DOCSIS® Timing in Converged Interconnect Network for Distributed Access Architectures

Jason Miller
Technical Marketing Engineer
Cisco

John Downey
CMTS Technical Leader
Cisco

Background

In DOCSIS, the only device that transmits on the downstream is the converged cable access platform (CCAP). On the upstream, however, there can be hundreds of modems needing to transmit data. The upstream is deterministic where individual modems are assigned capacity at a given time by the CCAP. To accomplish this deterministic behavior, the upstream is divided into time segments referred to as minislots. The CCAP will assign minislots to modems based on their upstream traffic requirements. Some minislots will also be reserved for contention-based traffic. Modems make requests for upstream transmission opportunities for "best effort" traffic in these contention-based minislots (non-best effort traffic may get dedicated minislots to transmit traffic or make requests).

The CCAP maps out upstream minislot assignments using upstream bandwidth allocation map (MAP) messages. These MAP messages are sent in the downstream and received by all modems. The modems look at these messages to determine which minislots are allocated to which modems and which are for contention-based activities. A modem will only transmit traffic on a minislot assigned to it (or on a contention slot if making a bandwidth request or other station maintenance activity). The MAP messages from the CCAP allocate approximately 2 milliseconds (ms) worth of time [low latency DOCSIS (LLD) will provide options to reduce this value below 2 ms].

It is important for the CCAP and each modem have the same concept of time. For example, in Figure 1, if the CMTS assigns minislots 16-20 to the blue modem and 21-27 to the red modem, both modems need to know when their minislots start in time. If the red modem's clock is off, it can start transmitting too early causing errors to both minislot 20 being used by the blue modem as well as its own minislot 21.

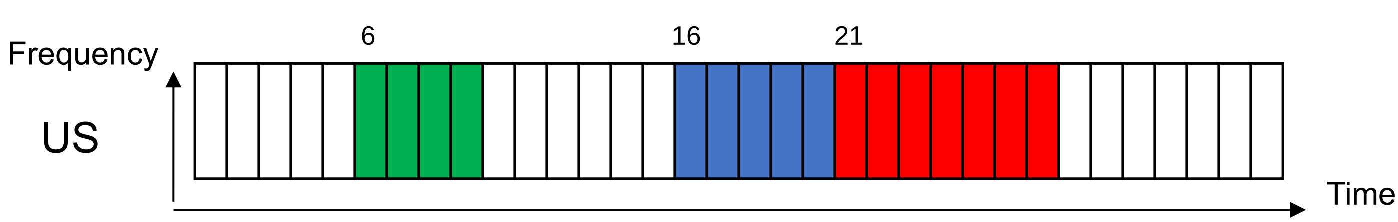

Figure 1

January 2020

Modems will be located different distances from the CCAP. Because it takes time for signals to travel through the hybrid fiber/coax (HFC) plant, modems need to adjust their clocking to account for this propagation delay. They do this by calculating a time offset during initial ranging.

The CCAP should assure that it does not assign a modem a time slot too quickly following a request where the modem has not had time to receive the MAP message and process it. If this happens, the modem learns of its minislot assignment after the minislot time has already passed and the modem is unable to transmit any traffic in the upstream during that period. To avoid this situation, the CCAP uses a MAP advance timer where it doesn't schedule traffic for a modem until a point in time later than the MAP advance timer. This timer can be static or dynamic. If dynamic, the timer value will be made up of downstream interleaving time, a CCAP network time, modem processing time, the time offset of the most distant modem, as well as a safety factor.

Remote PHY

With remote PHY (R-PHY) architectures, the upstream and downstream physical later (PHY) components are no longer in the CCAP device and are instead located in the remote PHY device (RPD). This enables several benefits outside the scope of this paper. A converged interconnect network (CIN) is used to link the RPDs to the CCAP (referred to as the CCAP core with R-PHY). This CIN is IP-based and may be dedicated to cable access or as the name implies, be shared by other applications.

