

Cisco Cable Modem Provisioning Scenarios

John J. Downey – Cisco Systems
CMTS Technical Leader
3/20/2020

Table of Contents

[Introduction](#)

[Setup Requirements and Specifications](#)

[First Time Provisioning](#)

[DS Frequency Scanning Tables](#)

[Other Considerations](#)

[IP Address Assignment](#)

[Max Devices](#)

[Oversubscription](#)

[Scenarios](#)

1. [One downstream feeding 12 nodes and 2 nodes per upstream port](#)
[Problem - DOCSIS config file lists the wrong DS frequency](#)
[Problem - US is wired to wrong blade and all ports use the same US freq](#)
[Problem - US is too noisy](#)
[Problem - Already provisioned and lose "station maintenance" because of disconnected US/DS](#)
[Problem - 3.75 dB of loss induced in the reverse path](#)
[Problem - CPE Power outage then back on](#)
2. [Dense mode configuration and the idea of load balancing](#)
[Problem - First time provisioning](#)
[Problem - RF outage, then reconnected](#)
3. [Two or more DS frequencies from the same CMTS](#)
[Problem - First time provisioning and you don't want the first DS frequency](#)
[Problem - Already provisioned and you want a different DS frequency from the original](#)
4. [Two or more DS frequencies from different CMTSs](#)
[Problem - First time provisioning and you don't want the first frequency](#)
5. [Two or more DS frequencies from different linecards, but separate US networks](#)
[Solution](#)

[Questions/Notes](#)

[Commands and Tips](#)

Introduction

There are many different scenarios and permutations in regards to physically wiring cable modem termination systems. We could have a sparse mode where upstream ports on the CMTS are kept separate, a dense mode where signal is fed to multiple upstream ports, multiple CMTSs on the same physical plant, and/or different card densities. These combinations will affect what happens when provisioning, maintaining, and troubleshooting.

The following five “CMTS / cable modem” combinations are an attempt to address the issues associated with these permutations. Each combination may have multiple scenarios and recommendations. We also try to address typical setup requirements, specifications, and default settings.

Setup Requirements and Specifications

1. If utilizing a C6U upconverter from Motorola/GI, make sure the freq. says 1.75 MHz below the center freq. and the input is about 20 dBmV. The GI C8U displays the correct center frequency. A EuroDOCSIS upconverter needs 36.125 MHz intermediate frequency (IF) input and the filter is better suited for the 6.952 symbol rate of the 8 MHz DS carrier. DOCSIS output spec. is 50 - 61 dBmV.
2. If using an MA4040D upconverter from Vecima (new name for WaveCom), make sure the center frequency is selected and the input is between 28-35 dBmV. If the IF at 44 MHz has an output power greater than 32 dBmV, appropriate padding is necessary. The latest linecard output is ~ 42 dBmV.
3. The CMTS upstream input is typically set for 0 dBmV and has an internal IF for the upstream of 70 MHz. Be very cautious on inserting high signals (>30 dBmV) at 17.5 MHz or 35 MHz as the 4th or 2nd harmonics, respectively, could be created and “blow out” the IF at 70 MHz. This is not a problem on the 5x20 and 28U linecards because new US Phy chips do not use a fixed IF. They use broadband direct sampling; the IF is digital. DOCSIS states < 35 dBmV total power per US port from 5-42 MHz.
4. DOCSIS frequencies are 88-860 MHz for DS and 5-42 MHz for US. The U cards can be configured for 5-65 MHz on the US. Ironically, the center frequency for the lowest DS would be 91 MHz, but that’s not a typical NTSC/NCTA channel; 93 MHz is. Also, 855 would be the highest NTSC/NCTA channel giving an upper bandedge of 858 MHz.
5. The cable modem (CM) has an output spec. of 8-58 dBmV for QPSK, up to 55 dBmV for 16-QAM, and 54 dBmV for 64-QAM. Cisco CMs transmit up to 60 or 61 dBmV.
6. The CM input spec, is –15 to +15 dBmV and the total input power should be less than 30 dBmV. Example, if you have 100 analog channels each at ~10 dBmV, that would equal $10 + 10 \cdot \log(100) = 30$ dBmV. A DS input of around –5 to + 5 dBmV average digital power seems optimum.
7. No more than 150 to 200 modems per US/MAC domain. If you’re going to do VoIP, you may want to cut this in half. However, advances in DOCSIS Phy technology may allow greater US aggregated bandwidth allowing more modems per US than currently recommended. Devices such as a digital settops requiring low bandwidth may also be installed and allow more devices to be installed. For guidelines to the maximum number of recommended users on an upstream or downstream port, refer to [What is the Maximum Number of Users per CMTS](#) document.

First Time Provisioning

The modem will scan for the DS frequency. There are about 20 frequency tables in the modem for scanning purposes. These ranges are listed below. Keep this in mind when deciding which frequency to use and also any potential ingress sources like off-air digital channels. The modem may also have EuroDOCSIS and special frequency tables included.

DS Frequency Scanning Tables

79 453000000 – 855000000 6000000	91 55752700 – 67753300 6000300
80 930000000 – 1050000000 6000000	92 79753900 – 85754200 6000300
81 111025000 – 117025000 6000000	93 175758700 – 211760500 6000300
82 231012500 – 327012500 6000000	94 121756000 – 169758400 6000300
83 333025000 – 333025000 6000000	95 217760800 – 397769800 6000300
84 339012500 – 399012500 6000000	96 73753600 – 115755700 6000300
85 405000000 – 447000000 6000000	97 403770100 – 595779700 6000300

86	123012500 – 129012500	6000000	98	601780000 – 799789900	6000300
87	135012500 – 135012500	6000000	99	805790200 – 997799800	6000300
88	141000000 – 171000000	6000000			
89	219000000 – 225000000	6000000			
90	177000000 – 213000000	6000000			

The modem will scan all the standard tables then move on to the HRC tables. In newer firmware the modem will recheck the original DS about every 120 seconds if it was already provisioned at one time. The modem will save the last 3 known good DS frequencies. 453 MHz is the default start for Cisco CMs. The CM locks on to the digital carrier center freq. and looks for the hexadecimal 1FFE MPEG-2 packet identifier (PID), which signifies DOCSIS. It waits for all upstream channel descriptors (UCDs), which are used for US freq, mod profile, channel width, etc. If it's the wrong UCD, the modem will eventually time out from being on the wrong US and try the next UCD till it finally connects. Some modems may actually listen to an upstream channel change (UCC) command sent by the CMTS on the DS to advise the CM which UCD it should be using.

