

Troubleshooting Slow Performance in Cable Modem Networks

Document ID: 12551

Introduction

Before You Begin

Conventions

Prerequisites

Components Used

Accurately Determining the Levels of Performance Being Achieved

Measuring the Correct Parts of the System

Determining the Download and Upload Rate

Potential Reasons for Poor Performance

Performance Restricted by DOCSIS Configuration File

Using a Sub-optimal Method for Rate Limiting

Upstream Channel Congestion

Downstream Channel Congestion

Backhaul Network or Internet Congestion

Noise and Errors on the Cable Plant

High CPU Usage on the CMTS

Under Powered or Malconfigured CPE Equipment

Conclusion

Related Information

Introduction

There are a number of issues that can affect the performance and speed of cable modems in a Data-over-Cable Service Interface Specifications (DOCSIS) system. This document seeks to address the major causes of slow throughput from the perspective of a cable service provider.

This document first looks at how to determine accurately what kinds of throughput levels an end user is achieving and how to make sure that the performance being measured is that of the cable network, rather than that of the broader Internet.

The next section looks at the most common potential reasons for slow performance and suggested resolutions. These issues include:

- Performance being restricted by the limits in the DOCSIS configuration file.
- Bursty or inconstant download performance caused by using a sub-optimal rate limiting scheme on the cable modem termination system (CMTS).
- Upstream and downstream channel congestion.
- Backhaul network or Internet congestion.
- Noise or errors on the cable plant.
- Under powered end user customer premises equipment (CPE).

Each of these individually or in combination can affect throughput and performance in a cable network.

This document does not discuss troubleshooting a complete loss of connectivity over the cable network or cable modems not coming online. Instead, refer to [Troubleshooting uBR Cable Modems Not Coming Online](#) for these kinds of problems. In addition, for end users of a uBR 900 or CVA 120 series cable modem

experiencing performance issues, the best starting place for troubleshooting this problem is [Beginners FAQ for uBR900 Series Cable Modem End Users](#).

Before You Begin

Conventions

For more information on document conventions, see the [Cisco Technical Tips Conventions](#).

Prerequisites

There are no specific prerequisites for this document.

Components Used

The information in this document is based on the software and hardware versions below.

- Cisco IOS® Software Release 12.1(9)EC for the uBR7200 and uBR7100 CMTS.
- The Cisco uBR7100, uBR7200, and uBR7200VXR suite of CMTS products.
- The information in this document is relevant for all other currently available releases of DOCSIS 1.0–based Cisco IOS software for Cisco brand CMTS equipment.

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

Accurately Determining the Levels of Performance Being Achieved

Measuring the Correct Parts of the System

There are a number of ways to gauge the speed and performance of a system, however it is important to understand exactly what parts of the system are being tested. Consider the diagram below.




Figure 1 (To see this diagram as a Flash animation click [here](#).)

In this diagram there are a number of components:

- The Hybrid Fiber Coax network between the end user and the CMTS.
- The Local CMTS network segment where the CMTS connects to the cable service provider's network.
- The cable service provider's internal network.
- The public Internet.

When performing a speed test between two points, the speed of all the network components between the two points is being measured.

For example, if performing a speed test between the CPE and Server 3, which is connected to the Internet through a 128 Kbps ISDN line, there never will be speeds of greater than 128 Kbps, even if the available bandwidth on the cable segment is greater than 128 Kbps.

The most accurate way to measure the performance of the cable segment itself is to perform a speed test between the CPE and Server 1, which is connected to the same network segment as the CMTS. This is because the only path data needs to travel over is the coaxial cable segment. The data also must travel across the local CMTS network segment, but it is presumed this segment is of a high bandwidth (FastEthernet or greater) and does not have a high level of congestion.

If for some reason, no server can be connected to the local CMTS network segment, then the next most accurate way to test the performance of the cable segment is to perform a speed test between the CPE and Server 2. This is an accurate measurement as long as there are adequately high speed and uncongested links within the cable service provider's internal network between the CMTS and the CPE.

The most inaccurate way to determine the performance of the cable segment is to perform a speed test between the CPE and a server on the public Internet. This is because there may be congested links in the public Internet between the CPE and the server, or there may be very low speed links in the path between the

CPE and the server on the Internet.

Determining the Download and Upload Rate

It is very important to be able to get an objective measurement of exactly what levels of upload and download throughput are being achieved before any conclusions can be made about whether a performance problem exists in a DOCSIS system.

The easiest way to determine the speed at which uploads and downloads are occurring is to upload or download a large file using FTP or HTTP between a CPE device connected to a cable modem, and a server behind the CMTS. Most FTP and HTTP clients are able to display the speed at which a download or an upload is occurring either during the transfer or once the transfer is complete. The transfer speed seen as a result of the FTP or HTTP operation is typically about 90 percent of the true total throughput attained. This is because the displayed FTP or HTTP transfer speed does not take into account extra IP and DOCSIS overhead that needs to travel between the CPE device and the CMTS.

There are more accurate methods of measuring throughput, for example by using third-party dedicated testing equipment, such as a Netcom Smartbits or a IXIA packet generator, however these systems are not always readily available or easily connected to a production cable network. It is worth noting that if throughput tests are being carried out in a lab environment, then using a dedicated device will reveal much more information than the simple FTP or HTTP download test.

Note: The FTP- or HTTP-based upload and download test is only reliable for testing speeds of around 3 Mbps or less. At higher speeds the processing power of the CPE device, server or Network Interface Cards (NIC) may become a limiting factor in the test. For testing speeds higher than about 3 Mbps, dedicated data throughput testing equipment should be used.

