

Understanding Data Throughput in a DOCSIS World

Document ID: 19220

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

Background Information

- Bits, Bytes, and Baud
- What is Throughput?
- Throughput Calculations

Limiting Factors

- Downstream Performance – MAPs
- Upstream Performance – DOCSIS Latency
- TCP or UDP?
- Windows TCP/IP Stack

Performance Improvement Factors

- Throughput Determination
- Increasing Access Speed
- Channel Width and Modulation
- Interleaving Effect
- Dynamic MAP Advance
- Concatenation and Fragmentation Effect
- Single Modem Speeds
- DOCSIS 2.0 Benefits
- Other Factors

Verifying the Throughput

Summary

Conclusion

Related Information

Introduction

Before you attempt to measure the performance of a cable network, there are some limiting factors that you should take into consideration. To design and deploy a highly available and reliable network, you must establish an understanding of basic principles and measurement parameters of cable network performance. This document presents some of those limiting factors and then discusses how to actually optimize and qualify throughput and availability on your deployed system.

Prerequisites

Requirements

Readers of this document should have knowledge of these topics:

- Data-over-Cable Service Interface Specification (DOCSIS)
- Radio Frequency (RF) technologies

- Cisco IOS® software command–line interface (CLI)

Components Used

This document is not restricted to specific software or hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Conventions

For more information on document conventions, refer to the Cisco Technical Tips Conventions.

Background Information

Bits, Bytes, and Baud


This section explains the differences between bits, bytes, and baud. The word *bit* is a contraction of **B**inary **digiT**, and it is usually symbolized by a lower case b. A binary digit indicates two electronic states: an on state or an off state, sometimes referred to as *1s* or *0s*.

A *byte* is symbolized by an upper case B, and it is usually 8 bits in length. A byte could be more than 8 bits, so an 8–bit word is more precisely called an *octet*. Also, there are two *nibbles* in a byte. A nibble is defined as a 4–bit word, which is half of a byte.

Bit rate, or throughput, is measured in bits per second (bps), and it is associated with the speed of a signal through a given medium. For example, this signal could be a baseband digital signal or, perhaps, a modulated analog signal that is conditioned to represent a digital signal.

One type of modulated analog signal is Quadrature Phase Shift Keying (QPSK). This is a modulation technique that manipulates the phase of the signal by 90 degrees to create four different signatures, as shown in Figure 1. These signatures are called *symbols*, and their rate is referred to as *baud*. Baud equates to symbols per second.

Figure 1 QPSK Diagram



QPSK signals have four different symbols; four is equal to 2^2 . The exponent gives the theoretical number of bits per cycle (symbol) that can be represented, which equals 2 in this case. The four symbols represent the binary numbers 00, 01, 10, and 11. Therefore, if a symbol rate of 2.56 Msymbols/s is used to transport a QPSK carrier, then it would be referred to as 2.56 Mbaud and the theoretical bit rate would be $2.56 \text{ Msymbols/s} \times 2 \text{ bits/symbol} = 5.12 \text{ Mbps}$. This is further explained later in this document.

You might also be familiar with the term *packets per second* (PPS). This is a way to qualify the throughput of a device based on packets, regardless of whether that packet contains a 64–byte or a 1518–byte Ethernet

frame. Sometimes the bottleneck of the network is the power of the CPU to process a certain amount of PPS and is not necessarily the total bps.

What is Throughput?

Data throughput begins with a calculation of a *theoretical* maximum throughput, then concludes with *effective* throughput. Effective throughput available to subscribers of a service will always be less than the theoretical maximum, and it is what you should try to calculate.

Throughput is based on many factors:

- total number of users
- bottleneck speed
- type of services accessed
- cache and proxy server usage
- MAC layer efficiency
- noise and errors on the cable plant
- many other factors

The goal of this document is to explain how to optimize throughput and availability in a DOCSIS environment and to explain the inherent protocol limitations that affect performance. If you want to test or troubleshoot performance issues, refer to Troubleshooting Slow Performance in Cable Modem Networks. For guidelines on the maximum number of recommended users on an upstream (US) or downstream (DS) port, refer to What is the Maximum Number of Users per CMTS?.

Legacy cable networks rely on polling or carrier sense multiple access collision detect (CSMA/CD) as the MAC protocol. Today's DOCSIS modems rely on a reservation scheme where the modems request a time to transmit and the CMTS grants time slots based on availability. Cable modems are assigned a Service ID (SID) that is mapped to class of service (CoS) or quality of service (QoS) parameters.