The timing element of DOCSIS needed for upstream scheduling is still present in R-PHY. The CCAP core will handle upstream scheduling and generation of the MAP messages. However, the downstream and upstream signals now physically originate and terminate on the RPD. The RPD needs to have the same concept of time as the CCAP core. Remote DOCSIS timing interface (R-DTI) is the specification from CableLabs that details how this timing takes place. For Ethernet based networks, IEEE 1588 based on precision time protocol (PTP) is used to achieve this timing. In the current implementation from Cisco, both the cBR-8 and the RPD will act as a slave device to a PTP master clock.

PTP Timing

PTP enables a slave clock to determine its time offset from a master clock (difference in time between the clocks) as well as the propagation delay in the transport network between the two clocks. The master and slave device will exchange messages that include time stamps before the slave runs an algorithm to determine these values. The formulas for this calculation assume a symmetrical connection between the two clocks. Non-symmetrical connections like EPON are listed in the R-DTI specification for use as a CIN but rely on a different timing method not currently supported by Cisco. This paper does not attempt to cover PTP in detail as more information is available online in documents like here.

Warning: *One of the main causes of DOCSIS issues in R-PHY is created by non-symmetrical PTP links which lead to clocking instabilities.*

The RPD should reach the master clock via the CIN. The cBR-8 can access the master clock via WAN interfaces on the supervisor physical interface card (PIC) or via digital physical interface card (DPIC) interfaces on the cable line card (the DPIC option was added in 16.8.1 release). The RPD should not pass through the cBR-8 to access the master clock in any 16.8.1 or later release. The RPD and the cBR-8 can only function as slave clocks in current software although the cBR-8 roadmap adds support for it as a grand master and boundary clock.

Note: *Once the cBR-8 is configured to use PTP for timing, all line cards will rely on this clock, even line cards with RF PICs. This means that PTP clock stability issues will impact all modem on a chassis, even those on integrated CCAP (I-CCAP) line cards when running a mix of cards in a chassis.*

cBR-8 PTP Configuration

The cBR-8 acts as the CCAP core to the RPD so will be responsible for the PTP configuration of both itself and associated RPDs. The cBR-8 uses profiles to get this PTP information to the RPDs. There are multiple options for PTP that are configurable. The servo tracking type for the cBR-8 should be set to R-DTI to expedite clock synchronization. The cBR-8 uses a user-defined loopback address as the source of its PTP packets. Assure there is a proper routing for the master clock to reach the loopback interface. If you are unable to ping the master clock with a packet sourced from the loopback address, PTP will not be operational. ITU-T G.8275.2 Telecom profile is recommended as it supports IPv4 and IPv6, and intermediate switching devices in the CIN do not need to be PTP aware. The PTP domain number is chosen by the user but needs to be the same for the cBR-8 and the RPDs. PTP packets will be marked with higher QoS by both the RPD and the cBR-8 for priority on the CIN. Differentiated services code point (DSCP) value 46 (express forwarding – EF) is used by default on the cBR-8. The current R-PHY specifications (CM-SP-R-PHY-I13-190912) recommend using DSCP 47 for PTP packets transmitted from the RPD. Cisco follows this recommendation by default. The specifications are expected to change to use DSCP 46 and Cisco will change to that as the default then. The DSCP value for PTP packets from the RPD can be changed to 46 as shown in the example below and this is recommended.

Below are two examples of PTP configuration for the cBR-8 and a profile for the RPDs. Cisco recommends configuring one clock as preferred when using multiple master clocks.

cBR-8 timing with single master clock source.

```
ptp clock ordinary domain 44
 servo tracking-type R-DTI
 clock-port slave_from_labmaster slave profile g8275.2
  delay-req interval -5
  sync interval -5
  sync one-step
  transport ipv6 unicast interface Lo0 negotiation
  clock source ipv6 2001:DB8:1::251
```

cBR-8 timing with multiple master clock sources. The clock named “slave_from_east” will be preferred as its local priority is lower than the default 128 used by the clock names “slave_from_west”.

```
ptp clock boundary domain 44
 servo tracking-type R-DTI
 clock-port slave_from_east profile g8275.2 local-priority 80
  delay-req interval -5
  sync interval -5
  sync one-step
  transport ipv4 unicast interface Lo0 negotiation
  clock source 100.200.0.100
 clock-port slave_from_west profile g8275.2
  delay-req interval -5
```

```
  sync interval -5
  sync one-step
  transport ipv4 unicast interface Lo0 negotiation
  clock source 100.200.0.200
```