In the following example, a simple FTP download and upload test is performed between a CPE device connected to a cable modem, and an FTP server on the cable service provider's network. The cable modem has downloaded a DOCSIS configuration file that allows a download speed of up to 256 Kbps and an upload speed of up to 64 Kbps. In this test, a 3 Mb file has been placed on the FTP server at IP address 172.17.110.132. The user of the CPE device is given a username and password in order to be able to log into the FTP server so they can download this file from the FTP server, and then upload it back to the FTP server. The command line FTP utility is used to perform the transfer. This utility is available in virtually all versions of Microsoft Windows and UNIX.

A similar test is conducted by having an HTTP web server set up in the service provider's network and performing an HTTP download.




Figure 2

Note:

!--- Comments are in blue.

C:\>**ftp 172.17.110.132**

!--- Initiate the FTP session to the server.

Connected to 172.17.110.132.
220 Solaris FTP server (SunOS 5.6) ready.
User (172.17.110.132:(none)): **anonymous**

!--- Enter the FTP server username.

331 Guest login ok, send your complete e-mail address as password.
Password: **user@samplenetwork.com.au**

!--- Enter the FTP server password.

230 User anonymous logged in.
ftp> **dir**

!--- View the contents of the current directory.

200 PORT command successful.
150 ASCII data connection for /bin/ls (64.104.207.118,1282) (0 bytes).
total 74932
-rw-r--r-- 1 root other 3276800 Oct 10 19:31 cable.txt

!--- A 3 M file that you can download.

226 ASCII Transfer complete.
ftp: 105 bytes received in 0.12 Seconds 2.46 Kbytes/sec.
ftp> **bi**

!--- Turn on Binary File transfer mode.

200 Type set to I.
ftp> **get cable.txt**

!--- Retrieve the file cable.txt and wait for it to download.

200 PORT command successful.
150 Binary data connection for cable.txt (192.168.1.13,3154) (3276800 bytes).
226 Binary Transfer complete.
ftp: 3276800 bytes received in 111.35 Seconds 29.43 Kbytes/sec.

*!--- Download complete. It seems that the download occurred
!--- at 29.43 Kbytes/sec, which equals 235 Kbits/sec. This is about 90 percent of
!--- the allowed 256 Kbps download rate for the modem being tested.*

ftp> **put cable.txt**

!--- Begin uploading the file. You need to make sure you have

```
!--- the correct access in order to upload a file to the FTP server or
!--- you may get an access-denied error.
```

```
200 PORT command successful.
150 Binary data connection for cable.txt (192.168.1.13,3157).
226 Transfer complete.
ftp: 3276800 bytes sent in 432.49 Seconds 7.58 Kbytes/sec.
```

```
!--- Upload Complete. Here you see the upload
!--- occurred at 7.58 Kbytes/sec,
!--- which is equivalent to 60.64 Kbits/sec. This
!--- is about 90 percent of the allowed
!--- 64 Kbps upload rate for the modem being tested.
```

```
ftp> quit
```

```
!--- Exit the FTP client application.
```

```
221 Goodbye.
```

While the FTP transfer is occurring, it is possible to monitor the progress of the test on the CMTS using the **show interface cable X/Y sid Z counters** command where cable X/Y is the cable interface the modem under test is connected to, and Z is the Service ID (SID) number of the modem under test. This command shows how many bytes are being transferred from or to a particular cable modem. For example, if the CPE being tested is behind a cable modem with MAC address **0001.9659.4461**.

First find the SID number of the modem being tested by using the **show cable modem** command. In this case the SID of the cable modem is **5**.

```
uBR7246-VXR# show cable modem 0001.9659.4461
Interface   Prim Online   Timing Rec   QoS CPE IP address   MAC address
           Sid  State    Offset Power
Cable3/0/U0 5   online   1996    0.25  5   2   10.1.1.24   0001.9659.4461
```

While the download or upload is progressing, clear all the packet counters on the CMTS back to zero using the **clear counters** command. Exactly when the counters are cleared, start a stopwatch or timer.

```
uBR7246-VXR# clear counters
```

```
!--- Reset packet counter to zero.
```

```
Clear "show interface" counters on all interfaces [confirm]
```

```
!--- Start the stopwatch when you hit Enter.
```

After the stopwatch or time reads exactly one minute, issue the **show interface cable X/Y sid Z counters** command. It may be best to type the command first and then hit Enter exactly when the timer indicates one minute. The test can be performed over a longer or a shorter period. The longer the test period, the more accurate the result, however make sure that the download or upload does not finish before the stopwatch timer reaches the specified time, otherwise the measurement is inaccurate.

```
uBR7246-VXR# show interface cable 3/0 sid 5 counters
```

!--- Hit enter when stopwatch is at exactly one minute.

```
Sid   Inpackets  Inoctets   Outpackets  Outoctets  Ratelimit  Ratelimit
      4019      257216     3368        1921488    0          DSPktDrop
      149
uBR7246-VXR#
```

In this case download speed is being tested. The output of the **show interface cable X/Y sid Z counter** command indicates that over a period of one minute, 1,921,488 bytes is downloaded by the cable modem. Converting 1,921,488 bytes into bits reveals:

8 bits per byte * 1,921,488 bytes = 15,371,904 bits.

Then, to find the download rate in bits per second, divide this total number of bits downloaded by the time it takes to download them in seconds.

15,371,904 bits / 60 seconds = 256 Kbps.

The download rate in this example is shown to be approximately 256 Kbps, which happens to be the allowed download rate for the cable modem under test.

In order to look at the upload speed using the **show interface cable X/Y sid Z counters** command, the **Inoctets** column should be used to determine the number of bytes sent in the upstream direction from the cable modem.

See the Cisco Broadband Cable Command Reference Guide for more information about the **show interface cable sid counters** command.

Potential Reasons for Poor Performance

Performance Restricted by DOCSIS Configuration File

The first piece of information that needs to be gathered when troubleshooting slow cable modem performance is the prescribed class of service throughput limitations of the cable modem. When a cable modem comes online, it downloads a DOCSIS configuration file that contains operational limits for the cable modem, including the maximum upload and download rates. Under normal circumstances, the cable modem is not allowed to exceed these rates.