In a bursty, time division multiplex access (TDMA) network, you must limit the number of total cable modems (CMs) that can simultaneously transmit, if you want to guarantee a certain amount of access speed to all requesting users. The total number of simultaneous users is based on a Poisson distribution, which is a statistical probability algorithm.

Traffic engineering, as a statistic used in telephony-based networks, signifies about 10 percent peak usage. This calculation is beyond the scope of this document. Data traffic, on the other hand, is different than voice traffic; and it will change when users become more computer savvy or when Voice over IP (VoIP) and Video on Demand (VoD) services are more available. For simplicity, assume 50 percent peak users \times 20 percent of those users actually downloading at the same time. This would equal 10 percent peak usage also.

All simultaneous users contend for the US and DS access. Many modems can be active for the initial polling, but only one modem can be active in the US at any given instant in time. This is good in terms of noise contribution, because only one modem at a time adds its noise complement to the overall effect.

An inherent limitation with the current standard is that some throughput is necessary for maintenance and provisioning, when many modems are tied to a single cable modem termination system (CMTS). This is taken away from the actual payload for active customers. This is known as *keepalive polling*, which usually occurs once every 20 seconds for DOCSIS but could occur more often. Also, per-modem US speeds can be limited by the Request-and-Grant mechanisms, as explained later in this document.


Note: Remember that references to file size are in bytes made up of 8 bits. Thus, 128 kbps equals 16 KBps. Likewise, 1 MB is actually equal to 1,048,576 bytes, not 1 million bytes, because binary numbers always yield a number that is a power of 2. A 5 MB file is actually $5 \times 8 \times 1,048,576 = 41.94$ Mb and could be longer

to download than anticipated.

Throughput Calculations

Assume that a CMTS card that has one DS and six US ports is in use. The one DS port is split to feed about 12 nodes. Half of this network is shown in Figure 2.

Figure 2 Network Layout



- 500 homes per node \times 80 percent cable take-rate \times 20 percent modem take-rate = 80 modems per node
- 12 nodes \times 80 modems per node = 960 modems per DS port

Note: Many multiple service operators (MSOs) now quantify their systems as Households Passed (HHP) per node. This is the only constant in today's architectures, where you might have direct broadcast satellite (DBS) subscribers buying high speed data (HSD) service or only telephony without video service.


Note: The US signal from each one of those nodes will probably be combined on a 2:1 ratio so that two nodes feed one US port.

- 6 US ports \times 2 nodes per US = 12 nodes
- 80 modems per node \times 2 nodes per US = 160 modems per US port.

Downstream

DS symbol rate = 5.057 Msymbols/s or Mbaud. A filter roll-off (alpha) of about 18 percent gives $5.057 \times (1 + 0.18) = \sim 6$ MHz wide haystack, as shown in Figure 3.

Figure 3 Digital Haystack



If 64-QAM is used, then $64 = 2$ to the 6^{th} power (2^6). The exponent of 6 means 6 bits per symbol for 64-QAM; this gives $5.057 \times 6 = 30.3$ Mbps. After the entire forward error correction (FEC) and Motion Picture Experts Group (MPEG) overhead is calculated, this leaves about 28 Mbps for payload. This payload is

further reduced, because it is also shared with DOCSIS signaling.

Note: ITU–J.83 Annex B indicates Reed–Solomon FEC with a 128/122 code, which means 6 symbols of overhead for every 128 symbols, hence $6 / 128 = 4.7$ percent. Trellis coding is 1 byte for every 15 bytes, for 64–QAM, and 1 byte per 20 bytes, for 256–QAM. This is 6.7 percent and 5 percent, respectively. MPEG–2 is made up of 188–byte packets with 4 bytes of overhead (sometimes 5 bytes), which gives $4.5 / 188 = 2.4$ percent. This is why you will see the speed listed as 27 Mbps, for 64–QAM, and as 38 Mbps, for 256–QAM. Remember that Ethernet packets also have 18 bytes of overhead, whether for a 1500–byte packet or a 46–byte packet. There are 6 bytes of DOCSIS overhead and IP overhead also, which could be a total of about 1.1 to 2.8 percent extra overhead and could add another possible 2 percent of overhead for DOCSIS MAP traffic. Actual tested speeds for 64–QAM have been closer to 26 Mbps.

