

# Cable Modem Provisioning Scenarios

Document ID: 49778

## Contents

### Introduction

#### Setup Requirements and Specifications

#### First Time Provisioning

Other Considerations

#### IP Address Assignment

Scenario 1

Scenario 2

Scenario 3

Scenario 4

Scenario 5

#### FAQ and Notes

#### Related Information

## Introduction

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

The five CMTS-and-cable-modem (CM) combinations in this document are an attempt to address the issues associated with these permutations. Each combination has multiple scenarios and recommendations. Typical setup requirements, specifications, and default settings are also addressed.

## Setup Requirements and Specifications

- If you are using a C6U upconverter from Motorola or General Instruments (GI), make sure that the frequency is set to 1.75 MHz below the center frequency and that the input is approximately 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. The DOCSIS output specification is 50 to 61 dBmV.
- If you are using an MA4040D upconverter from VCom (formerly known as WaveCom), make sure that the center frequency is selected and that the input is between 28 and 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 approximately 42 dBmV.
- The CMTS upstream input is typically set for 0 dBmV and has an internal IF of 70 MHz for the upstream. Be very cautious when inserting high signals (greater than 30 dBmV) at 17.5 MHz or at 35 MHz, as the 4th or 2nd harmonics (respectively) could be created and could blow out the IF at 70 MHz. This is not a problem on the MC5x20U and MC28U linecards, because new US physical layer (PHY) chips do not use a fixed IF. They use broadband direct sampling; the IF is digital. DOCSIS specifies less than 35 dBmV of total power per US port from 5 to 42 MHz.
- DOCSIS frequencies are 88 to 860 MHz for DS and 5 to 42 MHz for US. Ironically, the center frequency for the lowest DS is 91 MHz, but that is not a typical National Television Systems Committee (NTSC) or National Cable and Telecommunications Association (NCTA) channel; 93 MHz is. Also, 855 MHz is the highest NTSC or NCTA channel, giving an upper bandedge of 858

MHz.

- The CM output specification is 8 to 58 dBmV for Quadrature Phase–Shift Keying (QPSK) and 8 to 55 dBmV for quadrature amplitude modulation 16 (16–QAM). Cisco CMs transmit up to 60 or 61 dBmV.
- The CM input specification is  $-15$  to  $+15$  dBmV, and the total input power should be less than 30 dBmV. For example, if you have 100 analog channels each at approximately 10 dBmV, that equals  $10 + 10 \times \log(100)$ , which equals 30 dBmV. A DS input of around  $-10$  to  $+5$  dBmV average digital power seems optimal.
- A general recommendation is to put no more than 150 to 200 modems per US or per MAC domain. If you are doing Voice–over–IP (VoIP), you might want to halve this limit. Advances in DOCSIS PHY technology, however, might allow greater US aggregated bandwidth, permitting more modems per US than is currently recommended. Devices such as a digital set tops requiring low bandwidth might also be installed, allowing more devices to be installed. For guidelines about the maximum number of recommended users on a US or DS port, refer to What is the Maximum Number of Users per CMTS?.

## First Time Provisioning

The modem scans for the DS frequency. There are approximately twenty frequency tables in the modem for scanning purposes, which are listed in Table 1. Keep this in mind when deciding which frequency to use; also keep in mind any potential ingress sources, such as off–air digital channels. The modem might also have EuroDOCSIS and special frequency tables included.

Table	Range (Hz)	Increments (Hz)
79	453000000	855000000
80	930000000	1050000000
81	111025000	117025000
82	231012500	327012500
83	333025000	600000000
84	339012500	399012500
85	405000000	447000000
86	123012500	129012500
87	135012500	135012500
88	141000000	171000000
89	219000000	225000000
90	177000000	213000000
91	55752700	67753300
92	79753900	85754200
93	175758700	211760500
94	121756000	169758400
95	217760800	397769800
96	73753600	115755700
97	6000300	403770100
98	601780000	799789900
99	805790200	997799800
	6000300	

The modem scans all of the standard tables before moving on to the HRC tables. In newer firmware, the modem rechecks the original DS approximately every 120 seconds, if it was already provisioned at one time. The modem saves the last three known good DS frequencies. 453 MHz is the default starting frequency for Cisco CMs. The CM locks on to the digital carrier center frequency 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 frequency, modulation profile, channel width, and so forth. If it receives the wrong UCD, the modem eventually times out due to being on the wrong US and it tries the next UCD until it finally connects. Some modems might actually listen to an upstream channel change (UCC) command sent by the CMTS on the DS to advise the CM as to which UCD it should be using.