Initially it is necessary to identify the MAC address of a cable modem having problems. Taking a modem with MAC address **0050.7366.2223** that is having problems with slow throughput, it is necessary to find out what class of service profile this cable modem is using by executing the **show cable modem <mac-address>** command as seen in the example below.

```
uBR7246-VXR# show cable modem 0050.7366.2223
Interface  Prim Online   Timing Rec   QoS CPE IP address  MAC address
      Sid  State      Offset Power
Cable3/0/U1 1  online    1548   0.75   5   0   10.1.1.10  0050.7366.2223
```

Here it shows this cable modem has a Quality of Service (QoS) profile of 5. In order to find out what downstream and upstream rates this QoS profile corresponds to, use the **show cable qos profile *profile-number*** command, where *profile-number* is the QoS profile of interest.

```
uBR7246-VXR# show cable qos profile 5
ID  Prio  Max      Guarantee  Max      Max      TOS  TOS  Create  B  IP prec.
      upstream  upstream  downstream tx  mask value by  priv  rate
```

		bandwidth	bandwidth	bandwidth	burst				enab	enab
5	0	64000	0	256000	1600	0x0	0x0	cm	no	no

Here it shows that QoS profile 5 corresponds to a service providing 256 Kbps in the downstream and 64 Kbps is the upstream. Any CPE connected to cable modems using QoS profile 5 are not able to exceed these limits. The QoS profile settings are determined by the contents of the DOCSIS configuration files downloaded by cable modems from the provisioning system's TFTP server, therefore QoS profile 5 in the system may not be the same as QoS profile 5 in the example shown above.

If an end user's download and upload performance correlate with the limits shown in their QoS profile, then they are getting the class of service and throughput levels for which the cable modem has been provisioned and configured. The only way to increase the upload and download throughput is to change the DOCSIS configuration file being downloaded by the cable modem to one that has higher throughput limits. See the document entitled Building DOCSIS 1.0 Configuration Files Using Cisco DOCSIS Configurator (registered customers only) for detailed instructions on how to create or modify a DOCSIS configuration file.

Using a Sub-optimal Method for Rate Limiting

When an end user is trying to download data from the Internet at a rate greater than their cable modem's DOCSIS configuration file allows, the CMTS must rate limit the traffic being sent to that user to ensure that the user does not consume more than their allowed share of bandwidth.

Similarly, when an end user tries to upload or send data to the Internet at a rate greater than what the DOCSIS configuration file allows, the cable modem itself should stop the excess traffic from traveling over the cable segment to the CMTS. If the cable modem, for some reason, fails to perform upstream rate limiting properly, then the CMTS explicitly forbids the cable modem from transmitting higher than the allowed rate. This behavior on the CMTS is to ensure that even a cable modem with "hacked" characteristics is unable to subvert the service provider assigned upload rate limits.

The default rate-limiting scheme used by the CMTS monitors the rate of traffic to or from each cable modem over every one-second period. If the cable modem sends or receives more than its per second quota in less than a second, then the CMTS does not allow any more traffic to flow to that cable modem for the rest of the second.

As an example, take a cable modem with a QoS profile allowing a download rate of 512 Kbps. If the cable modem downloads 512 kilobits (64 kilobytes) within the first half of a second, then for the next half of the second, the cable modem is not allowed to download anything. This type of rate limiting behavior may have the effect of a bursty download pattern that seems to stop and start every second or two.

The best downstream rate-limiting scheme to use is the **token bucket** rate-limiting algorithm with **traffic shaping**. This rate-limiting scheme has been optimized to allow for a smooth web browsing experience at a steady rate, while at the same time ensuring that end users are not allowed to exceed the prescribed download rate as specified in the DOCSIS configuration file.

The way this scheme works is to measure the rate at which a cable modem is downloading or uploading data each time a packet is sent to or from the cable modem. If sending or receiving the packet in question causes the modem to exceed its allowed transfer rates, then the packet is buffered or cached in CMTS memory until the CMTS can send the packet without exceeding the downstream bandwidth limit.

Note: If the downstream traffic rate consistently exceeds the allowed downstream rate for the cable modem, then packets are eventually dropped.

By using this smoother method of rate limiting and shaping, most TCP-based Internet applications such as HTTP web browsing and FTP file transfers operate more smoothly and efficiently than when using the default

rate-limiting scheme.

The token-bucket rate-limiting-with-traffic-shaping scheme can be enabled on the downstream path on a cable interface by issuing the following cable interface configuration command:

```
uBR7246-VXR(config-if)# cable downstream rate-limit token-bucket shaping
```

Note: It is highly recommended to enable **token-bucket shaping** on the user's CMTS. This command is supported as of Cisco IOS Software Releases 12.0(5)T1 and 12.1(1)EC1.

The token-bucket with traffic shaping scheme also can be applied to upstream ports, but since it is the responsibility of the cable modems to perform upstream rate limiting, the upstream rate-limiting scheme applied to the CMTS normally will not have any impact on the performance of a system.

```
uBR7246-VXR(config-if)# cable upstream 0 rate-limit token-bucket shaping
```

See the Cisco Broadband Cable Command Reference Guide for more information about the **cable downstream rate-limit** and **cable upstream rate-limit** commands.

Users can view how severely the CMTS is rate limiting traffic to a particular cable modem by using the **show interface cable X/Y sid <Z> counters** command, where cable X/Y is the cable interface to which the cable modem is connected, and Z is the SID number of the modem being observed. This command shows the number of times the CMTS has dropped a downstream packet or refused to allow an upstream packet due to the modem exceeding its allowed throughput limits. If no value is specified for Z, then counter information for all cable modems connected to interface cable X/Y are displayed.

```
uBR7246-VXR# show interface cable 3/0 sid 5 counters
Sid   Inpackets  Inoctets   Outpackets  Outoctets  Ratelimit  Ratelimit
      BReqDrop  DSPktDrop
5     150927    9662206   126529     72008199   0          5681
```

The **Ratelimit DSPktDrop** field shows how many times the CMTS has dropped packets destined for the cable modem due to the modem trying to exceed its allowed downstream throughput.