The latest versions of CPE IOS code have essentially three scanning algorithms. 1) scan NTSC, 2) scan selective European center frequencies and 3) do an exhaustive scan that looks for a DOCSIS DS at every 250 kHz and/or 1 MHz frequencies. This could take a long time.

There are 4 standard SA frequency plans.

Plan "a" – North America (NTSC STD and HRC)

Plan "b" – International Hybrid (PAL I, PAL G, NTSC)

Plan "c" – Euro-DOCSIS Frequency Plan with Dual Mode (PAL G, PAL I and NTSC)

Plan "d" – Japan (NTSC Japan)

Tip: Provisioning could be faster if you setup a modem in the warehouse before taking out to the customer's house. After provisioned, be sure to pull the power plug first so the DS and some US parameters are cached. It may also be faster to re-provision a modem by pulling the power to the modem or clear the modem interface via console or CLI command. That way it starts scanning the original frequency table again. It's also recommended to shut down US ports that are not being used so CMs don't needlessly range on an US port that isn't needed or used.

Depending on the modem, the US level starts around 6 dBmV and increments by 3 dB till it hits the CMTS within -25 to +25 dBmV. The modem will use a temporary SID of 0. Once in range, the modem is told to power adjust to its required level (usually 0 dBmV CMTS input, but can be set between -10 to +25 dBmV). This finalizes Ranging 1 {init(r1)}, and then Ranging 2 {init(r2)} commences by fine-tuning the modem in 1 dB increments. The CMTS can track in .25 increments, but the modem can only change in 1 dB increments. Init(r1) is in contention time so collisions could occur. Modems attempt to initialize during the cable insertion interval. Once init(r2) is reached, the modem will get another temporary SID that it usually will keep after full registration. Init(r2) and other provisioning steps are done during reserved times based on the modem's SID. Ranging completes and the CMTS and CM are synchronized.

Other Considerations

Using the example QoS Profile below can cause certain issues:

cable qos profile 6 max-burst 255

cable qos profile 6 max-downstream 64

cable qos profile 6 guaranteed-upstream 64

cable qos profile 6 max-upstream 64

1. The max-burst is in bytes and should be set between 1522 and 4096 depending on the linecard.
2. The default cable interface config setting of "cable downstream rate-limit token-bucket shaping max-delay 128" is optimized for Downstream rate limits that are higher than 85 kbps. $1/.128 = 7.81$ packets per second (PPS) on the DS. If sending 1518-byte packets at 7 PPS, that would equal $1518*8*7 = 85$ kbps. The key word "shaping" is on by default in BC code, but not in EC code. If a service class is offered with downstream rates lower than 85 kbps, there could be issues with dropped packets. Either set the "shaping max-delay" to 256 ms or turn off the "shaping" feature. Turning off the "shaping" feature can lead to erratic traffic patterns on the downstream. This command is relevant for the VXR chassis, but not the uBR10K.
3. A guaranteed US rate of 64 kbps and using QPSK at 1.6 MHz (2.56 Mbps total rate) will only allow 40 CMs to come online because Admin Control is on by default at 100% in some BC codes. $2.56 \text{ Mbps} / 64 \text{ kbps} = 40$.

IP Address Assignment

The next step is IP address assignment. Most systems set up a non-routable address space for the modems such as a 10-net and a public addressing network for the CPEs such as a 24 or 66-net. The command, **cable dhcp giaddr policy**, is used to tell the customer PCs to use the secondary IP address pool. Some configurations will rely on Option 82 to achieve this and leave the command at: **dhcp giaddr primary**.

Tip: To display CPEs associated with a specific modem, use the following commands: "**sh cable modem <client_ip_address>**", or this one; **show int cable x/y modem 0**. The following link explains how to control how many CPEs are allowed to connect to one CM: [How MAX-CPE in DOCSIS File and CMTS Works](#).

Max Devices

There are too many variables to put a real number on the max number of devices supported on a uBR10K chassis. SID space is 8175 per mac domain, but each service flow gets a SID, so assuming an average of 2.3 SIDs per CM (some are dynamic flows and others are nailed up for signaling and BE), that would be 3554 CMs per DS. That in essence could be $3554*5*8 = 142\text{K}$!! We support more service flows if used for multicast.

We publish 45,000 devices supported if using a PRE1 and 64,000 if using a PRE2 or 4 because of memory limitations on the PREs. This assumes one CPE behind the CM. **Note:** 12.2(33)SCC IOS, also called "Bighorn" and later, require more memory per device and necessitates a PRE upgrade to PRE4 to satisfy >30,000 devices per chassis.

Now, to be realistic, service tiers for DOCSIS will dictate how much oversubscription you will really do. But, with e-qam DSs, that may increase now. Assuming 3 DSs (1 linecard DS and 2 e-qam DSs), with a 10 Mbps tier being prevalent very soon in many markets, that could be $37*3/10*100 = 1110$ CMs per service group. This could be $1110/4 = 275$ per US.

This number has been decreasing also because of US tiers increasing to 2 Mbps and higher. The US can be increased with ATDMA, D2.0, but maybe not entirely. If we had each node feeding 2 US channels with 1 being 9 Mbps and the other being 27, that's $36 \text{ Mbps} / 2*10 = 120$ CMs per 2 US ports leading to 240 per DS leading to $5*240*8 = \sim 9,600$ CMs per chassis.

So, you can run the numbers many different ways. If many devices are running low traffic like settop boxes using DOCSIS Settop Gateway (DSG), then you could put many more on the chassis. We have a

system with 1000 devices on an US port, but lots of settop boxes. We have also seen some 10Ks with PRE2s with 60,000 subs.

The majority though, would configure 7+1 redundancy for 7 active linecards and maybe only 4 DSs per linecard used so the 5th DS associated US ports can be used by the first 4 DSs. This could be 20 US ports * 7 cards * 200 subs per US = 28,000 subs.

Each MC3G60 linecard supports 12,000 CMs/DOCSIS devices, 64K DS service flows and 128K US service flows. This number includes 24K CCF flows and a 32K limit on SFIDs per mac domain. The recommended limit for the 7225 VXR is 5000 devices and 10,000 devices for the 7246 VXR.


Oversubscription

Having 100 Mbps from a 200 Mbps “pipe” may only allow a 1:50 oversell, hence only about $200/100 \times 50 = 100$ people with this tier would share the 200 Mbps pipe. That would be $200/100 = 2$ Mbps average rate. Typical oversubscription for the DS is 1:100, but when the tier starts to reach 50% of the pipe, then stat muxing drops. It is conservative to drop from 100 to 50 oversell. The real test is to track actual utilization over time; hours, days, weeks and decide if your oversubscription assumptions are legitimate. “Powerboost” can skew these numbers also.