The latest versions of customer premises equipment (CPE) Cisco IOS® Software code have essentially three scanning algorithms:

- Scan NTSC.
- Scan selective European center frequencies.
- Do an exhaustive scan that looks for a DOCSIS DS at every frequency that is divisible by 250 kHz or 1 MHz, which could take a long time.

**Tip:** Provisioning can be faster if you setup a modem in the warehouse before taking it out to the customer's house. After provisioned, be sure to pull the power plug so that the DS parameters and some of the US

parameters are cached. It might also be faster to re-provision a modem by pulling the power to the modem or by clearing the modem interface using console or command-line interface (CLI) commands. That way, it starts scanning the original frequency table again. It is also recommended that you shut down the US ports that are not being used so that the CMs do not needlessly range on them.

Depending on the modem, the US level starts at approximately 6 dBmV and increments by 3 dB until it hits the CMTS within 5 to +25 dBmV. The modem uses a temporary Service ID (SID) of 0. Once in range, the modem is told to power adjust to its required level: usually, this is 0 dBmV CMTS input, but it can be set between -10 and +25 dBmV). This finalizes Ranging 1 (R1, init(r1)), and then Ranging 2 (R2, init(r2)) commences by fine-tuning the modem in 1 dB increments. The CMTS can track in 0.25 dB 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 gets another temporary SID that it usually keeps 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 this example quality of service (QoS) profile 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
```

- The max-burst is in bytes, and it should be set between 1522 and 4096, depending on the linecard.
- The default cable interface configuration setting of **cable downstream rate-limit token-bucket shaping max-delay 128** is optimized for DS rate limits that are higher than 85 kbps.  $1 / 0.128 = 7.81$  packets per second (PPS) on the DS. If sending 1518 byte packets at 7 PPS, that equals  $1518 \times 8 \times 7 = 85$  kbps. The key word **shaping** is on by default in BC code, but not in EC code. If a service class is offered with DS rates lower than 85 kbps, there can 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 DS. This command is relevant for the VXR chassis, but not for the uBR10k.
- A guaranteed US rate of 64 kbps using QPSK at 1.6 MHz, which yields a 2.56 Mbps total rate allows only forty CMs to come online, because Admission Control is on by default at 100 percent 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 nonroutable address space for the modems (such as a 10-net) and a public addressing network for CPE (such as a 24-net or 66-net). The **cable dhcp giaddr policy** command is used to tell the customer PCs to use the secondary IP address pool. Some configurations rely on Option 82 to achieve this goal and leave the command as **dhcp giaddr primary**.

**Tip:** To display the CPE associated with a specific modem, issue the **show cable modem client\_ip\_address** command or issue the **show interface cable x/y modem 0** command. Using the max-cpe Command in the DOCSIS and CMTS explains how to control the number of CPE that are allowed to connect to one CM.

## Scenario 1

One DS frequency is feeding twelve nodes, and one US frequency with two nodes per port is feeding six US ports (typical setup).

This diagram shows half of this setup:



### Problem DOCSIS Configuration File Lists the Wrong DS Frequency

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

### Solution

Either leave the DS frequency empty in the DOCSIS configuration file or configure it correctly. The DS frequency listed in the uBR configuration 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 become an issue when multiple DS frequencies are on the same plant. This command is intended for use in scenarios where the modem can see two different DS frequencies from the same CMTS, but it is only wired to one US or to multiple USs from the same MAC domain. The DS frequency in the uBR configuration also has an effect when doing N+1 redundancy. The external upconverter that has Simple Network Management Protocol (SNMP) capability needs to learn the DS frequency from the uBR configuration when a failover occurs.

**Tip:** It is recommended that you allow all modems to register and have non-paying customers download a disable configuration file in which Network Access is set to False. To convert a non-paying modem to a paying modem, update the database to give the modem a normal configuration file and then do one of these things:

- Bounce the modem using SNMP.
- Issue the **clear cable modem {mac-address | ip-address} reset** command. A new command exists to remove a modem from the CMTS database: **clear cable modem {mac-address | ip-address} delete**.
- Tell the user to power cycle the modem.

### Problem First-Time Provisioning, US Is Wired to the Wrong Blade or Card, and All Ports Use the Same US Frequency

The modem scans for DS and locks on. It then acquires a UCD and a time slot to transmit. The US transmission could affect existing modems throughput and takes up contention time for provisioning other modems. R1 starts, but never finishes because of T3 timeout or R1 failure. It starts scanning DS again, locks back on the 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 the effects of incorrect wiring on real traffic.

## Solution

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