The **Ratelimit BReqDrop** field shows how many times the CMTS has refused to let a cable modem send a packet in the upstream path due to the modem trying to exceed its allowed upstream throughput. In most circumstances this counter always should remain at 0. If it rises significantly above zero, it may be that the cable modem being observed is not performing upstream rate limiting properly.

Note: The values displayed by the **show interface cable X/Y sid Z counters** command may be reset to zero by issuing the **clear counters** command as seen in the example below.

```
uBR7246-VXR# show interface cable 3/0 sid counters
Sid   Inpackets  Inoctets   Outpackets  Outoctets  Ratelimit  Ratelimit
      BReqDrop  DSPktDrop
1      7          1834       7           1300       0          0
2     2052      549150     0            0          0          0
3      2          1244       2            708        0          0
4      2          1244       2            714        0          0
5    160158   10253220   134294      76423270   0          6023
6      2          1244       2            712        0          0
7      9          1906       4            858        0          0
9      6          1076       3            483        0          0
12     616       165424     0            0          0          0
uBR7246-VXR# clear counters
Clear "show interface" counters on all interfaces [confirm] <press enter here>
uBR7246-VXR# show interface cable 3/0 sid counters
Sid   Inpackets  Inoctets   Outpackets  Outoctets  Ratelimit  Ratelimit
```

					BWReqDrop	DSPktDrop
1	0	0	0	0	0	0
2	0	0	0	0	0	0
3	0	0	0	0	0	0
4	0	0	0	0	0	0
5	111	7104	92	52728	0	6
6	0	0	0	0	0	0
7	0	0	0	0	0	0
9	0	0	0	0	0	0
12	0	0	0	0	0	0

See the Cisco Broadband Cable Command Reference Guide for more information about the **show interface cable sid counters** command.

Upstream Channel Congestion

The upstream channel is normally the most precious resource in a cable system. Currently, most cable service providers use a 1.6 MHz channel width and Quadrature Phase Shift Keying (QPSK) modulation in the upstream path. This equates to approximately 2.5 Mbps in total available upstream bandwidth for all users connected to the one upstream channel. It is important to ensure that the upstream channel does not become over used or congested, otherwise all users on that upstream segment suffer poor performance.

The upstream usage for a particular upstream port can be obtained by executing the CMTS command **show interface cable X/Y upstream <Z>**, where *cable X/Y* is the downstream interface number and *Z* is the upstream port number. If *Z* is omitted, information for all upstreams on interface cable *X/Y* will be displayed. See the Cisco Broadband Cable Command Reference Guide for more information about the **show interface cable upstream** command.

```
uBR7246-VXR# show interface cable 6/0 upstream 0
Cable6/0: Upstream 0 is up
Received 71941 broadcasts, 27234 multicasts, 8987489 unicasts
0 discards, 140354 errors, 0 unknown protocol
9086664 packets input, 4394 uncorrectable
122628 noise, 0 microreflections
Total Modems On This Upstream Channel : 359 (354 active)
Default MAC scheduler
Queue[Rng Polls] 0/64, fifo queueing, 0 drops
Queue[Cont Mslots] 0/104, fifo queueing, 0 drops
Queue[CIR Grants] 0/64, fair queueing, 0 drops
Queue[BE Grants] 0/64, fair queueing, 0 drops
Queue[Grant Shpr] 0/64, calendar queueing, 0 drops
Reserved slot table currently has 0 CBR entries
Req IEs 64609697, Req/Data IEs 0
Init Mtn IEs 521851, Stn Mtn IEs 569985
Long Grant IEs 2781600, Short Grant IEs 2067668
Avg upstream channel utilization : 18%
Avg percent contention slots : 77%
Avg percent initial ranging slots : 2%
Avg percent minislots lost on late MAPs : 0%
Total channel bw reserved 37858000 bps
CIR admission control not enforced
Admission requests rejected 0
Current minislot count : 7301855 Flag: 0
Scheduled minislot count : 7301952 Flag: 0
```

On the upstream port seen in the example, the upstream usage is currently 18 percent and there are 359 modems connected to this upstream.

If upstream channel usage is consistently above 75 percent during the peak usage time, end users begin to suffer issues such as latency, slower "ping" times, and a generally slower Internet experience. If upstream

channel usage is constantly above 90 percent during the peak usage time, end users experience an extremely poor level of service because a large portion of the end user's upstream data will have to be delayed or discarded.

Upstream channel usage changes during the day as different users have an opportunity to use their cable modem, so it is important to monitor the upstream usage during the busiest times of the day rather than at low usage times.

Ways of relieving upstream congestion include:

- **Reducing the number of cable modems per upstream** If there are too many cable modems connected to a particular upstream, or if users on a particular upstream are heavy users of upstream bandwidth, the best solution is to move some users on the congested upstream port to an under used upstream port, or to a completely new upstream port. This is typically accomplished by moving a fiber node from one upstream combining group to another, or splitting an upstream combining group into two separate combining groups. For more information, refer to [What is the Maximum Number of Users per CMTS](#).
- **Increasing the upstream channel width** This involves a **rigorous and thorough** analysis of the upstream spectrum to find a wide enough band with adequate Signal-to-Noise Ratio (SNR) characteristics to support the increased channel width. The upstream channel width should not be changed without careful planning because this change potentially can affect other services in the user's cable system. The upstream channel width may be changed by using the cable interface command **cable upstream Z channel-width <new-channel-width>** where Z is the upstream port number and new channel width is one of 200000, 400000, 800000, 1600000 (the default) or 3200000. An example follows.

```
uBR7246-VXR(config-if)# cable upstream 0 channel-width 3200000
```

See the Cisco Broadband Cable Command Reference Guide for more information about the **show interface cable upstream** command.

