

Optimized CMTS Commands

John J. Downey – Cisco Systems
CMTS Technical Leader
6/29/05

1. Map Advance

This info is explained in the white paper; Understanding Map Advance:

http://www.cisco.com/en/US/tech/tk86/tk89/technologies_tech_note09186a00800b48ba.shtml

Digest version: Multiply 16 times the miles of fiber for the farthest modem on that particular US. Add 300 to that and use for the max-delay in the map advance command. Example; 25 miles from CMTS to CM. $25 \times 16 + 300 = 700$. Use the command: `cab map-advance dynamic 600 700`.

2. cable insertion-interval automatic 60 480

This interface command is used to control the frequency of insertion intervals. It affects all upstream ports of a MAC domain. It may help modems come online faster. The minimum of 60 ms means when not much traffic is passing or when many modems are offline, the insertion interval will be every 60 ms giving modems plenty of opportunities to do initial maintenance and come online. If there is a lot of US traffic and only a few modems offline, the insertion interval will be every 480 ms (~1/2 sec) so as not to waste too much time for contention and allocate more time for real data traffic. That is, for 60-480, the IM is scheduled every 60 ms when there is little traffic and scales up linearly to 480 when there is a lot of traffic. The general idea is that when there is not much traffic then not many modems have registered and so by scheduling IM often we don't lose anything, and possibly gain by letting them come online faster.

If the insertion interval is hard-set to 60, then each reserved period will be used for IM and won't be used for anything else.

Note: "Shut" USs that aren't being used or have no modems online.

The DOCSIS 1.1 scheduler reserves space every 60 msec for "something". That something can be an IM slot, or it can be a DOCSIS 1.0 CM that wants to schedule a data grant that can not be fragmented and could not find any other place (because there was a lot of UGS on the US).

If you have UGS + DOCSIS 1.0 modems and you set the insertion interval to 60 ms, then you will run into problems. If you have UGS and 1.0 modems using all available minislots, you're gonna have problem with IM no matter what settings you use!

Do "show int cab x/y mac-sched x | inc Initial" in two, 60 sec intervals and count how many IMs were generated in between. Also, use "sh cont cx/y uz" to display the current IM time. Run the command below a few times to verify IM slots are incrementing.

```
ubr7246#sh int c6/0 mac-s 0 | inc Mtn
      Init Mtn Slots 7125, Stn Mtn Slots 3131259
ubr7246#sh int c6/0 mac-s 0 | inc Mtn
      Init Mtn Slots 7133, Stn Mtn Slots 3131260
```

Q&A

[Why is the lowest setting only 100 when you hard-set it \(without auto\)?](#)

I guess it's an attempt to prevent the user to shoot himself in the foot. The IM is 2 ms long; if you schedule it every 20 ms you end up using up (a fixed) 10% of your BW for IM. Since if you set it to 60, it reserves it, it would be an issue. I would set it to auto 60 480 like the default so it can breath depending on utilization and it's still an increment of 60 in case the US is totally full of UGS.

[Is there a reason why you specify 60 480?](#)

Since the 1.1 scheduler already allocates reserved time every 60 msec anyway, why not make the init maintenance an increment of 60.

[It appears from documentation that the default is 50 500.](#)

That was an old default. The default now is auto 60 480.

[Curious about the slight differences. I guess I would be interested in the logic for choosing a slightly lower or higher value than 50.](#)

I guess making the first number lower is fine, since it only occurs when there are not many modems online or passing traffic anyway. Therefore, every 20 msec (50 per sec) would use more time slots that aren't reserved in the scheduler, but no ones using them anyway. The one limitation is the IM must be aligned with UGS. UGS grants could be 10, 20, or 30 msec. Scheduling an IM every 10 msec would affect UGS flows that are every 20 msec. Since most systems will use 20 msec voice packetization, using the lowest number of 20 or an increment of 20 for the highest number would be optimum. If using 20 and 30 msec voice, then a common increment of 60 would be used for IM.

3. cable up<#> data-backoff 3 5 & cable up<#> range-backoff 3 6

Cable modems may fail to range RF US levels correctly and cycle to maximum power. This increases their connection time drastically and some may not come to maintenance state for hours. Try putting the following commands on the upstream interfaces:

```
cable up<#> data-backoff 3 5
cable up<#> range-backoff 3 6
```

The **cab up x range-backoff** default values may be too small and need to be changed from **automatic**. Once these changes are applied and tested, the cable modems may be able to range the RF US levels immediately after re-booting. This could reduce the connection time to less than a few minutes. Once a collision occurs, an algorithm forces the modems to retry after backing off exponentially, decreasing the probability of another collision. The amount of backoff is controlled by the interface command, cable upstream x range-backoff 3 6. 3 means 2 to the 3rd power, which = 8. 6 means 2 to the 6th power, which = 64. The modem will randomly backoff between 8 and 64 initial maintenance opportunities. The maintenance opportunities are controlled by the command, cab insertion interval auto 60 480. If the current IM is 250 msec, the CM could back off between 2 and 16 seconds before trying to range again.

