) + network_delay (500)+ max_tmoff (119) = 2899 µs.

If the distance between the cBR-8 and RPD combined with the CIN devices delays exceeds the network delay default value of 500 µs, late MAP messages are possible.

There are two methods to deal with the default network delay parameter when this represents a problem, and both are set per RPD on the cBR-8:

- Statically configure the delay.
- Set the cBR-8 to measure and adjust the delay periodically.

The network delay can be statically configured per RPD on the cBR-8 as shown here:

```
cbr8#conf t
cbr8(config)#cable rpd <name>
cbr8(config-rpd)#core-interface <interface_name>
cbr8(config-rpd-core)#network-delay static <CIN_delay_in_us>
```

For dynamic network delay, the cBR-8 relies on a latency measurement feature called DEPI Latency Measurement (DLM), which determines the one-way delay in the downstream path.

The cBR-8 sends out a DLM packet with its time stamp, then the RPD marks its time stamp on the DLM packet when received, and forward it back to the cBR-8.

Note that Cisco supports the required option where the RPD marks the packet closest to its ingress interface, not to the egress.

The cBR-8 takes the average of the last 10 DLM values and use it as the network delay value in the MAP advance calculation. The time stamps from both the cBR-8 and the RPD are based on the PTP reference clocks.

Warning: If PTP is unstable, so are the DLM values and ultimately the MAP advance timer.

By default DLM is disabled, and it can be enabled with the **network-delay dlm <seconds>** command as

shown. Once enabled, a DLM packet is sent to the RPD periodically in accordance with the configured value.

There is also a **measure-only** option that can be added, which does just measure the CIN delay without adjustment of the network delay value.

It is recommended to enable DLM at a minimum in the measure-only parameter, in order to monitor the CIN delay.

<#root>

```
cbr8#conf t
cbr8(config)#cable rpd <name>
cbr8(config-rpd)#core-interface <interface_name>
cbr8(config-rpd-core)#
```

network-delay dlm

<interval_in_seconds> [

measure-only

]

Usage:

cbr8#

show cable rpd a0f8.496f.eee2 dlm

```
DEPI Latency Measurement (ticks) for a0f8.496f.eee2
  Last Average DLM:              481
  Average DLM (last 10 samples): 452
  Max DLM since system on:       2436
  Min DLM since system on:       342
  Sample #     Latency (usecs)
  x------------x------------
  0            52
  1            41
  2            48
  3            41
  4            41
  5            44
  6            40
  7            45
  8            44
  9            41
```

More information on this feature can be found here; DEPI Latency Measurement

The MAP advance safety can also be changed manually in the cable interface configuration (default values are 1000 µs for the safety factor and 18000 µs for the max map advance):

```
cbr8#conf t
cbr8(config)#interface Cable1/0/0
cbr8(config-if)# cable map-advance dynamic 1000 18000
```

```
OR (if a mac-domain profile is used)

cbr8#conf t
cbr8(config)# cable profile mac-domain RPD
cbr8(config-profile-md)# cable map-advance dynamic 1000 18000
```

⚠ **Caution**: Very short CIN delays can also cause late MAP messages

There have been issues observed with dropped upstream DOCSIS traffic when the MAP advance timer is under 2500 µs.

Some modems can take longer to process MAP messages, and that extra delay can cause late MAP messages for those modems (the RPD possibly does not show late MAP counts if it was able to get the message in time).

A low MAP advance timer is possible with very low DLM values, or with low manual network delay or MAP advance safety configuration. Network delay values in the MAP advance calculation can be as low as 30 µs (even if the DLM average is lower).

It is recommended to either use DLM “measure-only” option or increase the safety factor for dynamic MAP advance until the MAP advance timer is over 2500 µs.

Potential Cause 2. Software Bug

In order to verify if a software bug causes a synchronization failure, it is recommended to open a service request with Cisco to further investigate the specific case.

The solution in case you hit a software defect is usually a software upgrade to one of the releases which contains the fix. Since there is a compatibility correlation between cBR-8 and RPD software releases, it is important to choose the right version for both devices.

In order to help choose the right Cisco IOS® XE for every RPD software, you can find the software version compatibilities between cBR-8 and RPD in the Cisco Remothe PHY Install and Upgrade Guides.