- **Changing the upstream digital modulation scheme to 16-Quadrature Amplitude Modulation (QAM)** Once again, this requires a **rigorous and thorough** analysis of the upstream spectrum in order to verify whether there is a frequency band in the upstream available that can support 16-QAM modulation. If this analysis is not performed properly, there is a risk that performance is further decreased or a complete upstream outage may occur. The upstream modulation scheme may be changed by creating an upstream modulation profile that uses 16-QAM modulation and then applying that to an upstream port. An example follows.

```
uBR7246-VXR(config)# cable modulation-profile 2 mix
```

```
!--- Create an optimized 16-qam/qpsk modulation profile.
```

```
uBR7246-VXR(config)# interface cable 6/0
```

```
uBR7246-VXR(config-if)# cable upstream 0 modulation-profile 2
```

See the Cisco Broadband Cable Command Reference Guide for more information about the **cable modulation-profile** and **cable upstream modulation-profile** commands. See also [Configuring Cable Modulation Profiles on Cisco's Cable Modem Termination Systems](#).

- **Reducing the allowed upstream throughput per cable modem** By reducing the maximum upstream transmit rate in the appropriate DOCSIS configuration files, cable modem users are not able to transmit at as high a rate in the upstream direction and upstream congestion is relieved. The negative aspect of this course of action is that cable modem users are limited to a slower class of service. See [Building DOCSIS 1.0 Configuration Files Using Cisco DOCSIS Configurator \(registered customers only\)](#).

Note: The measures discussed in this section do not significantly increase the performance of an already uncongested system.

Downstream Channel Congestion

The downstream channel has significantly more bandwidth to share than an individual upstream channel, therefore the downstream is not usually as subject to congestion as the upstream. Still, more users typically share a downstream channel than any single upstream channel, so if the downstream channel becomes congested, all users connected to the downstream segment experience reduced performance.

The following table shows the total available downstream bandwidth associated with the four possible downstream modulation schemes available in DOCSIS systems.

Downstream Modulation Scheme	Available Downstream Bandwidth
64-QAM North American DOCSIS	27 Mbps
256-QAM North American DOCSIS	38 Mbps
64-QAM Euro DOCSIS	38 Mbps
256-QAM Euro DOCSIS	54 Mbps

The majority of DOCSIS cable systems currently deploy 64-QAM North American DOCSIS and therefore have 27 Mbps available per downstream channel.

Downstream channel usage can be determined by issuing the **show interface cable X/Y** command, where *cable X/Y* is the cable interface being observed. The displayed output rate in bits per second should be compared to the available downstream bandwidth as seen in the table above.

In the following example, an interface using North American DOCSIS and 64-QAM digital modulation is analyzed.

```
uBR7246-VXR# show interface cable 3/0
Cable3/0 is up, line protocol is up
  Hardware is BCM3210 ASIC, address is 0005.5fed.dca4 (bia 0005.5fed.dca4)
  Internet address is 10.1.1.1/24
  MTU 1500 bytes, BW 27000 Kbit, DLY 1000 usec,
    reliability 255/255, txload 9/255, rxload 5/255
  Encapsulation MCNS, loopback not set
  Keepalive not set
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, output 00:00:00, output hang never
  Last clearing of "show interface" counters 00:45:01
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue :0/40 (size/max)
  5 minute input rate 587000 bits/sec, 228 packets/sec
  5 minute output rate 996000 bits/sec, 239 packets/sec
    85560 packets input, 8402862 bytes, 0 no buffer
    Received 1013 broadcasts, 0 runts, 0 giants, 0 throttles
    247 input errors, 35 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
    65912 packets output, 38168842 bytes, 0 underruns
    0 output errors, 0 collisions, 0 interface resets
    0 output buffer failures, 0 output buffers swapped out
```

The first component of this output to note is the bandwidth of the interface indicated by the **BW** parameter. In Cisco IOS Software Releases 12.1(8)EC and later, this value is adjusted automatically according to the downstream modulation scheme and version of DOCSIS being used. In revisions earlier than Cisco IOS Software Release 12.1(8)EC, this value must be configured manually using the cable interface command **bandwidth <bandwidth-in-kilo-bits-per-second>** or otherwise it remains at the default value of 27000 Kbps.

The second component to note is the transmission load as indicated by the **txload** parameter. This parameter gives a metric out of 255 where 0/255 means that no traffic is flowing in the downstream direction to 255/255, which means data is traveling in the downstream at the maximum possible rate (in this case at 27000 Kbps). If this parameter is running consistently at greater than approximately 75 percent during the peak usage time (for example, greater than 191/255), end users start to experience slower Internet access and higher latency.

The third component to note is the **output rate**, which shows the average downstream throughput rate in bits per second. If this number consistently exceeds approximately 75 percent of the available downstream bandwidth during the peak usage time, end users start to experience slower Internet access and higher latency.

By default, these statistics are calculated over a five-minute moving average. (Refer to Understanding the Definition of bits per second (bits/sec) from the show interfaces Command Output for details of how the average is calculated.) The period over which this average is calculated can be reduced to as little as 30 seconds by issuing the **cable interface** command **load-interval 30**. By lowering this period to 30 seconds, a more accurate an up-to-date value is calculated for each of the parameters discussed in this section.

Downstream channel usage changes during the day as different users have an opportunity to use their cable modem, so it is important to monitor downstream usage during the busiest times of the day rather than at low usage times.