RPD PTP profile which will be referenced by individual RPD configurations on the cBR-8.
Note: *A gateway option can be added after the IP address of the master clock if using a Layer 2 (L2) CIN – not shown, and a Layer 3 CIN is recommended.*

```
ptp r-dti 10
 profile G.8275.2
 ptp-domain 44
 clock-port 1
   clock source ip 100.200.0.100
   clock source ip 100.200.0.200 alternate
   transport dscp 46
```

The PTP clock on both the cBR-8 and the RPD should phase sync with the master before troubleshooting any DOCSIS issues. There are a variety of commands that can show this state along with counts of packets. You want to see packets incrementing for "sync", "delay request" and "delay response". The two commands below are most commonly used on the cBR-8.

```
cbr8#show platform software ptpd stat stream 0
LOCK STATUS : PHASE LOCKED
SYNC Packet Stats
  Time elapsed since last packet: 0.0
  Configured Interval : -5, Acting Interval -5
  Tx packets : 0,  Rx Packets : 93758390
  Last Seq Number : 27388,  Error Packets : 0
Delay Req Packet Stats
  Time elapsed since last packet: 0.0
  Configured Interval : 0, Acting Interval : -4
  Tx packets : 46690369, Rx Packets : 0
  Last Seq Number : 0, Error Packets : 0
Delay Response Packet Stats
  Time elapsed since last packet: 0.0
  Configured Interval : -4, Acting Interval : -4
  Tx packets : 0, Rx Packets : 46657844
  Last Seq Number : 28735, Error Packets : 0
<output omitted>
Current Data Set
  Offset from master :  -0.000000284
  Mean Path Delay    :  +0.000010287
  Forward Path Delay :  +0.000010003
  Reverse Path Delay :  +0.000010601
  Steps Removed 1
<output omitted>
```

```
cbr8#show ptp clock running domain 44 << 44 is the configured domain
                      PTP Ordinary Clock [Domain 44]
State          Ports          Pkts sent      Pkts rcvd      Redundancy Mode
PHASE_ALIGNED  1              46706857       143389129      Hot standby

                       PORT SUMMARY
                                                            PTP Master
```

```
Name       Tx Mode   Role  Transport  State       Sessions   Port Addr
ADVA_IPv6 unicast   slave Lo15886    Slave        1         2001:DB8:1::251

                     SESSION INFORMATION
ADVA_IPv6 [Lo15886] [Sessions 1]
 Peer addr         Pkts in    Pkts out   In Errs    Out Errs
 2001:DB8:1::251  143389129  46706857   0          0
```

The following RPD commands are used to see the RPD's PTP status:

```
R-PHY#show ptp clock 0 state
apr state        : PHASE_LOCK
clock state      : SUB_SYNC
current tod      : 1563227244   Mon Jul 15 21:47:24 2019
active stream    : 0
==stream    0    :
  port id        :              0
  master ip      :   192.169.75.251
  stream state   :       PHASE_LOCK
  Master offset  :              -45
  Path    delay  :             1158
  Forward delay  :             1119
  Reverse delay  :             1198
  Freq offset    :           387714
  1Hz  offset    :               50
```

```
R-PHY#show ptp clock 0 statistics
<output omitted>
streamId  msgType           rx           rxProcessed      lost          tx
 0        SYNC              8585001      8584995          0              0
 0        DELAY REQUEST     0            0                0        8585000
 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    8584998      8584998          5              0
 0        P-DELAY FOLLOWUP  0            0                0              0
 0        ANNOUNCE          536571       536571           0              0
 0        SIGNALING         5593         5593             0           5591
 0        MANAGEMENT        0            0                0              0
   TOTAL                    17712163     17712157         5        8590591
```