The **cab up x data-backoff** command's algorithm is used to help alleviate multiple collisions of upstream requests. In real-world testing, it has been found that the old defaults of cab data-backoff 0 4 would result in multiple collisions because the CMs would all make requests based on the first number and not between the two numbers. Making the first number 3, gives a random backoff between 1 and 8 and if another collision occurs, the back off will be from 8 to

32. Hopefully, the CM is smart enough to piggyback subsequent requests so they are carried in the DOCSIS header of a reserved data grant and not sent in contention time.

Both the above commands (or at least the functionality of the commands) are required by DOCSIS. The range and data-backoff commands, in and of themselves, do not help alleviate collisions; but allow manipulation of the algorithm, which is used to calculate re-transmit times. This somewhat, pseudo-random re-calculation alleviates the likelihood of another collision. Run the command below a few times to verify UCDs are incrementing.

```
ubr7246#sh cont c6/0 u0 | inc UCD
    UCD Count = 8351
ubr7246#sh cont c6/0 u0 | inc UCD
    UCD Count = 8352
```

Q&A

[Do you recommend Automatic or hard coded figures? If Automatic, does the command actually show up in the running config?](#)

Hardcoded, because we did some testing and found that the default settings are not optimized. If using IOS 12.2(15)BC2 or later the default settings for data-backoff reflect the above recommendations of “3 5”.

Look at “sh controllers cx/y uz” to see what the current settings are.

```
ubr10k#sh cont c5/0/0 u0
Cable5/0/0 Upstream 0 is up
Frequency 23.984 MHz, Channel Width 3.200 MHz, QPSK Symbol Rate 2.560 Msps
Spectrum Group is overridden
US phy SNR_estimate for good packets - 31.7884 dB
Nominal Input Power Level 0 dBmV, Tx Timing Offset 1683
Ranging Backoff Start 3, Ranging Backoff End 5
Ranging Insertion Interval automatic (60 ms)
Tx Backoff Start 3, Tx Backoff End 6
```

[Perhaps just choosing the default is best option for 12.2\(15\)BC1b and up, going forward.....thinking out loud.](#)

Probably, since the defaults use [data-backoff 3 5](#) in BC2 code.

[What is the worst case scenario that a cable modem in a highly populated plant will experience between a Range Request / Response session?](#)

The CM will range on all UCDs that are present. This is why I make sure US ports are shut off that aren't being used. We've seen a case one time where an Eriksson CM got stuck in init(r1) and wouldn't let other modems come online. So, the answer to your question is infinity! In reality, the CM should stop trying to range and scan DS for another carrier. Usually the ranging is in 3 dB increments and will go from 6 to 61 dB or so in about 30 secs. Then move to another UCD.

[I want to understand how long this process could take in a populated system.](#)

If there are a lot of modems online and transmitting data, we limit the amount of init maintenance opportunities to allocate more time for real data transfer. The least amount of slots will be about every .5 seconds. Once the CM gets the correct UCD it should have no problem

talking with the CMTS in 10 seconds or less. We have a range back-off algorithm also. Some people set the command to range back-off 3 6 which means 2^3 to 2^6 . This means to backoff if you have a collision with another CM ranging to a random init maintenance slot anywhere between 8 to 64 init slots. If each slot is .25 seconds apart, that could be 16 seconds to start ranging after a collision with another CM ranging.

4. Station Maintenance (SM)

SM on Cisco universal broadband routers is one second per modem up to 20 modems. IOS before 15BC was 25. For example, if there are only four modems on a particular Media Access Control (MAC) domain (one downstream and all its associated upstreams), each modem will be polled every four seconds. Once you have 20 or more modems, it stays at 20 seconds. This feature can be turned off for lab testing with the hidden, global, test command: Router# test cable minimum-poll off, then the rate can be set with: Router(config-if)# cable polling msec . The default is 20,000 milliseconds. If you have 5 modems, you can still set the polling to 20 seconds for a lab environment.

When the default of test cable minimum-poll on is used, the station maintenance period can be changed with the interface command cable polling-interval <10 - 25000> in milliseconds. This is a hidden, interface command and thus, unsupported. It may be advantageous to set this to 15 seconds whenever there are more than 1500 devices on a DS.