In this table you can find a summary of the software version compatibility between cBR-8 and RPD, with the available software versions at the time of writing:

Version Compatibility Between Cisco cBR-8 and Cisco RPD	
Cisco cBR-8 Release Version	Compatible RPD Release Version
Cisco IOS® XE Everest 16.6.x	Cisco 1x2 RPD Software 2.x
Cisco IOS® XE Fuji 16.7.x	Cisco 1x2 RPD Software 3.x
Cisco IOS® XE Fuji 16.8.x	Cisco 1x2 RPD Software 4.x

Version Compatibility Between Cisco cBR-8 and Cisco RPD	
Cisco cBR-8 Release Version	Compatible RPD Release Version
Cisco IOS® XE Fuji 16.9.x	Cisco 1x2 RPD Software 5.x
Cisco IOS® XE Gibraltar 16.10.1c	Cisco 1x2 RPD Software 6.1, 6.2, 6.3
Cisco IOS® XE Gibraltar 16.10.1d	Cisco 1x2 RPD Software 6.4, 6.5, 6.7
Cisco IOS® XE Gibraltar 16.10.1f	Cisco 1x2 RPD Software 6.6, 6.7
Cisco IOS® XE Gibraltar 16.10.1g	Cisco 1x2 RPD Software 7.1, 7.2, 7.3, 7.4.x, 7.5
Cisco IOS® XE Gibraltar 16.12.1	Cisco 1x2 RPD Software 7.1, 7.2, 7.3, 7.4.x, 7.5
Cisco IOS® XE Gibraltar 16.12.1w	Cisco 1x2 RPD Software 7.1, 7.2, 7.3, 7.4.x, 7.5
Cisco IOS® XE Gibraltar 16.12.1x	Cisco 1x2 RPD Software 7.6, 7.7
Cisco IOS® XE Gibraltar 16.12.1y	Cisco 1x2 RPD Software 7.8, 7.8.1, 8.2
Cisco IOS® XE Gibraltar 16.12.1z	Cisco 1x2 RPD Software 8.3, 8.4, 8.5
Cisco IOS® XE Gibraltar 17.2.1	Cisco 1x2 RPD Software 8.1, 8.2, 8.3, 8.4, 8.5

Upstream Latency

As discussed in the previous section, long CIN delays can cause late MAP message issues, and can be addressed with the MAP advance timer increase. This in turn creates a longer request-grant delay, which leads to increased upstream latency.

Since steady upstream traffic flows use piggy-back requests, upstream traffic speed test can appear normal, and also voice flows with Unsolicited Grant Service (UGS) does not get impacted, as no requests are needed.

However, downstream TCP traffic speeds can be reduced due to lack of timely upstream ACKs. Although it is possible to address processing and queuing delays on the CIN, it is not likely to make signals travel faster over a given distance.

Cisco developed DOCSIS Predictive Scheduling (DPS) to reduce request-grant delay in R-PHY applications with longer CIN delays. DPS proactively provides grants to modems based on historical usage, to minimize

request-grant delay.

DPS is a pre-standard scheduling method, similar to Proactive Grant Service (PGS) described in the recent Low Latency DOCSIS (LLD) specification. However, DPS can be enabled per interface and is applied to all best effort upstream service flows. PGS is applied to traffic as a service flow type, so it requires changes to the modem configuration file.

DPS can be enabled with the interface command: cbr8(config-if)#**cable upstream dps**

Although DPS has been available since R-PHY support was added to the cBR-8, it is not an officially supported feature at this time. Nevertheless, DPS can be effective to resolve slow TCP downstream throughput associated with delayed ACKs.

Out of Order Layer 2 Tunneling Protocol (L2TP) Packets

On the RPD console, type this command multiple times in order to verify if the counters “SeqErr-pkts” and “SeqErr-sum-pkts” are positive and increase, which is an indication of L2TP out of order packets:

```
<#root>

R-PHY#

show downstream channel counter dpmi

Chan Flow_id SessionId(dec/hex)      Octs       Sum-octs   SeqErr-pkts SeqErr-sum-pkts
0    0       4390912    / 00430000   328        22770      0           1
0    1       4390912    / 00430000   25074      1179672    0           1
0    2       4390912    / 00430000   6022168    271459412  0           1
0    3       4390912    / 00430000   0          0          0           0
```

Potential Cause 1. Load balance

In some particular conditions, like links congestion in the CIN for example, load balance can contribute to the problem of packets received out of order at the destination.

If you have the possibility, in order to check if load balance triggers this problem, you can test to enforce a single path where the load balance is configured. If this resolves the out-of-order packets problem you have the confirmation of the trigger, and can further investigate the root cause in your network.