In the very unlikely event that all 960 modems download data at precisely the same time, they will each get only about 28 kbps. If you look at a more realistic scenario and assume a 10 percent peak usage, you get a theoretical throughput of 280 kbps as a worst–case scenario during the busiest time. If only one customer is online, the customer would theoretically get 26 Mbps; but the US acknowledgments that must be transmitted for TCP limits the DS throughput, and other bottlenecks become apparent (such as the PC or the Network Interface Card [NIC]). In reality, the cable company will rate–limit this down to 1 or 2 Mbps, so as not to create a perception of available throughput that will never be achievable when more subscribers sign up.

Upstream

The DOCSIS US modulation of QPSK at 2 bits/symbol gives about 2.56 Mbps. This is calculated from the symbol rate of $1.28 \text{ Msymbols/s} \times 2 \text{ bits/symbol}$. The filter alpha is 25 percent, which gives a bandwidth (BW) of $1.28 \times (1 + 0.25) = 1.6 \text{ MHz}$. Subtract about 8 percent for FEC, if it is used. There is also approximately 5 to 10 percent of overhead for maintenance, reserved time slots for contention, and acknowledgments (acks). Thus, there is about 2.2 Mbps, which is shared amongst 160 potential customers per US port.

Note: DOCSIS layer overhead = 6 bytes per 64–byte to 1518–byte Ethernet frame (could be 1522 bytes, if VLAN tagging is used). This also depends on the maximum burst size and whether concatenation or fragmentation are used.

- US FEC is variable: $\sim 128 / 1518$ or $\sim 12 / 64 = \sim 8$ or ~ 18 percent. Approximately 10 percent is used for maintenance, reserved time slots for contention, and acks.
- BPI security or Extended Headers = 0 to 240 bytes (usually 3 to 7).
- Preamble = 9 to 20 bytes.
- Guardtime ≥ 5 symbols = ~ 2 bytes.

Assuming 10 percent peak usage, this gives $2.2 \text{ Mbps} / (160 \times 0.1) = 137.5 \text{ kbps}$ as the worst–case payload per subscriber. For typical residential data usage (for example, web browsing) you probably do not need as much US throughput as DS. This speed might be sufficient for residential usage, but it is not sufficient for commercial service deployments.

Limiting Factors

There is a plethora of limiting factors that affect real data throughput. These range from the Request–and–Grant cycle to DS interleaving. Understanding the limitations will aid in expectations and optimization.

Downstream Performance – MAPs

The transmission of MAP messages sent to modems reduce DS throughput. A MAP of time is sent on the DS, to allow modems to request time for US transmission. If a MAP is sent every 2 ms, it adds up to $1 / 0.002s = 500$ MAPs/s. If the MAP takes up 64 bytes, that equals $64 \text{ bytes} \times 8 \text{ bits per byte} \times 500 \text{ MAPs/s} = 256 \text{ kbps}$. If you have six US ports and one DS port on a single blade in the CMTS chassis, this is $6 \times 256000 \text{ bps} = \sim 1.5$ Mbps of DS throughput used to support all of the modems MAP messages. This assumes that the MAP is 64 bytes and that it is actually sent every 2 ms. In reality, MAP sizes could be slightly larger, depending on the modulation scheme and the amount of US bandwidth that is used. This could easily be 3 to 10 percent overhead. Further, there are other system maintenance messages that are transmitted in the DS channel. These also increase overhead; however, the effect is typically negligible. MAP messages can place a burden on the Central Processing Unit (CPU), as well on DS throughput performance, because the CPU needs to keep track of all of the MAPs.

When you place any TDMA and standard code division multiple access (S-CDMA) channel on the same US, the CMTS must send double maps for each physical port. Thus, DS MAP bandwidth consumption is doubled. This is part of the DOCSIS 2.0 specification, and it is required for interoperability. Furthermore, US channel descriptors and other US control messages are also doubled.

Upstream Performance – DOCSIS Latency

In the US path, the Request-and-Grant cycle between the CMTS and the CM can only take advantage of every other MAP at most, depending upon the Round Trip Time (RTT), the length of the MAP, and the MAP advance time. This is due to the RTT that is affected by DS interleaving and the fact that DOCSIS only allows a modem to have a single Request outstanding at any given time, as well as a Request-to-Grant latency that is associated with it. This latency is attributed to the communication between the CMs and the CMTS, which is protocol-dependent. In brief, CMs must first ask permission from the CMTS to send data. The CMTS must service these Requests, check the availability of the MAP scheduler, and queue it up for the next unicast transmit opportunity. This back-and-forth communication, which is mandated by the DOCSIS protocol, produces such latency. The modem might miss every other MAP, because it is waiting for a Grant to come back in the DS from its last Request.