Station maintenance is a maximum of every 15 seconds when Hot-standby Connection-to-Connection Protocol (HCCP) is configured for N+1 availability. Once one maintenance message is lost, it goes into a quick mode where a maintenance message is sent every 1 second. After 16 total messages are missed, the modem is considered off-line. If a modem doesn't receive a station maintenance message because of a DS problem within its T4 timer (30 sec min to 35 sec max), it will go offline and restart scanning DS.

Tip: Use Show Cable Hop (SCH) to see the current station maintenance period.

Upstream Port	Port Status	Poll Rate (ms)	Miss Poll Count	Min Poll Sample	Missed Poll Pcnt	Hop Thres	Hop Period (sec)	Corr FEC Errors	Uncorr FEC Errors
Cable3/0/U0	33.008 MHz	789	* * *	set to fixed frequency	* * *			0	9
Cable4/0/U0	down	1000	* * *	frequency not set	* * *			0	0

Take the Poll Rate number, divide by 1000, and multiply by the number of modems registered in that MAC domain. Example: SCH says 789, and I have 19 modems on C3/0. $(789/1000)*19 = 14.99$, or roughly 15 seconds per modem (calculations done with HCCP on this box). Station maintenance at a rate of once every 15 seconds*19 modems = 1.27 station maintenance instances per second. If the CMTS sends a station maintenance instance to each modem once every 25 sec*1500 cable modems, this would equal 60 station maintenance instances/sec being generated by the CMTS.

If the DS is disconnected, the modems will lose DS sync. The default for DS sync from the CMTS is usually every 10 msec. The modem itself may have a timer for the DS lock and could be vendor specific. DOCSIS 1.0 specifies 600 ms as DS sync loss, **but it doesn't specify what the CM should do after sync loss**. Most current modems can endure up to 5-10 seconds of DS sync loss without re-registering.

Modems will also lose their station maintenance when the DS is disconnected. Modems are typical polled every 20 sec. Since the modem has a 30 sec T4 timer, it could potentially have a 10 second window before the CM T4 timeouts.

Another issue could be added attenuation. More than a few dB of padding (US and/or DS) could cause CM resync. This could be vendor specific as well.

If the US is disconnected, the modem also loses Station Maintenance. The modem's T4 timer will reset every time the DS is received, so a modem shouldn't fail from a T4 timeout when the US is disconnected. Once the modem receives the DS poll, it sends a range response and sets a T3 timer = 200 msec. This timer was only used for init maintenance of 1.0 CM, but for all station maintenance on 1.1 and above modems. A T3 timeout is possible from bad time offsets, noisy US, or a disconnected US. The CMTS sends a poll every second after missing the first one. After 16 attempts, the CMTS then declares the modem offline.

Once a timeout occurs, the modem will either try a new UCD and/or start rescanning DS frequencies. There are more timers added for DOCSIS 2.0.

5. cable sync-interval

```
ubr7246(config-if)#cab sync-interval ?  
<1 - 200> sync interval - msec
```

This downstream command is used to specify the interval between successive sync message transmissions from the Cisco CMTS to the modems in that MAC domain. The default is 10 ms. It has been found in lab and field testing that an interval of between 2 ms and 5 ms may work best when using an N+1 solution. The quicker interval may allow some modems to sync to the Protect downstream signal where they might have fallen offline with an interval of 10 ms. Testing should be done with your specific plant to verify which setting is optimal.

Commands to add:

cable default-phy-burst - allows max burst with 2000 as default and 0 to turn it off
cable downstream channel-id - helps with downstream override
cable downstream frequency - needed for N+1, built-in UPx, and override features
cable downstream if-output - helps with troubleshooting
cable downstream interleave-depth - Helps with impulse noise on DS with 32 default.
cable downstream override - Helps move modems to a different DS quicker. Only supported within a 28U card, not across cards. The USs must be the same channel width, minislot size, and mod profiles.
cable downstream rate-limit (shaping, max-delay, etc) - not useful for 10K, but good for VXR
cable upstream admission-control - helps with registration limitation
cable upstream minislot-size - helps keep granularity of "slices". Try to keep 16-B per minislot.
cable upstream power-adjust (continue, threshold, and noise) – refer to:
http://www.cisco.com/offer/sp/pdfs/cable/cable_land/returnpath_wp54.pdf
cable upstream rate-limit - helps control US traffic
cable upstream snr-polling - ??
cable upstream timing-adjust - ??
cable upstream unfrag-slot-jitter - ??

docsis-mode	Upstream channel DOCSIS mode
New	
ubr7246-2(config-if)#cab upstream 0 docsis-mode ?	
atdma	DOCSIS 2.0 ATDMA-only channel
tdma	DOCSIS 1.x-only channel
tdma-atdma	DOCSIS 1.x and DOCSIS 2.0 mixed channel

equalization-coefficient Enable equalization coefficient

Not on by default, but some modem vendors may have problems with pre-equalization anyway. I think this is specifically Pre-EQ.

fragment-force Force fragmentation for upstream performance optimization

Miss-spelling on optimization

ubr7246-2(config-if)#cab upstream 0 fragment-force ?