Potential Cause 2. Packet Drops

1. Check for any increased errors and drops on the cBR-8 counters on the DPIC card interface where the RPD is connected, with the use of the **show interface** command as shown here.

```
<#root>

cbr8#

sh run | s cable rpd SHELF-RPD0

cable rpd SHELF-RPD0
 description SHELF-RPD0
 identifier a0f8.496f.eee2
[…]
```

```
 core-interface Te6/1/2
[…]
cbr8#
```

show interface Te6/1/2

```
TenGigabitEthernet6/1/2 is up, line protocol is up
  Hardware is CBR-DPIC-8X10G, address is cc8e.7168.a27e (bia cc8e.7168.a27e)
  Internet address is 10.27.62.1/24
  MTU 1500 bytes, BW 10000000 Kbit/sec, DLY 10 usec,
     reliability 255/255, txload 90/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full Duplex, 10000Mbps, link type is force-up, media type is SFP_PLUS_10G_SR
  output flow-control is on, input flow-control is on
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:01, output 00:00:05, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/375/0/22 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 1002000 bits/sec, 410 packets/sec
  5 minute output rate 3535163000 bits/sec, 507528 packets/sec
     88132313 packets input, 26831201592 bytes, 0 no buffer
     Received 0 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 watchdog, 229326 multicast, 0 pause input
     179791508347 packets output, 164674615424484 bytes, 0 underruns
     0 output errors, 0 collisions, 1 interface resets
     13896 unknown protocol drops
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier, 0 pause output
     0 output buffer failures, 0 output buffers swapped out
```

2. Check on the RPD side if there are errors, dropped, and out of order packets on the interfaces and downstream counters.

<#root>

```
R-PHY#
```

show interface info

```
eth0      Link encap:Ethernet  HWaddr A0:F8:49:6F:EE:E4
          inet addr:192.168.1.1  Bcast:192.168.1.255  Mask:255.255.255.0
          inet6 addr: fe80::a2f8:49ff:fe6f:eee4/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:303 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:0 (0.0 B)  TX bytes:44034 (43.0 KiB)
          Memory:1ae2000-1ae2fff
vbh0      Link encap:Ethernet  HWaddr A0:F8:49:6F:EE:E2
          inet addr:10.7.62.7  Bcast:10.7.62.255  Mask:255.255.255.0
          inet6 addr: fe80::a2f8:49ff:fe6f:eee2/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:1174200 errors:0 dropped:0 overruns:0 frame:0
          TX packets:593404 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
```

```
          RX bytes:90888838 (86.6 MiB)  TX bytes:52749774 (50.3 MiB)
vbh1      Link encap:Ethernet  HWaddr A0:F8:49:6F:EE:E3
          inet6 addr: fe80::a2f8:49ff:fe6f:eee3/64 Scope:Link
          UP BROADCAST MULTICAST  MTU:1500  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:24 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:0 (0.0 B)  TX bytes:2438 (2.3 KiB)

R-PHY#
```

show downstream channel counter

```
------------------- Packets counter in TPMI -------------------

Level     Rx-pkts    Rx-sum-pkts
Node Rcv  4673022    2108792873
Depi Pkt  1696       774495

Port Chan SessionId(dec/hex)    Rx-pkts    Rx-sum-pkts
DS_0 0    4390912    /0x00430000 49032      22125274
DS_0 1    4390913    /0x00430001 49025      22116541
[…]
US_0 0    13893632   /0x00D40000 12193      5502543
US_0 1    13893633   /0x00D40001 12193      5501739
[…]

Port      Rx-pkts    Rx-sum-pkts Drop-pkts  Drop-sum-pkts
DS_0      3095440    1396529318  0          0
US_0      49215      22207507    0          0
US_1      0          4679        0          0

------------------- Packets counter in DPMI -------------------

Field        Pkts       Sum-pkts
Dpmi Ingress 12275995   1231753344
Pkt Delete   0          0
Data Len Err 0          0

Chan Flow_id SessionId(dec/hex)      Octs       Sum-octs   SeqErr-pkts SeqErr-sum-pkts
0    0       4390912    / 0x00430000   75         130496     0           1
0    1       4390912    / 0x00430000   15657      7208826    0           1
0    2       4390912    / 0x00430000   3181212    1431951867 0           1
0    3       4390912    / 0x00430000   0          0          0           0
[…]

------------------- Packets counter in DPS -------------------

Chan Tx-packets Tx-octets  Drop-pkts  Tx-sum-pkts Tx-sum-octs Drop-sum-pkts
0    50316      3273636    0          22126173    1439340721  0
1    50311      3272896    0          22117442    1438506648  0
2    50311      3272640    0          22121500    1438772715  0
3    50309      3272640    0          22122038    1438807607  0
[…]
```