Warning: *There have been issues with clock stability on the RPDs when experiencing large changes in network delay between PTP master and the RPD. The RPD can fall back to freerun timing causing multiple issues including modems going offline. RPD releases V6.7 and above filter out large jitter packets and adjust the delay threshold to improve PTP stability when this jitter is relatively high (5 ms or more). Larger jitter values in excess of 100 ms have still caused issues and the fix will be in future RPD code release.*

MAP Advance

The MAP advance timer is used by the CCAP to prevent scheduling a modem to transmit on an upstream minislot that occurs at a time before the modem can receive and process the MAP message (referred to as late MAP messages). This MAP advance timer, which is based in microseconds (µs), 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 interleaving (680 µs with SC-QAM using 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). A shorter MAP advance timer is desirable as it reduces the request-grant delay for upstream traffic. With I-CCAP, the main dynamic adjustment is related to time offset of the most distant modem from the cBR-8.

However, with R-PHY, the CCAP internal delay is now replaced by a "network delay" which defaults to 500 µs. The speed of light through a fiber cable will vary with the velocity of propagation for a given cable type. If the velocity of propagation is assumed to be 67% (velocity factor of 0.67), light will travel approximately 200 km in one millisecond (ms). The default network delay of 500 µs (0.5 ms) will then approximate the one-way propagation delay through about 100 km of fiber. Note that the network delay will also need to account for any queuing and processing delay on intermediate equipment in the CIN like switches and routers, so the distance will be somewhat less. It also will account for processing delays on the CCAP itself. Depending on the CIN design, this value may be larger than the default setting and can change over time.

For OFDMA upstream channels, there is an additional time equal to 3 times the OFDMA frame duration added to the MAP advance to allow for processing delays. The OFDMA frame duration varies with subcarrier size, symbols per frame, and cyclic prefix. The formula for frame duration is symbols per frame x ((1/subcarrier size) + (cyclic prefix samples / 102.4 MHz)). If using 25 kHz subcarrier size, 9 symbols per frame and cyclic prefix of 96 samples, the OFDMA frame duration = 9 x ((1/25 kHz) + (96 / 102.4 MHz)) = 368 µs. The additional time added to the MAP advance would then be 3 times this value or 1104 µs.

The MAP advance values for an upstream can be displayed with the following command.

```
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
```

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

This is the same output for an OFDMA channel. Note that the cyclic prefix, symbols per frame and subcarrier information is also included as part of the output. There are 8 subcarriers per minislot for 50 kHz and 16 subcarriers per minislot for 25 kHz.

```
cbr8#show controllers upstream-Cable 2/0/6 us-channel 12 cdm-ump
<output omitted>
nom_map_adv_usecs 4161, max_map_adv_usecs 4984 mtn_map_adv 8984
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 277
<output omitted>
OFDMA channel config :
cp 96          rolloff 64          sc per ms   16 (subcarrier per minislot)
symbols per frame       9
```

MAP advance = map_adv_safety (1000) + cm_map_proc (600) + intlv_delay (680) + network_delay (500) + max_tmoff (277) + OFDMA processing delay (1104 - see calculation above) = 4161 µs

Early DOCSIS specifications assume the one-way transit delay from the CCAP PHY to the most distant modem is 0.8 ms or less. The DOCSIS 3.1 specification reduced this assumed value to 0.4 ms or less but these times are not hard requirements and Cisco continues to use the 0.8 ms value. This one-way transit delay continues to be observed between the RPD, where the CCAP PHY now resides, and the most distant modem.