## Problem US Is Too Noisy

R1 completes with a high enough level for the modem and CMTS to talk. R2 instructs the modem to a lower level. It goes back a forth a few times, then it 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, change power levels, change bandwidth from 3.2 MHz to 200 kHz, or hop to a different frequency that is either programmed (32 spectrum groups) or determined by the S-card. All of this can be accomplished by tracking carrier-to-noise ratio (CNR) or signal-to-noise ratio (SNR), uncorrectable or correctable forward error correction (FEC) errors, station maintenance, and time or day. The drawback to this is that more bandwidth must be allocated for backup. The benefit is that you could run hotter levels (3 dB), because some of the power that is allocated for frequency is not being used.

## Solution

Refer to Determining RF or Configuration Issues on the CMTS. Also refer to How to Increase Return Path Availability and Throughput and Upstream FEC Errors and SNR as Ways to Ensure Data Quality and Throughput.

## Problem Already Provisioned and Lose Station Maintenance Because of Disconnected US or DS

Station maintenance on Cisco Universal Broadband Routers is one second per modem, up to twenty modems (in Cisco IOS Software Releases prior to 13BC, up to twenty-five modems). For example, if there are only four modems on a particular MAC domain (one DS and all of its associated USs), each modem is polled every 4 seconds. Once you have twenty or more modems, it stays at 20 seconds. This feature can be turned off for lab testing with the hidden, global test command **test cable minimum-poll off**, then the rate can be set with the **cable polling msec** command. The default for msec is 20000 milliseconds. If you have five 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 **cable polling msec** interface command, where msec is a value from 10 to 25000 milliseconds. This is a hidden interface command and, thus, is not supported. It may be advantageous to set this to 15 seconds whenever there are more than 1500 devices on a DS.

Station maintenance occurs at 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 sixteen total messages are missed, the modem is considered offline. If a modem does not receive a station maintenance message within its T4 timer (30 to 35 seconds), it will go offline and restart scanning DS.

**Tip:** Issue the **show cable hop** command to see the current station maintenance period.

Upstream Port	Port Status	Port	Poll Rate (ms)	Missed 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

Divide the Poll Rate value by 1000, then multiply the result by the number of modems registered in that MAC domain. For example, suppose that the **show cable hop** command shows 789 milliseconds and there are nineteen modems on the Cable3/0 interface. That equals  $789 \text{ ms} / 1000 \text{ ms/sec.} \times 19$ , which equals 14.99 seconds, or approximately 15 seconds per modem (calculations done with HCCP on this system). Station maintenance at a rate of once every 15 seconds for 19 modems equates to 1.27 station maintenance instances per second. If the CMTS sends a station maintenance instance to each modem once every 25 seconds for 1500 cable modems, this is equivalent to 60 station maintenance instances per second being generated by the CMTS. To clear the counters, issue the **clear cable hop** command in 15BC2 code, or issue **clear interface cablex/y** in earlier code.

If the US or DS is disconnected, the modem might timeout (with a T3 or T4 timer) or the modem itself might have a timer for the DS lock that could be vendor-specific. DOCSIS 1.0 specifies 600 ms as DS synchronization loss, but it does not specify what the CM should do after synchronization loss. Most CMs do not re-register immediately after synchronization loss, but they usually have a limit of approximately 6 to 10 seconds. T3 is a timer for ranging response from the CMTS, and T4 is a station maintenance timer. Depending on where a modem is in its station maintenance timer, it could get a T4 timeout within 5 seconds or 30 seconds. Once a timeout occurs, the modem either tries a new UCD or starts rescanning DS frequencies or both. There are more timers added in DOCSIS 2.0.

### **Solution**

Reconnect the US or DS cabling.

### **Problem Someone Induces 3.75 dB of Loss in the Reverse Path**

Depending on the Cisco IOS Software Release that you are using, 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  $\pm 1$  dB. The flap list also has a range that can be set from  $\pm 1$  to  $\pm 10$  dB for reporting purposes.

**Note:** Never set the Power-Adjust Threshold range to 0: modems will never successfully setup 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  $\pm 1$  dB might be sufficient, but a range of  $\pm 2$  dB might be justified for temperature swings that you do not want to track.

Because the loss was 3.75 dB, the CMTS instructs the modem to change by 3 or 4 dB, making the CMTS input  $\sim .75$  or  $+0.25$  dBmV (within the  $\pm 1$  dB range). The modems that are already maxed out are instructed to increase power indefinitely, so 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  $\$$  from nominal and can be increased to  $^{\text{TM}}0$ . Modems between the continue range and the threshold range are commanded to change level during their station maintenance even though they can not, but they are allowed to stay online. When you issue a **show cable modem** command, you see a ! next to the level for each modem that has maxed out. Modems that are outside the continue range 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.

### **Solution**

Remove some attenuation, change the CMTS US power level to  $>$  dBmV, or increase the **power-adjust continue** command to 6.