A MAP interval of 2 ms results in 500 MAPs per second / 2 = ~ 250 MAP opportunities per second, thus 250 PPS. The 500 MAPs is divided by 2 because, in a real plant, the RTT between the Request and the Grant will be much longer than 2 ms. It could be more than 4 ms, which will be every other MAP opportunity. If typical packets made up of 1518-byte Ethernet frames are sent at 250 PPS, that would equal about 3 Mbps because there are 8 bits in a byte. So this is a practical limit for US throughput for a single modem. If there is a limit of about 250 PPS, what if the packets are small (64 bytes)? That is only 128 kbps. This is where concatenation helps; see the Concatenation and Fragmentation Effect section of this document.

Depending on the symbol rate and modulation scheme used for the US channel, it could take over 5 ms to send a 1518-byte packet. If it takes over 5 ms to send a packet US to the CMTS, the CM just missed about three MAP opportunities on the DS. Now the PPS is only 165 or so. If you decrease the MAP time, there could be more MAP messages at the expense of more DS overhead. More MAP messages will give more opportunities for US transmission, but in a real hybrid fiber-coaxial (HFC) plant, you just miss more of those opportunities anyway.

Fortunately, DOCSIS 1.1 adds Unsolicited Grant Service (UGS), which allows voice traffic to avoid this Request-and-Grant cycle. Instead, the voice packets are scheduled every 10 or 20 ms until the call ends.

Note: When a CM is transmitting a large block of data US (for example, a 20 MB file), it will piggyback bandwidth Requests in data packets rather than use discrete Requests, but the modem still has to do the Request-and-Grant cycle. Piggybacking allows Requests to be sent with data in dedicated time slots, instead

of in contention slots, to eliminate collisions and corrupted Requests.

TCP or UDP?

A point that is often overlooked when someone tests for throughput performance is the actual protocol that is in use. Is it a connection-oriented protocol, like TCP, or connection-less, like User Datagram Protocol (UDP). UDP sends information with no regard to received quality. This is often referred to as best-effort delivery. If some bits are received in error, you make do and move on to the next bits. TFTP is another example of this best-effort protocol. This is a typical protocol for real-time audio or streaming video. TCP, on the other hand, requires an acknowledgment to prove that the sent packet was correctly received. FTP is an example of this. If the network is well maintained, the protocol might be dynamic enough to send more packets consecutively before an acknowledgment is requested. This is referred to as increasing the window size, which is a standard part of the transmission control protocol.

Note: One thing to note about TFTP is that, even though it uses less overhead because it uses UDP, it usually uses a step ack approach, which is terrible for throughput. This means that there will never be more than one outstanding data packet. Thus, it would never be a good test for true throughput.

The point here is that DS traffic will generate US traffic in the form of more acknowledgments. Also, if a brief interruption of the US results in the drop of a TCP acknowledgment, then the TCP flow will slow down. This would not happen with UDP. If the US path is severed, the CM will eventually fail the keepalive polling, after about 30 seconds, and it will start to scan DS again. Both TCP and UDP will survive brief interruptions, because TCP packets will get queued or lost and DS UDP traffic will be maintained.

The US throughput can limit the DS throughput as well. For example, if the DS traffic travels through coaxial or over satellite, and the US traffic travels through telephone lines, then the 28.8 kbps US throughput can limit the DS throughput to less than 1.5 Mbps, even though it might have been advertised as 10 Mbps maximum. This is because the low speed link adds latency to the acknowledgment US flow, which then causes TCP to slow down the DS flow. To help alleviate this bottleneck problem, Telco Return takes advantage of Point-to-Point Protocol (PPP) and makes the acknowledgments much smaller.

MAP generation on the DS affects the Request-and-Grant cycle on the US. When TCP traffic is handled, the acknowledgments must also go through the Request-and-Grant cycle. The DS can be severely hampered, if the acknowledgments are not concatenated on the US. For example, gamers might be sending traffic on the DS in 512-byte packets. If the US is limited to 234 PPS and the DS is 2 packets per acknowledgment, that would equal $512 \times 8 \times 2 \times 234 = 1.9$ Mbps.