Ways of relieving downstream congestion include:

- **Reducing the number of cable modems per downstream** If there are too many cable modems connected to a particular downstream, or if users on a particular downstream are heavy users of downstream bandwidth, then the best solution is to move some users on the congested downstream channel to another downstream channel. This typically is accomplished by splitting a group of downstream fiber nodes associated with the downstream into two separate groups and assigning each of the new groups separate downstream channels. Refer to What is the Maximum Number of Users per CMTS.
- **Changing the downstream digital modulation scheme to 256-QAM** This action requires a **rigorous and thorough** analysis of the downstream spectrum in order to verify whether the system can support a 256-QAM signal. If this analysis is not performed properly, there is a risk that performance will be further decreased or a complete downstream outage may occur. The downstream modulation scheme may be changed by issuing the cable interface command as seen below.

```
uBR7246-VXR(config-if)# cable downstream modulation 256qam
```

See the Cisco Broadband Cable Command Reference Guide for more information about the **cable downstream modulation** command.

- **Reducing the allowed downstream throughput per cable modem** By reducing the maximum downstream transmit rate in the appropriate DOCSIS configuration files, cable modem users are not able to download at as high a rate in the downstream direction and downstream congestion is relieved. The negative aspect of this course of action is that cable modem users are limited to a slower class of service. Refer to Building DOCSIS 1.0 Configuration Files Using Cisco DOCSIS Configurator (registered customers only) .

Note: The measures discussed in this section do not significantly increase the performance of an

already uncongested system.

Backhaul Network or Internet Congestion

In some cases, performance problems may not be a result of issues on the cable plant or the CMTS, but may be related to congestion or problems in the backhaul network that the CMTS uses to connect to the Internet, or within parts of the Internet itself.

The easiest way to determine if backhaul network congestion is a problem is to connect a workstation to the same network segment as the CMTS and try to browse the same websites as end users behind cable modems are trying to reach. If performance is still slow, there is a performance problem in the network not related to the CMTS or the cable segment. If performance from the local CMTS network segment is significantly better than for users connected to cable modems, then focus efforts back on the CMTS and the cable segment.

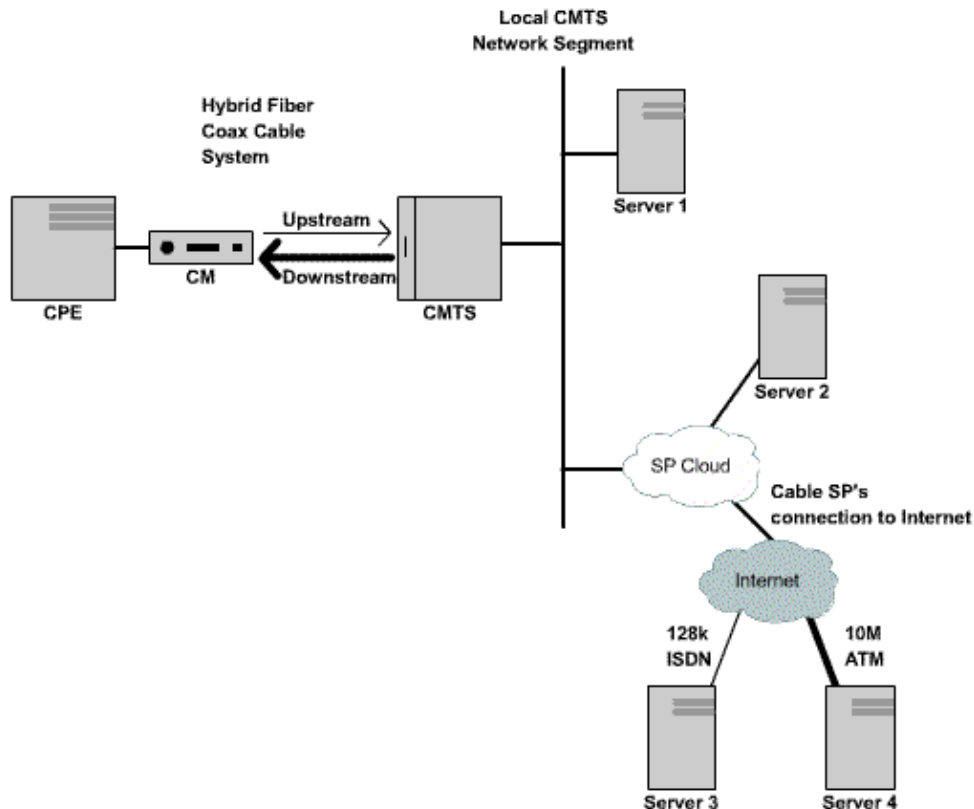


Figure 3

In the above network, if Server 1, which is connected to the same network segment as the CMTS, is getting slow performance when browsing the Internet, then the source of the problem is not the CMTS. Instead, the bottleneck or performance issue somewhere else. In order to determine where the problem is, performance tests are carried out between Server 1 and various other servers within the Internet Service Provider (ISP) network and the public Internet.

Noise and Errors on the Cable Plant

If there is an excessive amount of noise or ingress in a cable system, then packets between cable modems and the CMTS can be corrupted and lost. This can lead to a significant degradation in performance.

Aside from a degradation in performance and throughput, some of the prime indicators of noise or radio frequency (RF) issues include:

- Cable modems sporadically dropping offline or getting stuck in the init(r1) or init(r2) states.
- A low estimated SNR as seen in the output of a **show controller cable X/Y upstream Z**, where *cable X/Y* is the cable interface being observed and *Z* is the upstream port being observed. The DOCSIS specification requires a Carrier-to-Noise Ratio (CNR) of at least 25 dB for all upstream signals. This equates to an SNR of approximately 29 dB. The Cisco CMTS is able to detect coherently QPSK upstream signals at much worse SNR levels, however all cable service providers should strive to meet the DOCSIS CNR requirements in their system. A sample **show controller cable X/Y upstream Z** output is shown below.

```
uBR7246-VXR# show controller cable 6/0 upstream 0
Cable6/0 Upstream 0 is up
  Frequency 25.200 MHz, Channel Width 1.600 MHz, QPSK Symbol Rate 1.280 Msps
  Spectrum Group is overridden
  SNR 28.6280 dB
  Nominal Input Power Level 0 dBmV, Tx Timing Offset 6446
  Ranging Backoff automatic (Start 0, End 3)
  Ranging Insertion Interval automatic (102 ms)
  Tx Backoff Start 0, Tx Backoff End 4
  Modulation Profile Group 1
  Concatenation is enabled
  part_id=0x3137, rev_id=0x03, rev2_id=0xFF
  nb_agc_thr=0x0000, nb_agc_nom=0x0000
  Range Load Reg Size=0x58
  Request Load Reg Size=0x0E
  Minislot Size in number of Timebase Ticks is = 8
  Minislot Size in Symbols = 64
  Bandwidth Requests = 0x37EB54
  Piggyback Requests = 0x11D75E
  Invalid BW Requests= 0x102
  Minislots Requested= 0x65B74A2
  Minislots Granted = 0x65B74A2
  Minislot Size in Bytes = 16
  Map Advance (Dynamic) : 2809 usecs
  UCD Count = 23068
```

In the example above, the estimated SNR reading is 28.628dB. This is adequate for QPSK upstream operation. Note that the SNR figure given in the output of this command is only an estimate and is no substitute for an SNR figure derived from a spectrum analyzer or other appropriate testing equipment. See the Cisco Broadband Cable Command Reference Guide for more information about the **show controllers cable upstream spectrum** command.

- A quickly incrementing number of **Corr Forward Error Correction (FEC)** and **Uncorr FEC** errors in the output of a **show cable hop** command. **Corr FEC** errors indicates data that was corrupted by upstream noise but was able to be recovered. **Uncorr FEC** errors indicate data that was corrupted by upstream noise and was not able to be recovered resulting in lost data and slower performance. A sample output from the **show cable hop** command is shown below.

```
uBR7246-VXR# show cable hop cable 3/0
Upstream Port      Port      Status      Poll Rate      Missed Count      Min Sample      Missed Pcnt      Hop Thres      Hop Period      Corr FEC      Uncorr FEC
Port              Status      (ms)        Poll          Count           Sample          Pcnt            Pcnt           (sec)          Errors       Errors
Cable3/0/U0      25.200 Mhz 34          * * *         set to fixed    frequency      * * *          196            55
Cable3/0/U1      25.200 Mhz 34          * * *         set to fixed    frequency      * * *          1655           160
Cable3/0/U2      25.200 Mhz 34          * * *         set to fixed    frequency      * * *          76525          9790
Cable3/0/U3      25.200 Mhz 34          * * *         set to fixed    frequency      * * *          501            77
Cable3/0/U4      admin down 34          * * *         interface is    down           * * *          0              0
Cable3/0/U5      admin down 34          * * *         interface is    down           * * *          0              0
```

In the example above, each active upstream port on cable 3/0 seems to have experienced packet loss due to noise. Upstream port 0 seems to be the least affected and upstream port 2 seems to be the most heavily affected. The important factor to note is how quickly the FEC Errors are incrementing rather

than the total number of errors. See the Cisco Broadband Cable Command Reference Guide for more information about the **show cable hop** command.

- A high number of "flap" events in the output of a **show cable flap-list** command. The flap statistics most pertinent to possible RF or noise problems are the **Miss** column, which indicates missed ranging requests, and the **P-Adj** column, which indicates rapidly varying upstream power levels. A sample output from the **show cable flap-list** command is shown below.

```
uBR7246-VXR# show cable flap-list
MAC Address      Upstream  Ins   Hit   Miss  CRC   P-Adj  Flap  Time
0000.d025.1b99  Cable3/0/U0  23   58   30   0   *27   77   Oct 23 03:08:23
0002.ddfa.0aa5  Cable3/0/U1  5    518  1260 0   0     131  Oct 23 03:09:43
0001.e659.43bd  Cable3/0/U1  541  342  1467 0   0     746  Oct 23 03:09:17
0001.7659.44c7  Cable3/0/U1  0    694  0     0   1     1    Oct 23 01:44:23
0050.9366.22d3  Cable3/0/U1  0    708  0     0   1     1    Oct 23 01:38:14
0001.f659.44e7  Cable3/0/U1  0    701  0     0   1     1    Oct 23 02:25:11
```

- Cable modems displaying a "*" or a "!----" in the output of a **show cable modem** or **show cable flap-list** command. A "*" indicates a cable modem that is rapidly varying its upstream power levels. This is indicative of a poor connection to the cable plant, a faulty reverse path amplifier, or rapidly changing cable plant attenuation due to temperature or other environmental effects. A "!----" indicates a cable modem that has reached its maximum upstream power level. This is indicative of too much attenuation between the cable modem and the CMTS, or a poor connection between the cable modem and the cable plant. A sample output from the **show cable modem** command is shown below.

```
uBR7246-VXR# show cable modem
Interface  Prim Online  Timing      Rec      QoS CPE IP address  MAC address
          Sid  State   Offset      Power
Cable3/0/U1 1  online  1549  !---- -1.00  5  0  10.1.1.10  005a.73f6.2213
Cable3/0/U0 2  online  1980   0.75  5  0  10.1.1.16  009b.96e7.3820
Cable3/0/U0 3  online  1981   *0.75  5  0  10.1.1.18  009c.96d7.3831
Cable3/0/U1 4  online  1924   0.25  5  0  10.1.1.24  000d.96c9.4441
Cable3/0/U1 5  online  1925   0.50  5  0  10.1.1.13  000e.96b9.4457
```

In the example above, the cable modem with MAC address 005a.73f6.2213 is transmitting at its maximum output power. This results in that modem not being able to transmit at the correct level. Consequently, this modem's upstream transmissions are not heard as clearly as transmissions from other modems. The cable modem with MAC address 009c.96d7.3831 has a rapidly varying power output due to varying cable system attenuation. See the Cisco Broadband Cable Command Reference Guide for more information about the **show cable modem** and **show cable flap-list** commands.