US oversubscription is typically 1:10, but that was when we had 10 Mbps pipes and only offered 512K to 1 Mbps US speeds. Now that we have increased US tiers to 2 or 3 Mbps and above, the stat muxing would drop. To partially offset this, we have tried to activate ATDMA to make the pipe at least twice as big, if not 3 times bigger. With this info, 1:10 is still a pretty good rule-of-thumb. Keep in mind that we are talking about the median tier, not the super high tier where only a few people subscribe. So, if the typical tier is 2 Mbps and we used 6.4 MHz at 16-QAM (18 Mbps usable pipe), then $18/2 \times 10 = 90$ subs per US.

Scenarios

1. One downstream (DS) frequency feeding 12 nodes and 1 upstream (US) frequency with 2 nodes per port feeding 6 US ports. (Typical Setup) Half of this is shown below.



A. Problem: - DOCSIS configuration file lists the wrong Downstream (DS) frequency

After the modem and CMTS are synchronized with levels and timing, it obtains the IP address through DHCP and its DOCSIS config file through TFTP. The modem starts rescanning because it's told to use a different DS frequency from the DOCSIS config file.

Solution: Either leave the DS freq. empty in the DOCSIS config file or program it correctly. The DS freq. listed in the uBR config file has little effect when using a uBR with an external UPx in this scenario.

Note: When the DS frequency and DS channel ID are set in the cable interface configuration, the “cable downstream override” command could be an issue when multiple DS frequencies are on the same plant. This command is meant for scenarios where the modem can see 2 different DS frequencies from the same CMTS, but is only wired to one US or multiple USs from the same mac domain. The DS freq. in the uBR config also has an affect when doing N+1 redundancy. The external upconverter with SNMP capability needs to know the DS frequency from the uBR configuration when a failover occurs.

Tip: It is recommended to allow all modems to register and non-pay customers to download a “disable” config file that signifies Network Access “False” in the file. To convert the non-pay to a paying sub, update the database to give a normal config file, and “bounce” the modem via SNMP, or use the **clear cable modem** <mac-address | ip-address> **reset** command, or simply let the user power cycle the modem. A new command exists to remove a modem from the CMTS database; **clear cable modem** <mac-address | ip-address> **delete**.

B. Problem - First time provisioning, US is wired to wrong blade/card, and all ports use the same US frequency

Modem scans for DS, locks on, and acquires UCD and time slot to transmit. The US transmission could affect existing modems' throughput and takes up contention time for other modems to provision. R1 starts, but never finishes because of a CMTS 16-second keepalive failure. It starts scanning DS again, locks back on original DS frequency and the process starts all over. Because the VXR is clocked from a single source, the initial maintenance slots are somewhat time aligned across linecards, helping to mitigate incorrect wiring from affecting “real” traffic.

Solution. Wire it up correctly the first time. Cisco currently has a feature called virtual interfaces that will allow up to 8 USs to be assigned to a DS within the new 5x20 or 28U cards so the user can decide what combinations of DS and USs to use.

C. Problem - US is too noisy

R1 completes with high enough level for the modem and CMTS to talk. R2 instructs the modem to lower level. It goes back and forth a few times then stays on the higher level to allow R2 to complete. As a consequence of the high noise, the Range(complete) fails and the modem starts rescanning DS.

Note: If an S-card is used in conjunction with spectrum management, the modem can change modulation profiles, power levels, bandwidth from between 3.2 MHz to 200 kHz, or hop to a different frequency that is either programmed (40 spectrum groups) or decided upon by the S-card. All of this can be accomplished by tracking CNR/SNR, Uncorr/Corr FEC errors, and station maintenance. The con to this is more bandwidth must be allocated for backup.

D. Problem - Already provisioned and lose "station maintenance" because of disconnected US/DS

Station maintenance on Cisco universal broadband routers is one second per modem up to 20 modems. IOS before 13BC 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. Station maintenance is a maximum of every 15 seconds when Hot-standby Connection-to-Connection Protocol (HCCP) is configured for N+1 availability and may not reset if HCCP is un-configured except when rebooted.

The automatic polling 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"** (10-30000). 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 upstream ranging-poll interval <20000 - 30000>** in milliseconds. This is a hidden, interface command and thus, unsupported. It may be advantageous to set this to 20 seconds whenever there are more than 1500 devices on a DS.

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 offline. If a CM doesn't receive a station maintenance message 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)	Missed Poll Count	Min Poll Sample	Missed Poll Pcnt	Hop Thres Pcnt	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. The counters can be cleared with, "clear cable hop" in 15BC2 code or clear interface cx/y in older code.

If the DS is disconnected, the modem may time out with a T4 timeout or 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 CMs don't re-register immediately after sync loss, but usually have a limit of about 6-10 seconds. T3 is a 200 msec CM timer for a ranging response from the CMTS and T4 is a 30 second CM timer for station maintenance. Depending on where a modem is in its station maintenance timer, it could get a T4 timeout within 10 seconds or 30 seconds. 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. A modem could get many T3 timeouts without actually going offline. This is usually attributed to US issues such as; time offset, noisy, and disconnected US.

E. Problem - Someone induces 4.75 dB of loss in the reverse path

Depending on the version of IOS, the CMTS can have a “Power Adjust Threshold” range (0 to 10 dB) around nominal that can be set to disregard small changes. The default range is +/- 1 dB. The flap-list also has a range that can be set +/-1 to +/- 10 dB for reporting purposes.

Note: Never set the Power-Adjust Threshold range to 0 because modems will never set up unless they hit the CMTS exactly at 0 dBmV and ranging opportunities will be taken up by modems continually changing levels. The Flap-List will be very active! The default range of +/- 1 dB may be sufficient, but a range of +/- 2 may be justified for temperature swings you don’t want to track.

Since the loss was 4.75 dB, the CMTS will instruct the CM to change by ~ 5 dB making the CMTS input +.25 dBmV and within the +/- 1 dB range. The modems that are already maxed out will be instructed to increase power every 20 seconds during station maintenance as long as they are within the “continue” range (this can be changed with the Power-Adjust Continue command). This command has a default setting of -4 from nominal and can be increased to -10. Modems between the continue range and the threshold range will be commanded to change level during its station maintenance even though it can’t, but permitted to stay online. When you do a Show Cable Modem command, you’ll see an “!” next to the level for that modem to signify that it has “maxed out”. Modems that are outside the “continue” range will try a few times, relock on the original DS, retry the level again, and then rescan DS. US power adjustments more than 5 to 6 dB can cause modems to re-acquire.