### **Problem CPE Power Outage, Then Back On**

Modems remember some of their last settings (DS frequency, US frequency, modulation, channel width, and US transmit power) to speed up a reboot. They do not remember time offsets, so when multiple modems re-transmit, their physical distance helps to 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.

## Solution

The amount of backoff is controlled by the **cable upstream x range-backoff 3 6** interface command. In that command, **3** means  $2^3$ , which equals 8. **6** means  $2^6$ , which equals 64. Thus, the modem randomly backoffs between 8 and 64 initial maintenance opportunities. The maintenance opportunities are controlled by the **cable insertion interval auto 60 480** command. This command allows the maintenance contention period to be adjusted automatically between 60 to 480 ms. If there are many modems offline, the initial maintenance will be every 60 ms, to help speed up provisioning. When only a few modems are offline, the initial maintenance could be every 480 ms, to allocate more time for reserved grants for real traffic.

## Scenario 2

One DS frequency is feeding twelve nodes, and three US frequencies with six nodes combined then split are feeding three US ports each, for a dense mode configuration and load balancing.



## Problem First-Time Provisioning

Ranging completes and the CMTS and CM are synchronized. The CM waits for all UCDs. If it receives the wrong UCD, the modem eventually times out due to being on the wrong US and it tries another UCD until it finally connects. After the modem and CMTS are synchronized with levels and timing, the modem obtains its IP address through DHCP and it obtains its DOCSIS configuration file through TFTP. From 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 frequency. In the DOCSIS configuration file, you can set the US Channel ID to 0 for telephony return, 1 for US 0, 2 for US 1, 3 for US 2, 4 for US 3, 5 for US 4, and 6 for US 5; or you can leave it blank.

**Note:** In later Cisco IOS Software Releases, the UCDs are sent in a pseudo-random order, so that the modems do not all pick the first UCD and do not 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.

Table 2 lists the pseudo-random order of the UCDs.

**Timeslot 1st Choice 2nd Choice 3rd Choice 4th Choice 5th Choice 6th Choice A 0 1 2 3 4 5 B 5 0 1 2 3 4**

**Tip:** Knowing the sequence of UCDs can help to determine the best way to physically combine US ports. If three US ports will be combined, combine even ports together (0, 2, and 4) and odd ports together (1, 3, and 5). If only two US ports are used, combine 0 and 3, 1 and 4, and 2 and 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. Issue the **test cable ucc cablex/y** {sid-number} {port-number} command. The modem should change US ports without rebooting. Because it could be time-consuming to test each one individually, it is advisable to write some type of PERL script.

**Note:** Like other test commands, this test command is not supported.

You can also issue the **cable modem** {mac-address | ip-address} **change-frequency** {channel-id} command, where channel 1 is US0, channel 2 is US1, and so forth. The problem with this command is that it forces the modem to re-acquire and, therefore, it goes offline first.

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

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

## Solution

Redo the segmentation to allow only four nodes to be combined, then split to two with two US frequencies. This allows some load balancing, less frequency use, and less noise funneling. Another possibility is to combine two nodes to one US port with one US frequency, but doing so does not allow load balancing.

It is also possible to set the Min US Throughput field in the DOCSIS configuration file and issue the **Admission Control %** command to only allow modems to provision on an US until the % of the total possible throughput is used.

Cisco IOS Software Releases later than 12.2(15)BC1 introduce a feature called Dynamic Load Balancing and can be configured to balance the modems based on modem count or on actual utilization or load.

## Problem RF Outage, Then Reconnected

Modems remember some of their last settings (DS frequency, US frequency, modulation, channel width, and US transmit power) for a speedier reboot. The modem scans for DS and periodically rechecks saved DS frequencies approximately 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 Cisco uBR7200 series to automatically vary (between 50 milliseconds and 2 seconds) the initial ranging times available to new cable modems that attempt to join the network. Using the **automatic** keyword with this command helps to bring a large number of modems online quickly (for example, after a major power failure). Because the DOCSIS 1.1 code reserves Init Maintenance every 60 ms, it may be wise to use 60 ms increments in the command (**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.



## Solution

These are some commands that are useful to mitigate this potential problem:

- **data-backoff**
- **range-backoff**

It is also recommended that you use an external DHCP server with a CNR greater than or equal to 5.0, for a more evenly distributed request and grant cycle and for faster re-provisioning.

**Note:** Cable modems might fail to range the RF US levels correctly and cycle to maximum power. This drastically increases their connection time, and some may not reach maintenance state for hours. Try issuing these commands on the upstream interfaces:


```
cable up x data-backoff 3 5
```

```
cable up x range-backoff 3 6
```

Upstream Range Backoff values might be too small and might need to be changed from the default (**automatic**). Once these changes are applied and tested, the cable modems might 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. You can issue the **cable up x data-backoff** command to help alleviate multiple collisions of upstream requests.

## Scenario 3

Two or more DS frequencies from the same CMTS.



### Problem First-Time Provisioning, but the First DS Frequency Is Not Wanted

A modem scans for DS and locks onto the first one that it detects, acquiring UCDs and a time slot to transmit. It fails US connections and starts scanning again for DS, continuing the process until it finds the correct DS. It locks onto the correct DS and receives a proper UCD. The modem obtains its IP address through DHCP and it obtains its DOCSIS configuration file through TFTP. The modem selects a new DS frequency, if the DOCSIS configuration file tells it to do so.

**Note:** If the **cable downstream override** cable interface command is issued (by default), it helps to quickly force modems to the proper DS frequency. This feature was implemented for the case where you may have multiple DS frequencies from the same CMTS, but the modem is only physically connected to one US. To work properly, the DS channel ID must be set, the DS frequency must be set, and the US channels must have the same settings (such as channel width, minislots, modulation profile, and so forth).

## Solution

Redo the combining or place notch filters at the modems to eliminate the chance of locking onto 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, then do wavelength division multiplexing (WDM) into the 1550 nm path after the erbium-doped fiber amplifier (EDFA), if one is present. Be sure that the light level is approximately 10 dB lower than the 1550, and keep in mind that fiber loss is different at both wavelengths. This solution requires US RF at the hub, though. See Scenario 5 for another example.

## Problem Already Provisioned, but a Different DS Frequency is Wanted Other than the Original

Issue the **shut** command and then issue the **no shut** command on the interface; or clear all of the modems, to reload a new DOCSIS configuration file with the specific DS frequency. Other ways to force the modem to download its new configuration file are to issue the **cable modem change-frequency** command or to clear the cable modems one at a time. You might need to make the new configuration file a different name than what was used originally.


## Solution

To force a modem to use a different DS frequency, issue this command:

```
ubr7246-2# cable modem 003.e350.97f5 change-frequency ?  
<54000000-1000000000> Downstream Frequency in Hz
```

## Scenario 4

Two or more DS frequencies from different CMTSs.



## Problem First-Time Provisioning, but the First DS Frequency Is Not Wanted

A modem scans for DS and locks onto the first one that it detects, acquiring UCDs and a time slot to transmit. It attempts registration on the first CMTS. Depending on the setup, it may either fail DHCP or the downloaded DOCSIS configuration file forces it to the correct DS frequency. The CM acquires the commanded DS frequency, UCDs and timeslots to transmit. Ranging completes and the CMTS and CM are synchronized. The modem obtains its IP address through DHCP and it obtains its DOCSIS configuration file through TFTP. If DHCP fails, it tries the other UCDs before rescanning DS.

## Solution

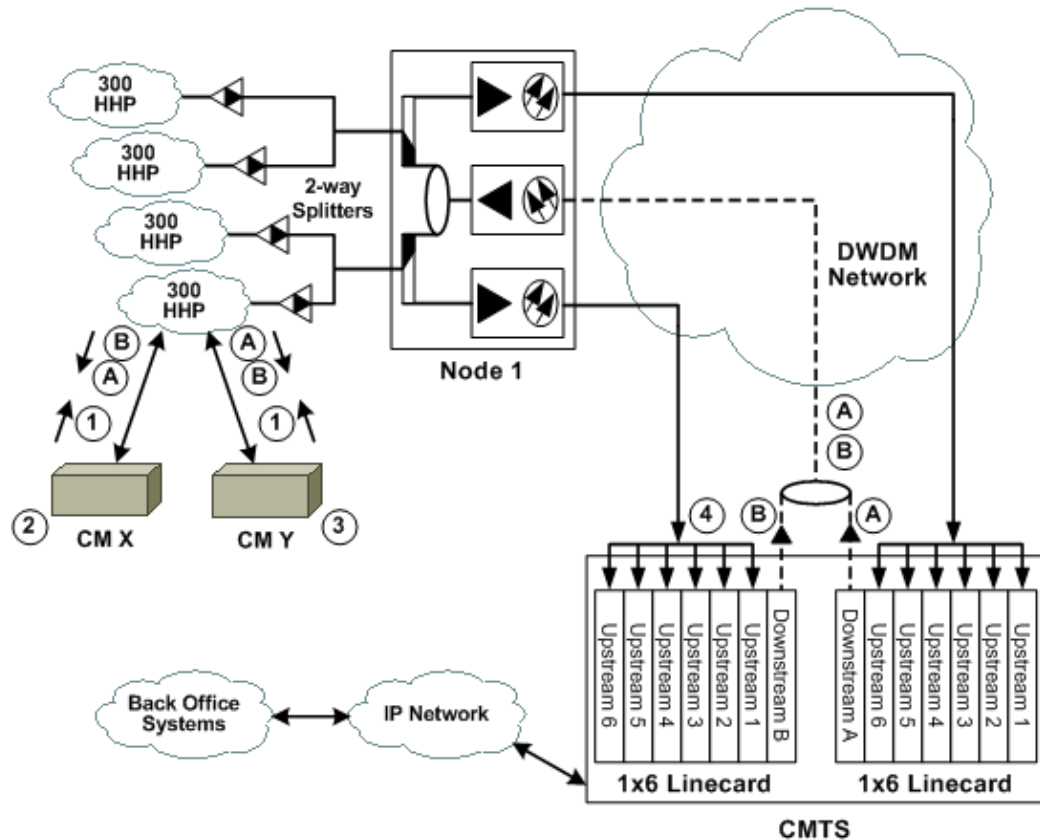
Issue the **no cable downstream override** cable interface command. This feature was implemented for the case where you may have multiple DS frequencies, but the modem is only physically connected to one US; it

is not intended for multiple-provider scenarios. If it is activated, a modem can lock to the correct DS frequency and transmit on the first UCD, hit both CMTSs, and one of the CMTSs sends the DS frequency override. Thus, it may start scanning on another DS frequency even though it did not have a chance to look at the other UCDs from the first DS frequency.

## Scenario 5

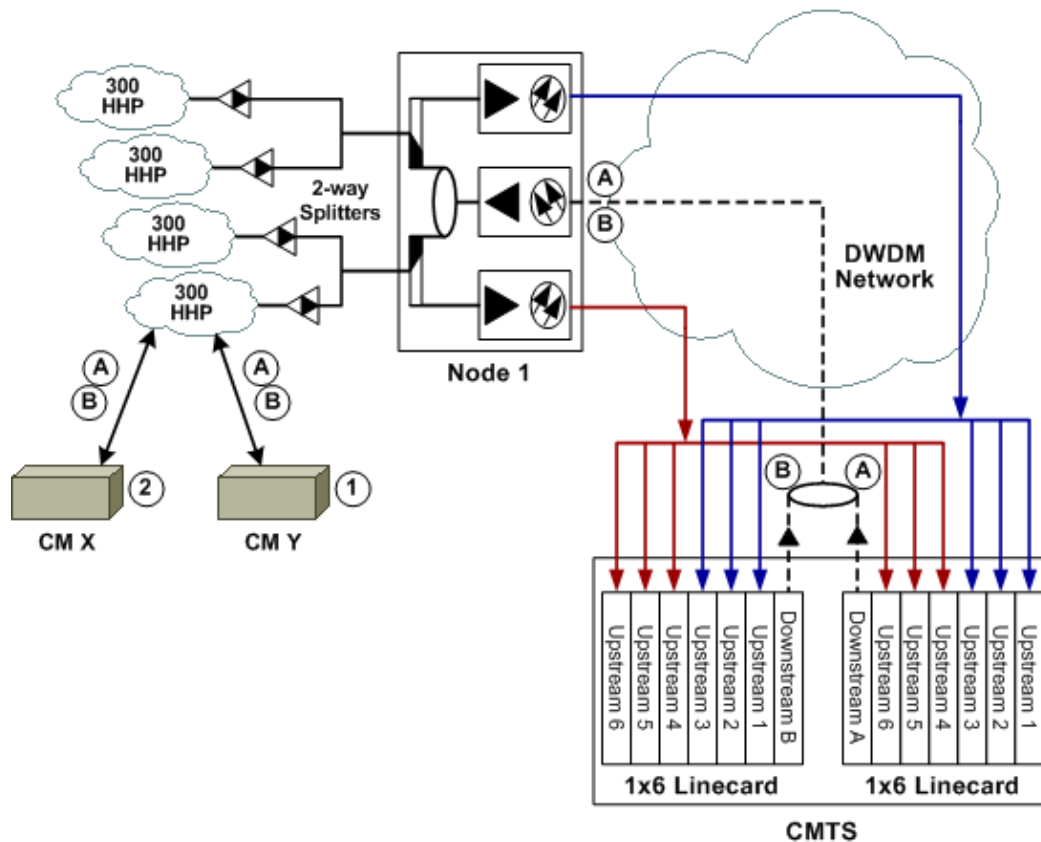
Two or more DS frequencies from different linecards, but separate US networks.

Current design indicates that a 1x6 line card is needed for one 600 HHP node, for 30 percent data penetration and 20 percent voice.



1. Two CMs, X and Y, are both physically connected to the same CMTS linecards (DSs A and B), but they are only connected to one linecard, with regards to US connectivity.
2. New CM (X) locks onto DS A and uses the map from A, which tells it to transmit on US 1 to perform initialization and ranging.
3. Existing CM (Y) transmits traffic on US 1 based on DS B map.
4. Transmit times coincide and data from CM Y is corrupted by transmissions from CM X.

## Solution



1. CMs X and Y are physically connected to both CMTS linecards A and B. Both CMs see both DS signals, but are physically only connected to three of the six US ports on each CMTS linecard.
2. CM X locks onto DS A and attempts to connect to an US channel. On DS A, CM X is only connected to US ports 1 through 3. Ports 4 through 6 share the same frequencies and the same map timing. Therefore, the only ports listening to these frequencies all share the same map timing from DS A; the transmissions from CM X can not interfere with any other port or linecard.

**US Port Frequency** A1 20 MHz A2 22 MHz A3 24 MHz A4 20 MHz A5 22 MHz A6 24 MHz B1 26 MHz B2 28 MHz B3 30 MHz B4 26 MHz B5 28 MHz B6 30 MHz

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 you issue the **cable downstream override** cable interface command (default) it helps to quickly force modems to the proper DS frequency. This topology occurs when the fiber node uses multiple US laser transmitters to segment the node into an asymmetrical topology, such as 1200 HHP per DS but 600 HHP per US.

## FAQ and Notes

### How can time offset be used to calculate the distance away?

The time offset is directly related to the physical distance away from the CMTS, using this equation:

$$((\text{tick} / 64) \times (\text{time offset} \text{ magic number}) \times c \times Vp) / 2$$