3. Check the downstream InterLaken counters multiple times in order to see if there are errors and if the counters increase. In order to do that, you need to enter the line card console interface as shown here.

```
cbr8#request platform software console attach 6/0
#
# Connecting to the CLC console on 6/0.
# Enter Control-C to exit the console connection.
#
Slot-6-0>enable
Slot-6-0#
Slot-6-0#test jib4ds show ilkstat ?
  <0-3>  ILK Link (0-BaseStar0, 1-BaseStar1, 2-Cpu0, 3-Cpu1)

Slot-6-0#
```

`test jib4ds show ilkstat 0`

```
Send Show-ilkstat IPC to CDMAN...Wait for output

Slot-6-0#
Jib4DS InterLaken Stats for BaseStar 0:
                    RX-Packets             RX-Bytes            TX-Packets             TX-Bytes
HUB Stats:         10425879607       14415939325556              75237425           8249683443
Chan    0:             4714787            360160866                109750             36594720
Chan    1:         10254597081       14397444921888                     0                    0
Chan    3:               63828             17214818                     0                    0
Chan    5:           166503829          18117169182              75127675           8213088761
PRBS  Err:                   0                    0                     0                    0
CRC32 Err:                   0                    0                     0                    0
CRC24 Err:                          0
Test-pattern-err:                   0

ILK Error log: ptr 0
Idx      Err1     Err2      Rst     Gtx0     Gtx1     Gtx2     Gtx3

Slot-6-0#
```

4. Take a modem connected to this RPD (DS channels bonded only), and send packets (for example, pings) to it, in order to check if sent packets match the injected downstream flow on the JIB module counters. Ensure that DS JIB sent all DS data packets for DEPI frame on the line card console. In this output, you can see how to see the packet sequence number from a modem service flow output. This sequence number increases for each data packet sent.

<#root>

```
Slot-6-0#
```

`show cable modem 2cab.a40c.5ac0 service-flow verbose | i DS HW Flow`

```
DS HW Flow Index            : 12473
Slot-6-0#
```

`test jib4ds show flow 12473`

```
Send Show-FLOW IPC to CDMAN flow 12473 seg 6...Wait for output

Slot-6-0#
Jib4DS Show Flow: [Bufsz 4400]: HW Flow id:12473 [0x30b9] for segment 0
Valid        : TRUE
DSID         :            3 [       0x3]
Priority     :            0
Bonding Group:           62 [      0x3e]
Channel      :        65535 [    0xffff]
```

```
DS-EH        :            3 [       0x3]
Data Prof 1  :            0 [         0]
Data Prof 2  :            0 [         0]
No Sniff Enabled.

Slot-6-0#
```

test jib4ds show dsid 3

```
Send Show-DSID 3 10 IPC to CDMAN...Wait for output

Slot-6-0#

Jib4DS DSID entry for DSID 3 [Bufsz 4400]:
SCC Bit         = 0x0
Sequence Number = 8
```

Send some pings to this modem from the cBR-8 command line, on another window:

```
cbr8#ping 10.0.0.3 rep 100
Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 10.0.0.3, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (100/100), round-trip min/avg/max = 4/7/27 ms
cbr8#
```

Check the delta after the test:

<#root>

```
Slot-6-0#
```

test jib4ds show dsid 3

```
Send Show-DSID 3 10 IPC to CDMAN...Wait for output

Slot-6-0#

Jib4DS DSID entry for DSID 3 [Bufsz 4400]:
SCC Bit         = 0x0
Sequence Number = 108
```

Calculate the delta after the test: the counter is 16 bit unsigned, so if the counter rolls over, the delta needs to be calculated with this formula:

```
(Initial Sequence Number + Number of Packets Sent) % 65536
```

Example:

```
Initial Sequence Number = 50967

Final Sequence Number = 2391

Packets sent: 1000000



(50967+1000000)%65536 = 2391 <== Good, no packet was dropped before DEPI frame.
```

As a consequence of the nature of the drops, the problem can be in the CIN (for example, bottleneck links, collisions, CRC errors) which need to be further investigated in the CIN network between the cBR-8 and RPD. If drops are observed in points 3 and 4 instead, it is recommended to engage Cisco for further investigation on the cBR-8.

PTP Loss or Unlocks Periodically

As you probably know, PTP is essential for normal RPD operations. Therefore, PTP packets must have high priority in QoS, and PTP packet drops are not a good sign.