The CMTS does not limit the max-power adjust value, it takes whatever the PHY chip reports in terms of power error. Once the CMTS is told by the PHY receiver chip, how much the power offset is, the CMTS will use that value to build a RNG-RSP to send to the CM. Same in the case of frequency offset and timing offset.

If the CMTS commands a CM to change the power by a certain level and if the CM is not capable of doing it in one step, the CM is expected to increase the level by the max value the CM can do in one step. In this case when the CMTS still sees a power difference it will command the CM to adjust by the difference.

After attempting to increase the power a certain number of times by commanding through the RNG-RSP messages and the power fails to reach a certain level, the code would detect that condition and it will declare that the CM has reached its maximum power level.

```
ubr203.cmts.stjmny#ping docsis 000c.e50f.7ed8 ver
Queueing 5 MAC-layer station maintenance intervals, timeout is 25 msec:
Reply from 000c.e50f.7ed8: 46 ms, tadj=0, padj=0.75, fadj=78
Reply from 000c.e50f.7ed8: 46 ms, tadj=-1, padj=0.75, fadj=83
Reply from 000c.e50f.7ed8: 46 ms, tadj=0, padj=0.75, fadj=63
Reply from 000c.e50f.7ed8: 50 ms, tadj=0, padj=0.75, fadj=78
Reply from 000c.e50f.7ed8: 46 ms, tadj=0, padj=0.75, fadj=87
```

tadj (TX timing offset adjustment), padj (TX power offset adjustment) and fadj (TX frequency offset adjustment) values displayed are based on the measurements done on the RNG-REQ packet by the PHY chip.

When a DOCSIS ping is done with out specifying the number of maintenance intervals to queue, CMTS will be providing 5 consecutive unicast polling opportunities to the modem, one per second. This will result in five RNG-REQ messages and corresponding five RNG-RSP messages.

In the ranging response messages, CMTS may or may not command the CM to change power level based up on whether the CM received power as measured on the previous RNG-REQ is outside the threshold value or not. In the above example since 0.75 is within the + or - 1 dB threshold (default setting) CMTS won't be asking for a power adjustment in the corresponding ranging response message. This can be verified by capturing the RNG-RSP messages using cable monitor.

In the case of timing offset and frequency offset as well, the CMTS may or may not command the CM to change the values based on whether the measured offset values are within the threshold or not. The timing adjust threshold is very small and frequency adjustment threshold is a relatively large value. These depend on the symbol rate and could vary for different card types. I don't have the absolute values.

Also if you enable the following debug commands, the output will show the timing, power and frequency errors measured on the ranging request packets. Debug output shows the threshold for the frequency offset as well.

```
debug cable mac-address 000f.6663.94ef verbose
debug cable tlvs
debug cable range
```

First of all, DOCSIS does not seem to have the max and min value settings provided. DOCSIS provides the following details:

1) In the Ranging Response message, the power level adjustment TLV can have an 8 bit signed value. So the numeric range is from -128 till +127. And this is represented as multiples of 0.25 dB. So as per the spec, the maximum possible negative value is -32 dB and maximum possible positive value is 31.75 dB.

2) Also the spec says, "When a CM is commanded with finer resolution than it can implement, it MUST round to the nearest supported step size. If the commanded step is half way between two supported step-sizes, the CM MUST choose the smaller step. For example, with a supported step resolution of 1 dB, a command to step ± 0.5 dB would result in no step, while a command to step ± 0.75 dB would result in a ± 1 dB step".

We did some lab tests on a uBR10K using a variable step attenuator and here is what we found so far.

1) When we removed the attenuation by 30 dB, the CMTS commanded the CM to adjust the power in three steps a) -45 times .25 dB b) -45 times .25 dB and c) -23 times .25 dB.

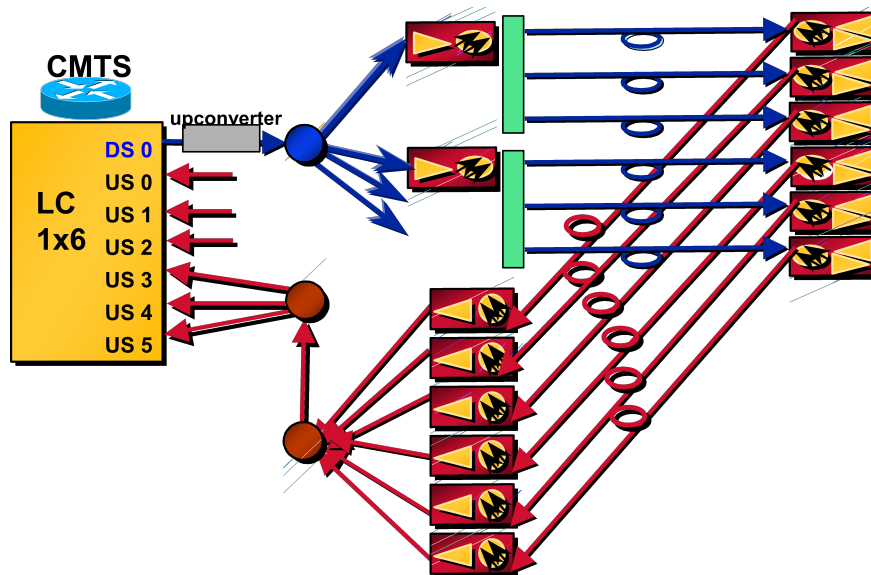
2) -0.5 dB is the smallest commanded power level change we could see on a rng-rsp and this was observed only after changing the US power adjust threshold to zero.

3) If we increase the attenuation above 10 dB, we see the modem go offline. Since the modem receive power level on the CMTS was around 0 dBmV, a sudden dip in power by -10 might have caused the signal to be too weak to be intelligible.

F. Problem - CPE Power outage then back on

Modems remember some of their last settings (DS freq, US freq, modulation, channel width and US transmit power) to speed up a reboot. They don't remember time offsets so when multiple modems re-transmit, their physical distance will help offset the possibility of collisions. 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 3 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**. This command allows the maintenance contention period to be adjusted automatically between every 60 msec to every 480 msec. If there are many modems offline, the initial maintenance will be every 60 msec to help speed up provisioning. When only a few modems are offline, the initial maintenance could be every 480 msec to allocate more time for reserved grants for "real" traffic.

2. One DS frequency feeding 12 nodes and 3 US frequencies with 6 nodes combined then split feeding 3 US ports each for a dense mode configuration and the idea of load balancing.