Windows TCP/IP Stack

Typical Window rates are 2.1 to 3 Mbps download. UNIX or Linux devices often perform better, because they have an improved TCP/IP stack and do not need to send an ack for every other DS packet that is received. You can verify if the performance limitation is inside the Windows TCP/IP driver. Often this driver behaves poorly during limited ack performance. You can use a protocol analyzer from the Internet. This is a program that is designed to display your Internet connection parameters, which are extracted directly from TCP packets that you send to the server. A protocol analyzer works as a specialized web server. It does not, however, serve different web pages; rather, it responds to all requests with the same page. The values are modified based on the TCP settings of your requesting client. It then transfers control to a CGI script that does the actual analysis and displays the results. A protocol analyzer can help you to check that downloaded packets are 1518 bytes long (DOCSIS Maximum Transmission Unit [MTU]) and to check that US acknowledgments run near 160 to 175 PPS. If the packets are below these rates, update your Windows drivers and adjust your UNIX or Windows NT host.

You can change settings in the Registry, to adjust your Windows host. First, you can increase your MTU. The

packet size, referred to as MTU, is the greatest amount of data that can be transferred in one physical frame on the network. For Ethernet, the MTU is 1518 bytes; for PPPoE, it is 1492; and for dial-up connections, it is often 576. The difference comes from the fact that, when larger packets are used, then the overhead is smaller, you have less routing decisions, and clients have less protocol processing and device interrupts.

Each transmission unit consists of header and actual data. The actual data is referred to as Maximum Segment Size (MSS), which defines the largest segment of TCP data that can be transmitted. Essentially, $MTU = MSS + \text{TCP/IP headers}$. Therefore, you might want to adjust your MSS to 1380, to reflect the maximum useful data in each packet. Also, you can optimize your Default Receive Window (RWIN) after you adjust your current MTU and MSS settings: a protocol analyzer will suggest the best value. A protocol analyzer can also help you ensure these settings:

- MTU Discovery (RFC1191) = ON
- Selective Acknowledgement (RFC2018) = ON
- Timestamps (RFC1323) = OFF
- TTL (Time to Live) = OK

Different network protocols benefit from different network settings in the Windows Registry. The optimal TCP settings for cable modems seem to be different than the default settings in Windows. Therefore, each operating system has specific information on how to optimize the Registry. For example, Windows 98 and later versions have some improvements in the TCP/IP stack. These include:

- Large window support, as described in RFC1323
- Selective Acknowledgments (SACK) support
- Fast Retransmission and Fast Recovery support

The WinSock 2 update for Windows 95 supports TCP large windows and time stamps, which means you could use the Windows 98 recommendations if you update the original Windows Socket to version 2. Windows NT is slightly different from Windows 9x in how it handles TCP/IP. Remember that, if you apply the Windows NT tweaks, you will see less performance increase than in Windows 9x, simply because NT is better optimized for networking.

However, to change the Windows Registry requires some proficiency with Windows customization. If you do not feel comfortable with editing the Registry, then you will need to download a ready to use patch from the Internet, which can automatically set the optimal values in the Registry. To edit the Registry, you must use an editor, such as Regedit (choose **START > Run** and type **Regedit** in the Open field).

Performance Improvement Factors

Throughput Determination

There are many factors that can affect data throughput:

- total number of users
- bottleneck speed
- type of services accessed
- cache server usage
- MAC layer efficiency
- noise and errors on the cable plant
- many other factors, such as limitations inside the Windows TCP/IP driver

The more users that share the pipe, the more the service slows down. Further, the bottleneck might be the web site that you are accessing, not your network. When you take into consideration the service in use, regular

e-mail and web surfing is very inefficient, as far as time goes. If video streaming is used, many more time slots are needed for this type of service.

You can use a proxy server to cache some frequently downloaded sites to a computer that is in your local area network, to help alleviate traffic on the entire Internet.

While reservation and grant is the preferred scheme for DOCSIS modems, there are limitations on per-modem speeds. This scheme is much more efficient for residential usage than it is for polling or pure CSMA/CD.

Increasing Access Speed

Many systems are decreasing the homes per node ratio from 1000 to 500 to 250 to passive optical network (PON) or fiber-to-the-home (FTTH). PON, if designed correctly, could pass up to 60 people per node with no actives attached. FTTH is being tested in some regions, but it is still very cost prohibitive for most users. It could actually be worse, if you decrease the homes per node but still combine the receivers in the headend. Two fiber receivers are worse than one, but the fewer homes per fiber, the less likely you will experience laser clipping from ingress.