The R-PHY specifications from CableLabs do not place a hard limit on the distance (or delay) between the CCAP core and the RPD. If this distance combined with CIN equipment delays exceeds the network delay default value of 500 µs, late MAP messages are possible. The MAP advance "safety" value provides some margin but can also be exceeded in CINs with larger delays. There are two methods to deal with the default network delay setting being an issue and both are set per RPD on the cBR-8. One option is to statically configure the delay and the other is to 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 below.

```
cable rpd parish_brewing_1
 identifier 0027.900a.4d18
 core-interface Te2/1/4
  principal
  rpd-ds 0 downstream-cable 2/0/31 profile 24
  rpd-us 0 upstream-cable 2/0/63 profile 2
  network-delay static 1000 << value in µs
 r-dti 1
 rpd-event profile 0
 rpd-55d1-us-event profile 0
```

The MAP advance safety can also be changed with the cable interface command but only up to 1500 µs

```
cbr8(config-if)#cable map-advance dynamic 1500
```

For dynamic network delay, the cBR-8 relies on a latency measurement specified in the original downstream external PHY interface (DEPI) specifications from CableLabs that came out with DOCSIS 3.0 in 2005. DEPI latency measurement (DLM) determines the one-way delay in the downstream path. The cBR-8 will send out a DLM packet with its time stamp. The RPD will mark its time stamp on the DLM packet when received and forward it back to the cBR-8. Note, Cisco supports the required option where the RPD marks the packet closest to it ingress interface not the option where it marks it closest to the egress interface. The cBR-8 will take the average of the last 10 DLM values and use that average as the network delay value in the MAP advance calculation. The time stamps from both the cBR-8 and the RPD will be based on the PTP reference clocks.

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

By default, DLM is disabled. It can be enabled with the `network-delay dlm <n>` command as shown below. Once enabled, a DLM packet is sent to the RPD every *n* seconds. It can be sent as often as one second as shown in the configuration. There is also a "measure-only" option that can be added that will just measure the CIN delay without adjusting the network delay value. It is recommended to enable DLM at a minimum in the "measure-only" setting to help monitor CIN delay.

```
cable rpd russian_river_5
 identifier 7872.5dd5.e24a
 core-interface Te2/1/0
  principal
  rpd-ds 0 downstream-cable 2/0/5 profile 7
  rpd-us 0 upstream-cable 2/0/7 profile 23
  network-delay dlm 1 << DLM enabled and measured every 1 second
 r-dti 1
 rpd-event profile 0
 rpd-55d1-us-event profile 0
```

In the case of a CIN with longer distances whose delay can change based on link outages like on a fiber ring, DLM has been shown to properly adjust the MAP advance timers so modems do not experience late MAP messages.

If DLM is enabled, this cBR-8 command will display DLM values (CIN is this example is ~1500 km).

```
cbr8#show cable rpd f4db.e6cd.d598 dlm
DEPI Latency Measurement (ticks) for f4db.e6cd.d798
  Last Average DLM:              83287  << values here in ticks (97 ns each)
  Average DLM (last 10 samples): 83287
  Max DLM since system on:       419269
  Min DLM since system on:       82107
  Sample #     Latency (usecs)
  x------------x------------
  0            8126
  1            8136         << values here in µs
  2            8124
  3            8153
  4            8127
  5            8147
  6            8121
  7            8129
  8            8127
  9            8140
```

Cisco enables downstream channels on the cBR-8 to be shared across multiple RPDs with a multicast feature called “virtual splitting”. These RPDs can be located different distances from the CCAP core and therefore have different network propagation delay values. The MAP advance timer for all upstream channels in this case will use the largest network delay either statically configured or dynamically calculated based on DLM.

Late MAP Messages

A late MAP message condition occurs when a modem receives a MAP message at a point in time after the timeslots 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 will be unable to perform station maintenance and will go offline.

With R-PHY, the cBR-8 sends MAP messages to the modems in a DEPI tunnel and to the RPD in an upstream external PHY interface (UEPI) tunnel. These messages have higher QoS marking with a DSCP value of 46 (EF). If a MAP message destined for the RPD gets dropped in the CIN, the RPD will not be able to able use those minislots and will count them as “unmapped”. If the MAP message arrives late at the RPD, it will initially count the minislots as unmapped and then after receiving the late MAP, will increment 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 with just 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 is using two ticks per minislot (popular for 6.4 MHz SC-QAM), a 2 ms MAP will have 160 minislots.