A. Problem - First time provisioning

Ranging completes and the CMTS and CM are synchronized. Waits for all UCDs. If it's the wrong UCD, the modem will eventually time out from being on the wrong US and try another UCD till it finally connects. After the modem and CMTS are synchronized with levels and timing, it obtains the IP address through DHCP and its DOCSIS config file through TFTP. On a provisioning point of view, you can do client-class processing to force a specific modem MAC address to a specific US. The modem starts transmitting on its required US freq. In the DOCSIS config file you can set the US Channel ID to 0 for telephony return, 1 for US1, , 6 for US5 or leave it blank.

Note: In newer IOS we send the UCDs in a pseudo-random order so the modems won't all pick the first UCD and provision on the same US when doing dense mode combining. This helps with load balancing across US ports. Besides combining the modems, noise and ingress are also combined and wreak havoc.

The following table lists the pseudo-random order of the UCDs.

Timeslot	Upstream Allocation Sequence					
	1st choice	2nd	3rd	4th	5th	6th
A	0	1	2	3	4	5
B	5	0	1	2	3	4
C	4	5	0	1	2	3
D	3	4	5	0	1	2
E	2	3	4	5	0	1
F	1	2	3	4	5	0

Tip: Knowing the sequence of UCDs, can help to determine the best way to physically combine US ports. If 3 US ports will be combined, combine even ports together (0, 2, & 4) and odd ports together (1, 3, & 5). If only 2 ports are used, combine (0 3), (1 4), and (2 5) for perfect balance.

If modems are already scattered between multiple USs, you can force specific modems over to a certain US port without shutting the interface or ports. Use the test command; **test cable ucc cx/y {sid #} {Port #}**. The modem is supposed to change US ports without rebooting. This could take you a while to do each one individually, unless you write some type of prl script. This command is also not supported such as other test commands.

You can also do, **cable modem {mac/IP-address} change-frequency {Ch ID #}**. Ch 1 = US0 Ch 2 = US1, etc. The problem with this command is that it will force the modem to re-acquire and thus, go offline first.

```
ubr7246-2#cable modem 003.e350.97f5 change-frequency ?
```

```
<1-6> Upstream Channel ID
```

Solution: Redo the segmentation to allow only 4 nodes to be combined then split to 2 with 2 US frequencies. This will allow some load balancing, less frequency use, and less noise funneling. Another possibility is to combine 2 nodes to 1 US port with 1 US freq., but load balancing wouldn't be achieved.

It's also possible to set the "Min US Throughput" field in the DOCSIS config file and use the command "Admission Control %" to only allow so many modems to provision on an US till the % of the total possible throughput is accounted for.

IOS > 12.2(15)BC1 introduces a feature called Dynamic Load Balancing and can be configured to balance the modems based on modem count or on actual utilization/load.

B. Problem - RF outage, then reconnected

Modems remember some of their last settings (DS freq, US freq, modulation, channel width and US transmit power) for a speedier reboot. The modem scans for DS and periodically rechecks saved DS frequencies about every 2 minutes. The CM locks on and goes through the normal provisioning setup. For catastrophic failures the exponential backoff ranging helps speed up the boot process by eliminating multiple collisions.

Note: The default insertion interval setting (automatic) configures the uBR to automatically vary (between 60 and 480 milliseconds) the initial ranging times available to new cable modems that attempt to join the network. The "automatic" keyword with this command helps to bring a large number of modems online quickly, for example, after a major power failure. Since the DOCSIS


1.1 code reserves Init Maintenance every 60 msec, the latest code uses 60 msec increments.
Example: “cable insertion interval automatic 60 480”.

Many times a power outage in the plant results in an RF outage to the modems causing a catastrophic failure. The limiting factor for modem reboot could be all the modems attempting to “talk” to the DHCP server for IP addresses. Some commands that are useful to mitigate this potential problem are; Data-Backoff, Range-Backoff, dhcp set max-dhcp-requests=50. It’s also recommended to use an external DHCP server with CNR >= 5.0 for a more evenly distributed request/grant cycle and faster re-provisioning.

Note: Cable modems may fail to range the 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

Upstream Range Backoff values may be too small and need to be changed from its default (automatic). Once these changes are applied and tested, the cable modems may be able to range the RF US levels immediately after the cable modem change-frequency command. This could reduce the connection time to less than a few minutes. The **cab up x data-backoff** command is used to help alleviate multiple collisions of upstream requests.

3. Two or more DS frequencies from the same CMTS



A. Problem - First time provisioning and you don’t want the first DS frequency

Modem scans for DS. Locks on first one it sees. Acquires UCDs and time slot to transmit. It fails US connections and starts scanning again for DS and continues process until it finds the correct DS. Locks on to correct DS and receives proper UCD. It obtains the IP address through DHCP and its DOCSIS config file through TFTP. The modem selects a new DS freq. if the DOCSIS config file tells it to.

Note: If the cable interface command “cable downstream override” command is invoked (by default) it’ll help quickly force modems to the proper DS frequency. This feature was implemented for the case where you may have multiple DS freq.s from the same CMTS, but the modem is only physically connected to one US. Some pre-requisites exist for this feature to work; DS channel IDs need to be unique, DS freq. configured, and all US parameters need to be the same.

Solution: Redo the combining or place notch filters at the modems to eliminate the chance of locking on to the wrong DS frequency. One could also narrowcast the DS into the plant farther downstream maybe


at the hub. If the hub is totally optical with no DS RF, you could put the DS into a 1310 nm Laser than WDM into the 1550 nm path after the EDFA, if present. Be sure the light level is about 10 dB lower than the 1550 and keep in mind that fiber loss is different at both wavelengths. This would require US RF at the hub though. Refer to #5 below for another example.

B. Problem - Already provisioned and you want a different DS frequency from the original

Shut/no shut the interface or clear all the modems to reload a new DOCSIS config file with the specific DS frequency. Some other ways to force the modem to download its new config file would be to execute the cable modem change-frequency command or clear the cable modem one at a time. You may need to make the new config file a different name then what was used originally.