In that equation,  $c$  is the speed of light in a vacuum (186,000 mi./sec. or 984e6 ft./sec.) and  $Vp$  is the velocity of propagation through the medium (0.82 for drop coaxial, 0.87 for hardline coaxial, or 0.67 for fiber). The

whole equation is divided by two because it is a round trip. For example:

$$((6.25e-6 \text{ sec.} / 64) \times (\text{time offset magic number}) \times 984e6 \text{ ft./sec.} \times Vp) / 2$$

You can assume that, whatever the time offset is in a 6-foot plant, that must be the magic number to subtract. For example, if the time offset says 3055, subtract 2800 and do the math from there to determine 6.9 miles. Different time offsets for different modems from the same vendor or different vendors all on the same 6-foot plant is possible.

### **Does interleaving change throughput? Interleaving is supposed to affect latency but does not add any overhead. Is the latency affecting DS or US throughput or both?**

Decreasing the interleave-depth value can affect upstream performance because it reduces the DS processing time, which affects US PPS 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. Therefore, it could slightly increase the upstream transmission rate (in PPS per modem) when the value is set to a lower number.

The default value is 32. As a workaround to impulse noise issues, you can increase it to 64 or to 128. By increasing this value, however, 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.

If you lower the interleave, it should theoretically increase the per-modem throughput, but real hybrid fiber-coaxial (HFC) plant delay might be limiting it anyway.

### **Why is the guard-t under Mod Profile set to default of 8 symbols?**

The guardtime (guard-t) might vary with the CMTS, depending on different vendors. The specification mentions that it needs to be greater than or equal to the duration of five symbols plus the maximum timing error created by both the CM and the CMTS.

It has been observed that, with the Cisco CMTS, the guard time is set to 8 for request, short, and long bursts and to 48 for initial and station bursts with QPSK and QAM. This sounds logical because you want a better chance of initializing and doing station maintenance and you want less overhead time with actual data traffic.

This guardtime is also different depending on the actual linecard. The MC5x20S uses a Texas Instruments (TI) upstream chip and requires a 22 symbol guardband, while the MC28U uses the new Broadcom chip and requires a variable guardband, depending on the burst size.

### **Is scrambler like randomization on an analog level or like Manchester coding on a data level? Is it for one s density or for the QAM chip to have different symbols?**

It is like Manchester coding on a data level and should never be turned off. You end up with the batman effect on the frequency trace, when viewing with peak-hold on a spectrum analyzer.

### **Is burst length in minislots or bytes? Is there also a command in the DOCSIS configuration file to set the max burst?**

The burst length is in bytes. 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.

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 (for instance, in the DOCSIS configuration file). If you set the value at 0 in the DOCSIS configuration file, then that burst length is variable (not fixed), and modems will be able to burst at what they request.

The value of 0 does not work for DOCSIS 1.1 modems. It must be 2000 or below. If it is set to 5000, concatenation is available for three 1518 byte Ethernet frames, but there is a problem in the Broadcom chip that does not allow it: it must be below 4096 bytes.