The RPD command below displays late MAP counts. The first value is the RF port on the RPD (0) while the second value is the upstream channel (3). The values for all but the last line will reset as the command is repeated.

```
R-PHY#show upstream map counter 0 3
Map Processor Counters
==============================================
Mapped minislots                   :    1217915
Discarded minislots (chan disable):          0
Discarded minislots (overlap maps):          0
Discarded minislots (early maps)  :          0
Discarded minislots (late maps)   :          0
Unmapped minislots                 :          4
Last mapped minislot               :   16714392
```

Congested CIN

The CIN can be viewed as an extension of the control plane of the CCAP core. Cisco usually avoids oversubscribing the control plane on the cBR-8 and would recommend the same in the CIN. If there is 1000 Mbps of DOCSIS and video traffic in the downstream for a given RPD, then that much capacity should be allocated in the CIN plus some additional for the L2TPv3 overhead used by the DEPI tunnels. There are no hard requirements to dedicate any capacity in the CIN, and in the end, it is up to the operator to determine any level of oversubscription they will allow and still provide a quality service to their end users.

If there is congestion in the CIN, then some packets can be delayed or lost. By default, the cBR-8 marks 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. The RPD currently uses DSCP 47 for PTP packets it sources per the RPHY specifications. The recommendation is to change this to DSCP 46 as mentioned in the PTP section on this document. Also, many PTP masters use a best effort (DSCP 0) for sourced PTP packets by default and this is also recommended to be changed to DCSP 46. The CIN should be QoS aware so these high priority packets experience minimum delay. Currently, DLM packet marking is not configurable and will use best effort (DSCP 0).

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 looking at interface queues on the cBR-8, RPD and CIN devices. If PTP or MAP messages are getting dropped or delayed, as evident with clocking instability or late MAP messages, then the CIN capacity or QoS design should 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 and will not be able to eliminate late MAP messages in this case.

There have also been times when the CIN experiences congestion leading 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 message or PTP problem were observed because these high priority packets were not delayed. Since DLM packets are marked as best effort traffic, this type of CIN jitter will cause spikes in the DLM values. If DLM is being used to dynamically adjust the network delay, this jitter will cause an unnecessary increase in the MAP advance timer leading to increased upstream request-grant delays. In this case, it is recommended to use a static network delay value. Cisco is also looking at options to enable DSCP values beyond just best effort on DLM in future releases. This will help reduce the upstream request-grant delay but may not address TCP throughput issues if ACKs are delayed crossing the CIN.

Very Short CIN Delays

There have been issues observed with dropped upstream DOCSIS traffic when the MAP advance timer is under 2500 µs. Some modems may take longer to process MAP messages and that extra delay may cause late MAP messages for that modem (the RPD may 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 above 2500 µs.

Long CIN Delays

Long CIN delays due to either long distances and / or high processing / queuing delays on CIN equipment will require larger MAP advance timers. This will in turn create a longer request-grant delay leading to increased upstream latency. Since steady upstream traffic flows use piggy-back requests, upstream traffic speed test can appear normal. This will also not impact voice flows using unsolicited grant service (UGS) as no requests are needed. However, downstream TCP traffic speeds can be reduced due to lack of timely upstream ACKs. Although it may be 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 will proactively provide grants to modems based on historical usage, minimizing 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 will require changes to the modem configuration file. DPS is 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. DPS has been very effective in resolving slow TCP downstream throughput associated with delayed ACKs. As one example, a customer was using a leased line to connect the CCAP core and an R-PHY shelf until a better option was available. The CIN delay was over 30 ms! The customer first had to enable DLM for the modems to even come online. Once online, modems were only able to get half of the provisioned downstream speed when running TCP traffic. Once DPS was enabled, modems were able to get expected downstream speeds.