Tip: To force a modem to use a different DS frequency, you can try the following command:
 ubr7246-2#cable modem 003.e350.97f5 change-frequency ?
 <54000000-1000000000> Downstream Frequency in Hz

4. Two or more DS frequencies from different CMTSs



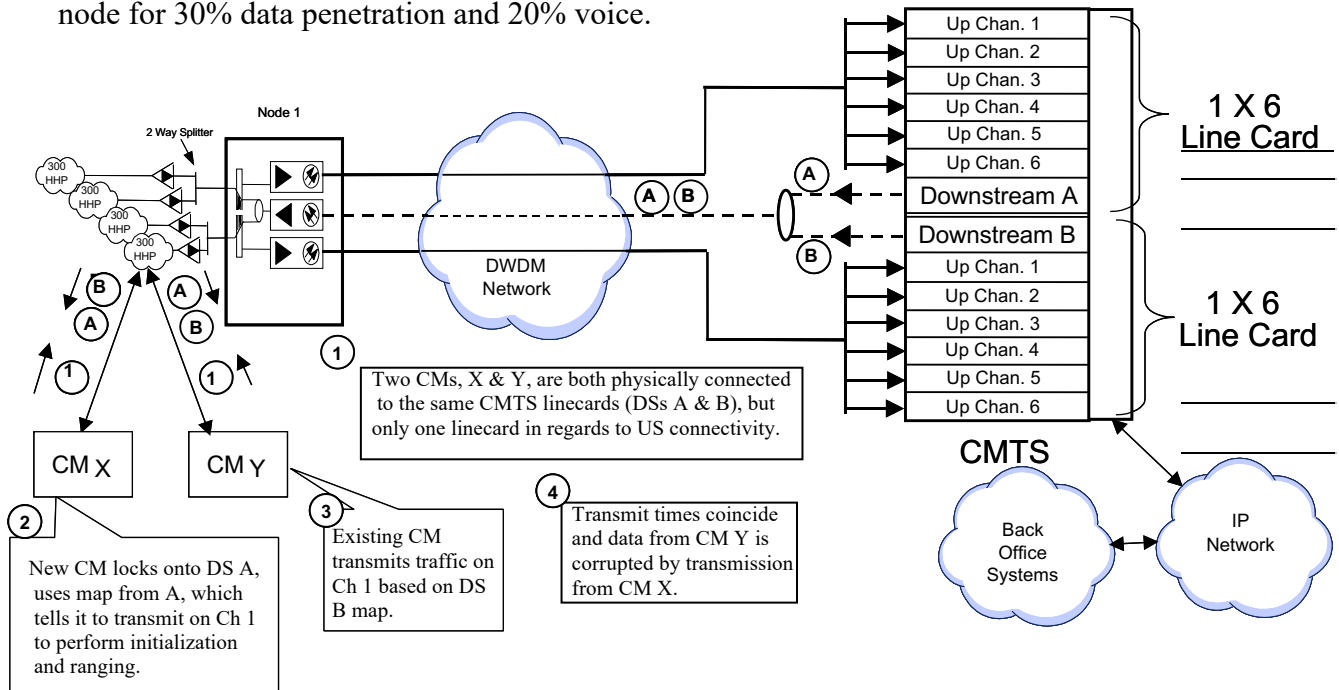
A. Problem - First time provisioning and you don't want the first frequency

Modem scans for DS. Locks on first one it sees. Acquires UCDs and time slot to transmit and attempts registration on the first CMTS. Depending on setup, it may either fail DHCP or the downloaded DOCSIS config file would force it to the correct DS frequency. The CM acquires the commanded DS freq., UCDs and timeslots to transmit. Ranging completes and CMTS and CM are synchronized. It obtains the IP address through DHCP and its DOCSIS config file through TFTP. If DHCP fails, it would try the other UCDs before rescanning DS.

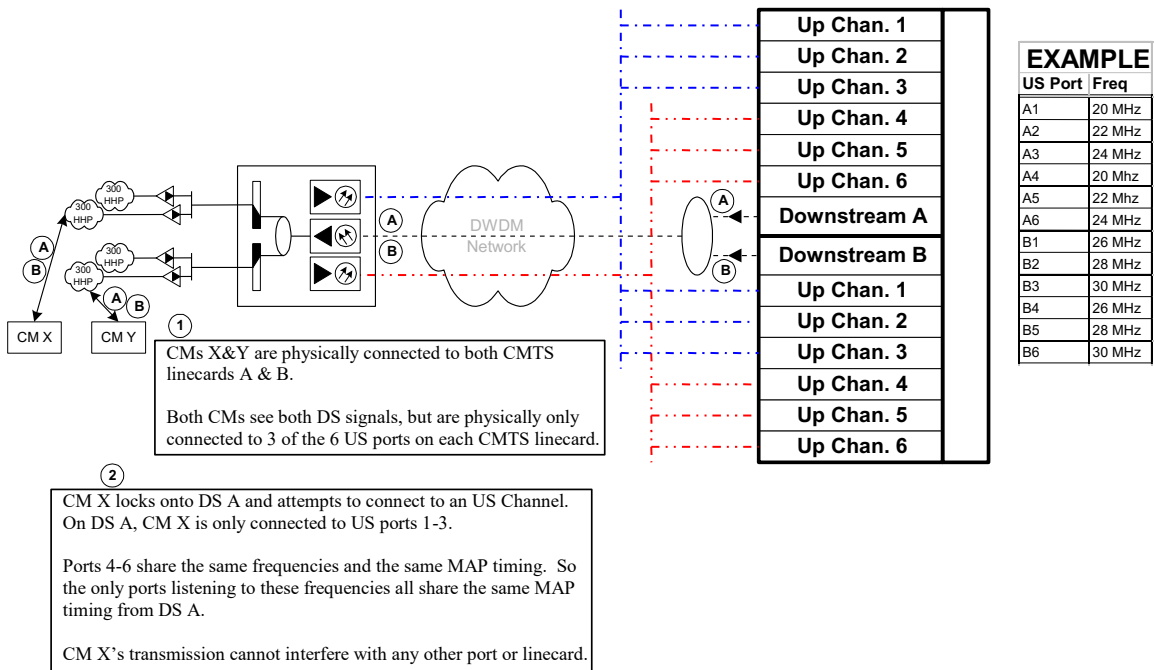
Tip: Invoke the cable interface command “no cable downstream override”. This feature was implemented for the case where you may have multiple DS freq.s, but the modem is only physically connected to one US. This was not intended for the multiple provider scenario. If it were activated, a modem could lock to the correct DS frequency and transmit on the first UCD, hit both CMTS and one of the CMTS would send the DS freq override and it may start scanning on another DS freq even though it didn't have a chance to look at the other UCDs from the first DS freq.

5. Two or more DS frequencies from different linecards, but separate US networks

Current design indicates that a 1x6 line card will be needed for one 600 HHP node for 30% data penetration and 20% voice.



Solution



This issue is mitigated by the fact that the linecards in the uBR7200 chassis are sourced from the same clock and are inadvertently synchronized, but the above solution gives extra assurance.

If the cable interface command “cable downstream override” command is invoked (by default) it’ll help quickly force modems to the proper DS frequency. This topology occurs when the fiber node utilizes

multiple US laser transmitters to segment the node into an asymmetrical topology such as 1200 HHP per DS, but 600 HHP per US.