The most obvious segmentation technique is to add more fiber optic equipment. Some newer designs decrease the number of homes per node down to 50 to 150 HHP. It does no good to decrease the homes per node if you just combine them again in the headend (HE) anyway. If two optical links of 500 homes per node are combined in the HE and share the same CMTS US port, this could realistically be worse than if one optical link of 1000 homes per node was used.

Many times, the optical link is the limiting noise contributor, even with the multitude of actives funneling back. You must segment the service, not just the number of homes per node. It will cost more money to decrease the number of homes per CMTS port or service, but it will alleviate that bottleneck in particular. The nice thing about fewer homes per node is that there is less noise and ingress, which can cause laser clipping, and it is easier to segment to fewer US ports later.

DOCSIS has specified two modulation schemes for DS and US and five different bandwidths to use in the US path. The different symbol rates are 0.16, 0.32, 0.64, 1.28, and 2.56 Msymbols/s with different modulation schemes, such as QPSK or 16-QAM. This allows flexibility to select the throughput required versus the robustness that is needed for the return system in use. DOCSIS 2.0 has added even more flexibility, which will be expanded upon later in this document.

There is also the possibility of frequency hopping, which allows a non-communicator to switch (hop) to a different frequency. The compromise here is that more bandwidth redundancy must be assigned and, hopefully, the other frequency is clean before the hop is made. Some manufacturers set up their modems to look before you leap.

As technology becomes more advanced, ways will be found to compress more efficiently or to send information with a more advanced protocol that either is more robust or is less bandwidth intensive. This could entail the use of DOCSIS 1.1 QoS provisioning, payload header suppression (PHS), or DOCSIS 2.0 features.

There is always a give-and-take relationship between robustness and throughput. The speed that you get out of a network is usually related to the bandwidth that is used, the resources allocated, the robustness against interference, or the cost.

Channel Width and Modulation

It would appear that the US throughput is limited to around 3 Mbps, due to the previously explained DOCSIS latency. It would also appear that it does not matter if you increase the US bandwidth to 3.2 MHz or the modulation to 16-QAM, which would give a theoretical throughput of 10.24 Mbps. An increase of the channel BW and modulation does not significantly increase per-modem transfer rates, but it does allow more modems to transmit on the channel. Remember that the US is a TDMA-based, slotted contention medium where time slots are granted by the CMTS. More channel BW means more US bps, which means more modems can be supported. Therefore, it does matter if you increase the US channel bandwidth. Also, recall that a 1518-byte packet only takes up 1.2 ms of wire time on the US and helps the RTT latency.

You can also change the DS modulation to 256-QAM, which increases the total throughput on the DS by 40 percent and decreases the interleave delay for US performance. Keep in mind, however, that you will disconnect all modems on the system temporarily, when you make this change.



Caution: Extreme caution should be used before you change the DS modulation. You should make a thorough analysis of the DS spectrum, to verify whether your system can support a 256-QAM signal. Failure to do so can severely degrade your cable network performance.



Caution: Issue the **cable downstream modulation {64qam | 256qam}** command to change the DS modulation to 256-QAM:

```
VXR(config)# interface cable 3/0  
  
VXR(config-if)# cable downstream modulation 256qam
```

For more information on US modulation profiles and return path optimization, refer to How to Increase Return Path Availability and Throughput. Also refer to Configuring Cable Modulation Profiles on Cisco's CMTS. Change **uw8** to **uw16** for the Short and Long Interval Usage Codes (IUC), in the default mix profile.



Caution: Extreme caution should be used before you increase the channel width or change the US modulation. You should make a thorough analysis of the US spectrum with a spectrum analyzer, to find a wide enough band that has an adequate carrier-to-noise ratio (CNR) to support 16-QAM. Failure to do so can severely degrade your cable network performance or lead to a total US outage.



Caution: Issue the **cable upstream channel-width** command to increase the US channel width:

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

Refer to Advanced Spectrum Management.

Interleaving Effect

Electrical burst noises from amplifier power supplies and from utility powering on the DS path can cause errors in blocks. This can cause worse problems with throughput quality than errors that are spread out from thermal noises. In an attempt to minimize the affect of burst errors, a technique known as interleaving is used, which spreads data over time. Because the symbols on the transmit end are intermixed then reassembled on the receive end, the errors will appear spread apart. FEC is very effective against errors that are spread apart. The errors caused by a relatively long burst of interference can still be corrected by FEC, when you use interleaving. Because most errors occur in bursts, this is an efficient way to improve the error rate.