Summary

DOCSIS has a timing element necessary for scheduling upstream traffic. With remote PHY deployments, this timing relies on PTP to synchronize clocks to a master source. Symmetrical routes between the master and slave clocks are required for PTP to work properly. The CIN that links the CCAP core and RPDs can introduce delays that must be accounted for in the MAP advance timer. Cisco recommends enabling DLM in "measure-only" mode as a minimum so CIN delays can be monitored. If CIN network delays exceed the default of 500 µs, then DLM should be enabled without this option so the MAP advance timer can be adjusted for the longer network delay. High jitter and / or drops in the CIN that impact high priority traffic like MAP messages should be resolved through proper QoS designs or increased CIN capacity. If high jitter and / or drops are just impacting best effort traffic with low QoS values, then it may be necessary to statically set the network delay to prevent unnecessary increases to the MAP advance timers as DLM packets can be delayed. This congestion will likely slow down TCP traffic leading to customer dissatisfaction, so should be avoided. Long CIN networks can lead to increased request-grant delays but DPS has been shown to eliminate these delays by providing upstream grants before they are requested by modems.

Summary of Configuration Recommendations

The MAP advance configuration is applied per MAC domain (cable interface) but is calculated on a per-upstream basis. Network delay is configured under the RPD and is specific to the CIN connection between that RPD and the CCAP core. Network delay is a variable used in the MAP advance calculation.

1. Cisco recommends enabling DLM in "`measure-only`" mode as a minimum so CIN delays can be monitored.
   ```
   cable rpd typical_rpd
     core-interface Te2/1/0
        network-delay dlm measure-only
   ```
   ! The default is 10 s but can be set as low as 1 s

   Monitor DLM with ***show cable rpd <mac> dlm***

2. If CIN delays exceed the default of 500 µs, then DLM should be enabled without the "measure-only" option so the MAP advance timer can be adjusted for this longer network delay. This is also recommended when secondary paths with significant delay differences from primary are used for redundancy.
   ```
   cable rpd long_rpd
     core-interface Te2/1/0
       network-delay dlm 1
   ```

3. High jitter and / or drops in the CIN that impact high priority traffic like MAP messages should be resolved through proper QoS designs or increased CIN capacity.

4. PTP packets sourced from the RPD will use DSCP of 47 by default as recommended by the current RPHY standards. This value can be changed to a more common DSCP 46 (EF).

```
ptp r-dti 10
 clock-port 1
   transport dscp 46
```

5. If high jitter and/or drops are just impacting best effort traffic with low QoS values, then it may be necessary to statically set the network delay to prevent unnecessary increases to the MAP advance timers as DLM packets can be delayed. This congestion will likely slow down TCP traffic leading to customer dissatisfaction, so should be avoided. The CIN jitter could also be attributed to CIN design flaws that should be resolved. Statically setting the network delay to the DLM value when the network is using the longest path and is stable can help mitigate modem issues until stabilized.

```
cable rpd jitter_rpd
  core-interface Te2/1/0
    network-delay static 2000
```

6. Long CIN designs can lead to increased request-grant delays but DPS has been shown to eliminate these delays by providing upstream grants before they are requested by modems. Configure DPS under the cable interface.

```
cbr8(config-if)#cable upstream dps
```

__Note__: As of the date this paper was published, DPS is not officially supported, but has been in cBR-8 code since R-PHY support

7. Aggressive map advance (< 2500 µs) could justify increasing the map advance "`safety`" to achieve at least 2500 to avoid the following modem problems.
   - Excessive flaplist "misses" compared to "hits"
   - Upstream bonding partial mode
   - Excessive modem T3 timeouts
   - Upstream voice issues

   Monitor MAP advance values and verify it is above 2500 µs for all upstreams with

   ```
   Cbr8show controllers cable x/y/z upstream | inc Dyn
   ```

   Configure MAP advance under the cable interface.

   ```
   cbr8(config-if)#cable map-advance dynamic 1200
   ```

   Range of 300 to 1500 µs (default 1000 µs)