Note: More details about identifying and resolving RF noise issues can be found in Determining RF or Configuration Issues On the CMTS and Connecting the Cisco uBR7200 Series Router to the Cable Headend.

High CPU Usage on the CMTS

In some circumstances a Cisco CMTS can become overloaded due to a sub-optimal configuration, over use of certain management functions, or a very high number of packets being routed by the CMTS.

The best way to determine the CPU usage of a Cisco CMTS is to execute the **show process cpu** command. The current CPU usage is indicated on the first line of the output of the command.

In the lines of output shown below the first line, each process running on the CMTS is shown along with the portion of the CPU being used by that process. This section of the **show process cpu** output is useful for determining if one particular process or function is the cause of high CMTS CPU.

```
uBR7246-VXR# show process cpu
CPU utilization for five seconds: 45%/21%; one minute: 45%; five minutes: 31%
PID Runtime(ms) Invoked uSecs 5Sec 1Min 5Min TTY Process
1 12 9220 1 0.00% 0.00% 0.00% 0 Load Meter
```

2	69816	18276677	3	21.79%	22.10%	9.58%	2 Virtual Exec
3	36368	5556	6545	0.00%	0.06%	0.05%	0 Check heaps
4	0	1	0	0.00%	0.00%	0.00%	0 Chunk Manager
5	96	1436	66	0.00%	0.00%	0.00%	0 Pool Manager
6	0	2	0	0.00%	0.00%	0.00%	0 Timers
7	0	2	0	0.00%	0.00%	0.00%	0 Serial Backgroun
8	0	1	0	0.00%	0.00%	0.00%	0 CMTS ping
9	17020	101889	167	0.00%	0.00%	0.00%	0 EnvMon
10	0	1	0	0.00%	0.00%	0.00%	0 OIR Handler
.							
<snip>							
.							
89	3304	81013	40	0.00%	0.00%	0.00%	0 PIM Process
90	12	769	15	0.00%	0.00%	0.00%	0 CEF Scanner
92	0	385	0	0.00%	0.00%	0.00%	0 DHCPD Timer
93	40	13058	3	0.00%	0.00%	0.00%	0 DHCPD Database

In the example above, the current CPU load on the CMTS is 45 percent/21 percent. This means that the total CPU usage is at 45 percent of the capacity of the system. In addition 21 percent of the CPU is being used to service interrupts. This second figure typically equates to the portion of the CPU being used to route packets and switch traffic through the CMTS.

If the five minutes CPU usage consistently is more than 80 percent during the peak usage time in the system, then end users may start to experience slower performance and increased latency. If the five minutes CPU usage is constantly more than 95 percent during the peak usage time, then take urgent action to ensure that the CMTS remains in a stable state.

Common strategies for reducing high CPU usage on the CMTS include:

- Upgrading to Cisco IOS Software Release 12.1(9)EC or later, activating the global configuration command **ip cef**, and making sure no interfaces on the CMTS have the command **no ip route-cache** configured. This typically leads to a 10 percent to 15 percent reduction in traffic-related CPU usage. Make sure that all of these steps are taken in conjunction.
- Making sure that Simple Network Management Protocol (SNMP) management stations are not being too aggressive in polling the CMTS. This leads to a high CPU usage in the IP SNMP process.
- Not running the **show tech** command several times in succession. This leads to an artificially high CPU usage in the Virtual Exec Process.
- Making sure that no **debug** commands are running on the CMTS.

For more information about High CPU usage on Cisco routers, including Cisco CMTS products, refer to [Troubleshooting High CPU Utilization on Cisco Routers](#).

Under Powered or Malconfigured CPE Equipment

In many cases, the cause of slow access to a cable network is a problem in the end user's CPE equipment. If only one or a handful of users are experiencing slow throughput, and the rest of the user population are experiencing no problem, then this is a strong indication that there may be a unique problem within that user's environment.

- **Under powered or overloaded CPE** If the end users complaining of difficulties are using antiquated CPE equipment, or equipment that may not be powerful enough to run their chosen operating system or Internet access software, then this end user will have difficulties. The only resolution if this is the case is for the end user to upgrade their CPE equipment.
- **Firewall or performance measurement software** If the end user is running any firewall, network performance measurement, or other similar software, a good troubleshooting step is to have the user turn off this software to see if it has any effect on performance. Often, these kinds of software can have a negative impact on performance.

- **Misconfigured TCP/IP settings** Most service providers require that end users have their CPE equipment acquire an IP address, network mask, default gateway, and DNS servers by way of Dynamic Host Configuration Protocol (DHCP). Ensure that any end users experiencing problems have their CPE devices configured to use DHCP to acquire all of these parameters.

If an end user claims to have none of the problems listed above, then confirm that the end user is not exceeding their maximum download or upload rate as per the sections above.

Conclusion

A DOCSIS cable network is a sophisticated system requiring proper planning and maintenance. Most performance issues in DOCSIS cable systems are a direct result of proper planning and maintenance not being performed. In today's Internet access market, where there are a variety of broadband Internet access alternatives, it is important that cable service providers quickly address any performance or congestion issues in their system before the problems become significant enough for end users to be affected noticeably, and, consequently, consider an alternative means of broadband access.

Related Information

- **Troubleshooting uBR Cable Modems Not Coming Online**
 - **Beginners FAQ for uBR900 Series Cable Modem End Users**
 - **Cisco Broadband Cable Command Reference Guide**
 - **Determining RF or Configuration Issues On the CMTS**
 - **Connecting the Cisco uBR7200 Series Router to the Cable Headend**
 - **Troubleshooting High CPU Utilization on Cisco Routers**
 - **Technical Support – Cisco Systems**
-

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

© 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Apr 24, 2008

Document ID: 12551