Note: If you increase the FEC interleave value, then you add latency to the network.

DOCSIS specifies five different levels of interleaving (EuroDOCSIS only has one). 128:1 is the highest amount of interleaving and 8:16 is the lowest. 128:1 indicates that 128 codewords made up of 128 symbols each will be intermixed on a 1 for 1 basis. 8:16 indicates that 16 symbols are kept in a row per codeword and are intermixed with 16 symbols from 7 other codewords.

The possible values for Downstream Interleaver Delay are as follows, in microseconds (μ s or usecs):

I (no. of taps)	I (increment)	64-QAM	256-QAM
8	16	220	150
16	8	480	330
32	4	980	680
64	2	2000	1400
128	1	4000	2800

Interleaving does not add overhead bits like FEC; but it does add latency, which could affect voice and real-time video. It also increases the Request-and-Grant RTT, which might cause you to go from every other MAP opportunity to every third or fourth MAP. That is a secondary effect, and it is that effect which can cause a decrease in peak US data throughput. Therefore, you can slightly increase the US throughput (in a PPS per modem way) when the value is set to a number lower than the typical default of 32.

As a workaround for the impulse noise issue, the interleaving value can be increased to 64 or 128. However, when you increase this value, performance (throughput) might degrade, but noise stability will be increased in the DS. In other words, either the plant must be maintained properly; or more uncorrectable errors (lost packets) in the DS will be seen, to a point where modems start to lose connectivity and there is more retransmission.

When you increase the interleave depth to compensate for a noisy DS path, you must factor in a decrease in peak CM US throughput. In most residential cases, that is not an issue, but it is good to understand the trade-off. If you go to the maximum interleaver depth of 128:1 at 4 ms, this will have a significant, negative impact on US throughput.

Note: The delay is different for 64-QAM versus 256-QAM.

You can issue the **cable downstream interleave-depth {8 | 16 | 32 | 64 | 128}** command. This is an example that reduces the interleave depth to 8:

```
VXR(config-if)# cable downstream interleave-depth 8
```



Caution: This command will disconnect all modems on the system, when it is implemented.

For US robustness to noise, DOCSIS modems allow variable or no FEC. When you turn off US FEC, you will get rid of some overhead and allow more packets to be passed, but at the expense of robustness to noise. It is also advantageous to have different amounts of FEC associated with the type of burst. Is the burst for actual data or for station maintenance? Is the data packet made up of 64 bytes or 1518 bytes? You might want more protection for larger packets. There is also a point of diminishing returns; for example, a change from 7 percent to 14 percent FEC might only give 0.5 dB more robustness.

There is no interleaving in the US currently, because the transmission is in bursts and there is not enough latency within a burst to support interleaving. Some chip manufacturers are adding this feature for DOCSIS

2.0 support, which could have a huge impact, if you consider all of the impulse noise from home appliances. US interleaving will allow FEC to work more effectively.

Dynamic MAP Advance

Dynamic Map Advance uses a dynamic look-ahead time, in MAPs, that can significantly improve the per-modem US throughput. Dynamic Map Advance is an algorithm that automatically tunes the look-ahead time in MAPs based on the farthest CM that is associated with a particular US port.

Refer to Cable Map Advance (Dynamic or Static?) for a detailed explanation of Map Advance.

To see if the Map Advance is Dynamic, issue the **show controllers cable slot/port upstream port** command:

```
Ninetail# show controllers cable 3/0 upstream 1

Cable3/0 Upstream 1 is up
Frequency 25.008 MHz, Channel Width 1.600 MHz, QPSK Symbol Rate 1.280 Msps
Spectrum Group is overridden
BroadCom SNR_estimate for good packets - 28.6280 dB
Nominal Input Power Level 0 dBmV, Tx Timing Offset 2809
Ranging Backoff automatic (Start 0, End 3)
Ranging Insertion Interval automatic (60 ms)
Tx Backoff Start 0, Tx Backoff End 4
Modulation Profile Group 1
Concatenation is enabled
Fragmentation 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 = 0xE224
Piggyback Requests = 0x2A65
Invalid BW Requests= 0x6D
Minislots Requested= 0x15735B
Minislots Granted = 0x15735F
Minislot Size in Bytes = 16
Map Advance (Dynamic) : 2454 usecs
UCD Count = 568189
DES Ctrl Reg#0 = C000C043, Reg#1 = 17
```

If you go to an interleave depth of 8, as mentioned earlier, you can further reduce the Map Advance because it has less DS latency.