Frequently Asked Questions (FAQs) & Notes

1. How can “time offset” be used to calculate the distance away?
 - a. The Offset time is directly related to the physical distance away from the CMTS.
 - b. $(\text{tick}/64) * (\text{"time offset"} - \text{magic number}) * (\text{speed of light in a vacuum}) * (\text{Velocity of propagation through the medium}) / (2 \text{ for round trip})$.
 - c. $6.25\text{e-}6 \text{ s} / 64 * (\text{"Time Offset"} - \text{magic number}) * 186,000 \text{ miles/sec. (or } 984\text{e}6 \text{ ft/s}) * V_p$ (.82 for drop coax. .87 for hardline coax, or .67 for fiber) / 2.
 - d. We could assume that whatever the time offset is in a 6-foot plant, that must be the **magic number** to subtract. Example: If the time offset says 3055, subtract 2800 and doing the math from there would = 6.9 miles. I’ve seen different time offsets for different modems from the same vendor or different vendors all on the same 6-foot plant. There is an Excel spreadsheet that does all this math.
2. Does interleaving change throughput? It’s my understanding that it affects latency, but it’s not adding any overhead. Is the latency affecting DS or US throughput or both?
 - a. Decreasing the interleave-depth value can affect upstream performance because it reduces the DS processing time, which affects US packet-per-second rates. Also, it is important to understand that it reduces the time between transmission of a MAP packet, which assigns upstream transmit opportunities, and its reception at the CM.
 - b. Therefore, it could slightly increase the upstream transmission rate (in pps per modem) when the value is set to a lower number.
 - c. The default value is 32. As a workaround to impulse noise issues, you can increase it to 64 or 128. However, by increasing this value you may see US performance degradation (speed) but it will increase noise stability in the downstream. In other words, either the plant has to be very clean, or the customer will see more uncorrectable errors in the downstream, to a point where modems start losing connectivity.
 - d. If you lower the interleave, it should theoretically increase the per modem throughput, but real HFC plant delay may be limiting this anyway.
3. When we set the US bandwidth, the minislot size would change according to certain parameters. Example, when using .4 MHz we had 16 ticks. When changing to 1.6 MHz the default of 8 ticks was not used because the 16 was already set and it doesn’t violate DOCSIS requirements of 32-256 symbols. Is that saying that when I need to send a 6-byte “Request” on the US, I may be using 32 bytes (16 ticks = 100 us, of 1.6 MHz at QPSK)?
 - a. A request is 6 bytes, but when including phy-layer overhead it becomes 16 bytes. So the answer is yes, we’re still wasting time on the wire if the minimum burst size is 32 bytes and we only need to send a 16-byte request.
 - b. Remember, Requests are in the contention/collision time, but can also be piggybacked with real data to eliminate the min-burst issue or collisions.
 - c. Concatenation may also be used to combine multiple grants into one burst. The biggest concatenated packet is also dictated by the bytes per minislot (255 minislots total).
 - d. Either way, it’s best to keep the minislot size down to 8 or 16 bytes for VoIP.

4. There's a field called guard-t under Mod Profile. Our default is 8 symbols. The DOCSIS spec is 5.
 - a. The guard time may vary with the CMTS depending on different vendors.
 - b. The spec mentions that it needs to be $>$ or $=$ to the duration of five symbols plus the max timing error created by both the CM and the CMTS.
 - c. I've noticed that with the Cisco CMTS, the guard time is set to 8 for request, short, and long bursts and 48 for initial and station bursts with QPSK and QAM. Sounds logical. You'd want a better chance of initializing and doing station maintenance and you'd want less overhead time with actual data traffic.
 - d. This guardtime is also different dependent on the actual linecard. The MC5x20S uses a TI upstream chip and requires a 22 symbol gaurdband, while the MC28U uses the new Broadcom chip and requires a variable gaurdband depending on the burst size.
5. Is scrambler like randomization on an analog level or like Manchester coding on a data level. Is it for ones density or for the QAM chip to have different symbols
 - a. It's like Manchester coding on a data level.
 - b. This should never be turned off. I believe you'll end up with the "batman" affect on the frequency trace when viewing with peak-hold on a spectrum analyzer.
6. Burst len; is this in minislots or bytes? Isn't there also a command in the DOCSIS config file to set the max burst?
 - a. Correct, it's in bytes.
 - b. Originally it was in minislots where 255 was a valid number. Currently 255 is not valid in DOCSIS. This value must be either 0 or a number greater than an Ethernet frame.
 - c. Burst length is a user-unique parameter and may vary for each user even when using the same burst type on the same channel as another user. Absence of this configuration setting implies that the burst size is limited elsewhere like in the DOCSIS config file. If in the DOCSIS config file you put the value of zero, then this means that burst length is variable, not fixed and modems will be able to burst at what they request.
 - d. The value of 0 doesn't work for DOCSIS 1.1 modems. It must be 2000 or below. If this is set to 5000, concatenation would be available for 3, 1518 byte Ethernet frames, but there is a problem in the Broadcom chip that doesn't allow it. It must be below 4096 bytes.
 - e. A number above 1522 will cap the modems' requests to a fixed limit.
 - f. The latest BC code has a command of cable default-phy-burst that is defaulted to 2000 bytes. It allows modems to come online when running DOCSIS 1.1 code with upstream concatenation activated even though the DOCSIS config file still has the max burst set at 0, which is normally illegal. The modems would normally get a reject(c) under the show cable modem command, but this new command overrides it.
 - g. The implementation of fragmentation allows modems to concatenate much more than previously allowed and the default-phy-burst command can be set to 0 to disable it.
7. What constitutes a short and a long burst?
 - a. If the minislot is selected for 8 ticks with QPSK at 1.6 MHz channel width, every minislot will be 16 bytes. $1.28 \text{ Msym/sec} * 2 \text{ bits/symbol} * 1 \text{ byte/8 bits} * 8 \text{ ticks/minislot} * 6.25 \text{ usec/tick} = 16 \text{ bytes/minislot}$.
 - b. The normal Max B size setting for a Short IUC in the modulation profile is 6 minislots. $16 * 6 = 96 \text{ bytes}$. Any burst 96 bytes or less will use a Short grant.
 - c. The short grant IUC is intended for TCP acknowledgements and 64-B Ethernet frames.