A number above 1522 will cap the modems requests to a fixed limit. The latest BC code has the **cable default-phy-burst** command, which defaults to 2000 bytes. It allows modems to come online when running DOCSIS 1.1 code with upstream concatenation activated, even though the DOCSIS configuration 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.

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.

### **What constitutes a short and a long burst?**

If the minislot is selected for 8 ticks with QPSK at 1.6 MHz channel width, every minislot will be 16 bytes:

$$1.28 \text{ Msym/sec.} \times 2 \text{ bits/symbol} \times 1 \text{ byte/8 bits} \times 8 \text{ ticks/minislot} \times 6.25 \text{ } \mu\text{s/tick} = 16 \text{ bytes/minislot}$$

The normal max-burst size setting for a Short Interval Usage Code (IUC) in the modulation profile is 6 minislots.  $16 \times 6 = 96$  bytes, thus any burst 96 bytes or less will use a short grant. The short grant IUC is intended for TCP acknowledgements and 64 byte Ethernet frames.

**A map is sent every 2 ms, which equals 500 maps/sec. A map is about 60 bytes and it changes size depending on the number of US ports on a blade or on the whole uBR. Thus, it is 500 maps/sec/US, so for a 1x6 card the DS overhead could be ~1.5 Mbps just for maps.**

Maps and Moving Picture Expert Group (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 on how many IUCs are in it. There is one IUC for each upstream transmission opportunity, be it a data packet, a request slot, maintenance interval, and so forth. The map size can change depending on the US modulation and bandwidth (BW) selected.

Maps can vary between 2 and 8 ms: 2 ms is the minimum used, while 8 ms relates to how long it takes to send a 1518 byte frame plus some other stuff. Smaller maps are better, as they decreases request-to-grant latencies.

Maps take a hit on the CPU as well as on the downstream BW. You could have  $500 \text{ maps/US} \times 6 \text{ USs} \times 4 \text{ linecards}$ , which equals 12000 per uBR. Typically, it is closer to 250 maps/sec/US.

**How will the distribution of the CMs over different Rx frequencies be done, when a group of three nodes is connected for dense mode combining with different frequencies? How are the load balancing and redundancy done in this case?**

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 depending on whether the modem has cached its last US frequency, a cable modem might always take the first available upstream channel or might choose randomly amongst the available options.

Newer Cisco IOS Software code sends the UCDs in a semi-random fashion, to help provision modems equally across the US ports. You should, however, be able to force certain modems to a specific US port through the DOCSIS configuration file.

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 threshold can be set between 10 percent and 1000 percent.

More USs available means more UCDs that the CMs must range on and, possibly, means they take longer to provision.

Cisco IOS Software Releases later than 12.2(15)BC1 introduce a feature called Dynamic Load Balancing and can be configured to balance the modems based on modem count or on actual utilization or load.

## Related Information

- [Broadband Cable Technology Support](#)
  - [Technical Support – Cisco Systems](#)
- 

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2014 – 2015 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

---

Updated: Apr 03, 2006

Document ID: 49778

---