On the RPD console, you can show the PTP statistics, and verify that the counters “DELAY REQUEST” and “DELAY RESPONSE” are closely matched. If you see a big gap instead, it can be an indicator of PTP drops in the network:

```
<#root>

R-PHY#

show ptp clock 0 statistics

 AprState    4  :
           2@0-00:06:25.877          1@0-00:06:16.234          0@0-00:03:42.629
           4@0-00:03:23.428
ClockState  5  :
           5@0-00:07:02.932          4@0-00:06:59.145          3@0-00:06:55.657
           2@0-00:06:26.657          1@0-00:06:25.834
BstPktStrm  1  :
           0@0-00:03:21.014
SetTime     1  :
  1000000000@0-00:03:24.776
StepTime    1  :
  -560112697@0-00:05:39.401
AdjustTime  44 :
          -8@0-00:52:03.776         -5@0-00:51:02.776          4@0-00:50:01.776
          -6@0-00:49:00.776         11@0-00:47:59.776          1@0-00:45:57.776
           5@0-00:44:56.776         -7@0-00:43:55.776        -22@0-00:42:54.776
streamId  msgType           rx          rxProcessed     lost        tx
0         SYNC              47479       47473           0           0
0         DELAY REQUEST     0           0               0           47473
0         P-DELAY REQUEST   0           0               0           0
0         P-DELAY RESPONSE  0           0               0           0
0         FOLLOW UP         0           0               0           0
0         DELAY RESPONSE    47473       47473           0           0
0         P-DELAY FOLLOWUP  0           0               0           0
0         ANNOUNCE          2974        2974            0           0
0         SIGNALING         34          34              0           32
```

```
0          MANAGEMENT          0          0          0          0
  TOTAL                       97960      97954      0          47505
```

Note: On cBR-8, PTP has the highest priority for clock, which means that once it is configured, it is used even for RF line cards. Therefore, an unreliable source would cause issues across the whole chassis.

For more information about the PTP clock configuration and troubleshoot, you can read the article PTP Design Recommendations For R-PHY Networks.

Congested CIN

The CIN can be viewed as an extension of the control plane of the CCAP core, so if there is 1000 Mbps of DOCSIS and video traffic in the downstream for a given RPD, then that much capacity must be allocated in the CIN, plus some additional capacity for the L2TPv3 overhead used by the DEPI tunnels.

If there is congestion in the CIN, then some packets can be delayed or lost.

Potential Cause 1. QoS

By default, the cBR-8 and the RPDs mark packets associated with PTP traffic and MAP messages with DSCP 46 (EF). Other DOCSIS control messages like upstream channel descriptors (UCD), modem bandwidth request and range response also use DSCP 46:

Item	Per-Hop-Behavior (PHB)	DSCP Value
DOCSIS data (L2TP)	Best Effort	0
PTP	EF	46
GCP	Best Effort	0
MAP/UCD (L2TP, DOCSIS control)	EF	46
BWR and RNG-REG	EF	46
Video	CS4	32
MDD (L2TP, DOCSIS control), voice	CS5	40

Source: Cisco Remote PHY Device Software Configuration Guide for Cisco 1x2 / Compact Shelf RPD Software 5.x

The CIN must be QoS aware so these high priority packets experience minimum delay.

Congestion that creates dropped packets or long queue delays has created PTP issues, late MAP messages and reduced throughput. These types of problems can be seen by observatiion of interface queues on the cBR-8, RPD and CIN devices.

If PTP or MAP messages get dropped or delayed, as evident with clock instability or late MAP messages, then the CIN capacity or QoS design has to be addressed, since these are sent with high priority.

DLM is not intended to handle short durations of jitter because the minimum polling cycle is one second, so it is not able to eliminate late MAP messages in this case.

Potential Cause 2. Delayed Best Effort Traffic

Currently, DLM packet marks are not configurable and uses best effort (DSCP 0). There have been cases in which the CIN experiences congestion which leads to long queue delay limited to best effort traffic.

This has typically shown up as reduced TCP downstream traffic rates, as CIN delays can create relatively large speed drops due to missed or delayed upstream ACKs.

In this case, no late MAP messages or PTP problems are observed, because these high priority packets are not delayed.

Since DLM packets are marked as best effort traffic, this type of CIN jitter can cause spikes in the DLM values. If DLM is used to dynamically adjust the network delay, this jitter can cause an unnecessary increase in the MAP advance timer, which leads to increased upstream request-grant delays.

In this case, it is recommended to use a static network delay value. Cisco also looks at options to enable DSCP values beyond just best effort on DLM in future releases. This can help reduce the upstream request-grant delay but possibly does not address TCP throughput issues if ACKs are delayed across the CIN.

Related Information

- DEPI Latency Measurement
- PTP Design Recommendations For R-PHY Networks
- Technical Support & Documentation - Cisco Systems