8. We send a Map every 2 ms, which equals 500 Maps/sec. A Map is about 60 bytes and it will change size depending on the number of US ports on a blade/whole uBR. It's 500 Maps/sec/US, so for a 1x6 card the DS overhead could be ~1.5 Mbps just for Maps.
 - a. Maps and MPEG frames are unrelated. All Ethernet packets are carried in the MPEG-TS payload. 184/4 bytes of each MPEG-TS frame create a continuous sequence of bytes that Ethernet packets overlay. A Map is an Ethernet packet. Its length depends upon how many IUCs (Interval Usage Codes) are in it. There is one IUC for each upstream transmission opportunity, be it a data packet, a request slot, maintenance interval, etc. The Map size can change depending on the US modulation and BW selected.
 - b. Our Maps vary between 2-8 ms. 2 ms is a minimum we use. The 8 ms is related to how long it takes to send a 1518-byte frame plus some other stuff. Smaller Maps are better as that decreases req-to-grant latencies.
 - c. Maps take a hit on the CPU as well as the downstream BW. We could have 500 Maps/US * 6 USs * 4 linecards = 12,000 per uBR. Typically, it'll be closer to 250 maps/sec per US.
9. How will the distribution of the CMs over different Rx frequencies be done, when a group of 3 nodes are connected for dense mode combining with different freqs? How will the load balancing and redundancy be done in this case?
 - a. Initially it is up to the cable modem to pick an upstream channel for which it has received a UCD message. Depending on vendor implementation or if the modem has cached its last US freq, a cable modem might always take the first available upstream channel or choose randomly amongst the available options.
 - b. The newer IOS code sends the UCDs in a semi-random fashion to help provision modems "equally" across the US ports. You should be able to force certain modems to a specific US port through the DOCSIS config file though.
 - c. When "admission control" is enabled in combination with minimum guaranteed upstream bandwidth per modem, the CMTS will not allow certain modems to provision when the configured admission control threshold is reached. This can be set between 10 and 1000%.
 - d. More USs available = more UCDs the CMs must range on and possibly take longer to provision.
 - e. IOS > 12.2(15)BC1 introduces a feature called Dynamic Load Balancing and can be configured to balance the modems based on modem count or on actual utilization/load.

Commands and Tips

For a complete list of very useful commands see the following url:

http://www.cisco.com/univercd/cc/td/doc/product/cable/cab_rout/cr72scg/cr72trbl.htm

Interface Bundling

Used to conserve IP addresses and thus, no need to subnet IP addresses on the uBR. Must assign a master interface and the rest will be slave interfaces. The IP address is assigned to the Master or subinterfaces on the Master. Some systems use subinterfaces for a type of MPLS VPN offering. Keep in mind that if the Master interface goes down, all sub-interfaces will fail as well. The latest IOS will allow CMs to stay online on a Slave interface if the Master fails, but no new modems can come online. There is a feature request DDTS to make the Master interface a loopback address that can never go down. It would be best to use the lowest number interface as the Master so modems can reach the dhcp server after a chassis failure. If the entire chassis is rebooted, it would be best for the Master to initialize first.

cable upstream <n> power-adjust noise 100

A modem receives periodic ranging about once every 20 seconds. Impulse noise, or certain kinds of interfering carriers will force the CMTS receiver chip to mis-read the actual upstream receive power. Often the read-error is 2-5 dB. The command **cable upstream <n> power-adjust noise 30** (default) means that if 30% (3 out of the last 10) Ranging Responses from the CMTS ordered the cable modem to adjust its power by MORE than the number specified in the "**cable upstream <n> power-adjust threshold x**" command, (x by default is 1 dB), then the CMTS should place that modem into a special mode called "power averaging". This is where you have the "*" next to the modem in the "**show cable modem**" display. Setting the averaging to 100% forces every single modem into averaging for all periodic ranging, and might work exceptionally well vs. incredible impulse noise and low C/N.

dhcp set max-dhcp-requests=50

After a HE reboot, all of the allocated buffers quickly fill up. After buffers are filled, the DHCP server throws away additional requests, accepting new ones only as it processes requests and frees up buffers. The first few requests that arrive are responded to quickly. The ones that come after, sit in the incoming buffer queue for several seconds. By the time the DHCP server processes and responds, the client that sent the request has timed out. The parameter max-dhcp-requests controls the number of buffers that the DHCP server allocates to hold incoming requests. I believe this is a DHCP server command.

Remote-Query

The CMTS will always poll only 4-5 CMs per second (depending on the load at CMTS) without any correlation to the querying interval configured. If the interval configured is 15 minutes and there are N modems on the chassis, the complete CM polling will take about $N / 5$ seconds. If $N / 5$ is greater than (15 minutes x 60) seconds, then we cannot finish all the polling in 15 minutes. However, if $N / 5$ is less than (15 minutes x 60) seconds then the polling completes within the specified interval and will commence again periodically at 15 minute intervals.

If there are some CMs that are not responding due to community string mis-configuration etc., then the CMTS retries for about 90 seconds and above calculation becomes: (where a is the number of CMs not responding) $(N - a) / 5 + a \times 90$ greater than/or less than polling interval (15 minutes x 60) seconds. The [CISCO-DOCS-REMOTE-QUERY-MIB ?](#) MIB internally performs a lookup on the CM database that has been populated by the remote-query feature. The purpose of this MIB is to allow an external management station to poll the CMTS of the CM information that the remote-query feature has filled in. If remote-query is not active, polling these MIBS will not return any values.

Addendum:

For the CMTS to place a modem in init(r1) state, it would have to know it is there and ranging. I would say that the CMTS has heard the modem do an initial maintenance, but I'm not sure of the exact time when it goes into init(r2).

Looking at a modem log, you can normally see it start at 6 dBmV and increment in 3 dB steps until the CMTS sees about -20 dBmV or so. At this point the CMTS will say init(r1). Somewhere in this range the CMTS will tell the modem (unicast station maintenance) to change its level to reach the CMTS at 0 dBmV in 1/4 dB increments. In this example, it could say 80 increments. I think once the CM responds with a SM burst, range-request or response? Then the CMTS will place it in init(r2).

The other time I have seen modems get stuck in this state is when the ranging response from the CMTS shows a different upstream channel ID than that of used on the ranging request. As you know, there is no provision in the Initial Ranging request for the CM to show which US it is using the UCD of. When wired in sparse mode, typically all of the UCD's burst profiles will look the same, so the CMTS will see all ranging requests no matter which is used by the CM. Then the CM will (or it should) use the UCD of the US it is wired to as specified in the rng-rsp message. It is my understanding that once the correct US is being used for the exchange by the CM, the CMTS shows r1 after the initial ranging request. Then as the CMTS is trying to fine tune the power/freq/timing/etc, and the ranging status in the ranging response message is ("1"; continue), it will show init(r2). Then after all adjustments have been made, and the CMTS is happy, it will send a ("3" complete) and the CMTS shows init(rc).