<0-4000> Fragmentation threshold

The 28U US chip is a Broadcom 3138 and supports up to 4000 bytes, but the 5x20 TI4522 only supports up to 2000-B fragments, whether it's because of the chip or the lookout or some other component.

ubr7246-2(config-if)#cab upstream 0 fragment-force 4000 ?

A carriage return after frag-force is a default of 2000 with 3 fragments.

<0-10> Number of equal size fragments

Default is 3. Values of 0 and 1 are just for troubleshooting.

ingress-noise-cancellation Enable or disable ingress noise cancellation

On by default and set at 200 msec. The 5x20 can't be turned off.

maintain-psd Maintain Power Spectral Density

This is not supposed to be there in BC2 code, I believe. What it does is change the US power level based on the channel width change the user makes. If you double the channel; width, the power level setting (cab u0 power-level) will go up by 3 dB.

ubr7246-2(config)#cab freq-range ?

european Set upstream freq range to 5-65 MHz

japanese Set upstream freq range to 5-55 MHz

north-american Set upstream freq range to 5-42 MHz

I'm glad to see that this is independent of the annex mode. I like the fact that I can do annex B and still use 5-65 MHz on the upstream if I want.

1. Here is a suggested ACL from PSIRT for W32.Blaster and vulnerable MS protocols. The reason for two of the same ACLs: Some users claim better performance if each interface has a separate ACL. (I.e. ACL 111 on inbound WAN and 112 on inbound cable)

! --- block Nachi / W32.Blaster related protocols

```
access-list 111 deny tcp any any eq 135
```

```
access-list 111 deny udp any any eq 135
```

! --- block other vulnerable MS protocols

```
access-list 111 deny udp any any eq 137
```

```
access-list 111 deny udp any any eq 138
```

```
access-list 111 deny tcp any any eq 139
```

```
access-list 111 deny udp any any eq 139
```

```
access-list 111 deny tcp any any eq 445
```

```
access-list 111 deny tcp any any eq 593
```

! --- block remote access due to W32.Blaster

```
access-list 111 deny tcp any any eq 4444
```

Final ACLs would look like this:

```
access-list 111 deny tcp any eq 1080
```

```
access-list 111 deny tcp any any eq 135
```

```
access-list 111 deny udp any any eq 135
```

```
access-list 111 deny udp any any eq 137
access-list 111 deny udp any any eq 138
access-list 111 deny tcp any any eq 139
access-list 111 deny udp any any eq 139
access-list 111 deny tcp any any eq 445
access-list 111 deny tcp any any eq 593
access-list 111 deny tcp any any eq 4444
access-list 111 permit ip any any
```

```
access-list 112 deny tcp any eq 1080
access-list 112 deny tcp any any eq 135
access-list 112 deny udp any any eq 135
access-list 112 deny udp any any eq 137
access-list 112 deny udp any any eq 138
access-list 112 deny tcp any any eq 139
access-list 112 deny udp any any eq 139
access-list 112 deny tcp any any eq 445
access-list 112 deny tcp any any eq 593
access-list 112 deny tcp any any eq 4444
access-list 112 permit ip any any
```

2. The ISP Security Essentials book recommends the following acl:

```
# access-list 150 deny ip host 0.0.0.0 any
# access-list 150 deny ip 10.0.0.0 0.255.255.255 255.0.0.0 0.255.255.255
# access-list 150 deny ip 127.0.0.0 0.255.255.255 255.0.0.0 0.255.255.255
# access-list 150 deny ip 169.254.0.0 0.0.255.255 255.255.0.0 0.0.255.255
# access-list 150 deny ip 172.16.0.0 0.15.255.255 255.240.0.0 0.15.255.255
# access-list 150 deny ip 192.0.2.0 0.0.0.255 255.255.255.0 0.0.0.255
# access-list 150 deny ip 192.168.0.0 0.0.255.255 255.255.0.0 0.0.255.255
# access-list 150 deny ip 224.0.0.0 31.255.255.255 224.0.0.0 31.255.255.255
# access-list 150 permit ip any any
```