Concatenation and Fragmentation Effect

DOCSIS 1.1 and some current 1.0 equipment supports a new feature called concatenation. Fragmentation is also supported in DOCSIS 1.1. Concatenation allows several smaller DOCSIS frames to be combined into one larger DOCSIS frame, and be sent together with one request.

Because the number of bytes being requested has a maximum of 255 minislots, and there are typically 8 or 16 bytes per minislot, the maximum number of bytes that can be transferred in one US transmission interval is about 2040 or 4080 bytes. This amount includes all FEC and physical layer overhead. Thus, the real maximum burst for Ethernet framing is closer to 90 percent of that, and it has no bearing on a fragmented grant. If you use 16-QAM at 3.2 MHz at 2-tick minislots, the minislot will be 16 bytes. This makes the limit $16 \times 255 = 4080$ bytes – 10% physical layer overhead = ~3672 bytes. To concatenate even more, you can change the minislot to 4 or 8 ticks and make the Max Concat Burst setting 8160 or 16,320.

One caveat is that the minimum burst that is ever sent will be 32 or 64 bytes, and this coarser granularity when packets are cut into minislots will have more round-off error.

Unless fragmentation is used, the maximum US burst should be set to less than 4000 bytes for the MC28C or MC16x cards in a VXR chassis. Also, set the maximum burst to less than 2000 bytes for DOCSIS 1.0 modems, if you do VoIP. This is because the 1.0 modems can not do fragmentation, and 2000 bytes is too long for a UGS flow to properly transmit around, so you could get voice jitter.

Therefore, while concatenation might not be too useful for large packets, it is an excellent tool for all those short TCP acknowledgments. If you allow multiple packets per transmission opportunity, concatenation increases the basic PPS value by that multiple.

When packets are concatenated, the serialization time of a bigger packet takes longer and affects RTT and PPS. So, if you normally get 250 PPS for 1518-byte packets, it will inevitably drop when you concatenate; but now you have more total bytes per concatenated packet. If you could concatenate four 1518-byte packets, it would take at least 3.9 ms to send with 16-QAM at 3.2 MHz. The delay from DS interleaving and processing would be added on, and the DS MAPs might only be every 8 ms or so. The PPS would drop to 114, but now you have 4 concatenated that makes the PPS appear as 456; this gives a throughput of $456 \times 8 \times 1518 = 5.5$ Mbps. Consider a gaming example where concatenation could allow many US acks to be sent with only one Request, which would make DS TCP flows faster. Assume that the DOCSIS configuration file for this CM has a Max US Burst setting of 2000 bytes, and assume that the modem supports concatenation: the CM could theoretically concatenate thirty-one 64-byte acks. Because this large total packet will take some time to transmit from the CM to the CMTS, the PPS will decrease accordingly. Instead of 234 PPS with small packets, it will be closer to 92 PPS for the larger packets. $92 \text{ PPS} \times 31 \text{ acks} = 2852 \text{ PPS}$, potentially. This equates to about 512-byte DS packets \times 8 bits per byte \times 2 packets per ack \times 2852 acks per sec = 23.3 Mbps. Most CMs, however, will be rate-limited much lower than this.

On the US, the CM would theoretically have $512 \text{ bytes} \times 8 \text{ bits per byte} \times 110 \text{ PPS} \times 3 \text{ packets concatenated} = 1.35$ Mbps. These numbers are much better than the original numbers that were obtained without concatenation. Minislot round-off is even worse when fragmenting, though, because each fragment will have round-off.

Note: There was an older Broadcom issue where it would not concatenate two packets, but it could do three.

To take the advantage of concatenation, you will need to run Cisco IOS Software Release 12.1(1)T or 12.1(1)EC or later. If possible, try to use modems with the Broadcom 3300 based design. To ensure that a CM supports concatenation, issue the **show cable modem detail**, **show cable modem mac**, or **show cable modem verbose** command on the CMTS.

```
VXR# show cable modem detail
```

Interface	SID	MAC address	Max CPE	Concatenation	Rx SNR
Cable6/1/U0	2	0002.fdfa.0a63	1	yes	33.26

To turn concatenation on or off, issue the **[no] cable upstream n concatenation** command, where *n* specifies the US port number. Valid values start with 0, for the first US port on the cable interface line card.

Note: Refer to History of the Maximum Upstream Burst Parameter for more information on DOCSIS 1.0 versus 1.1 and the concatenation issue with maximum burst size settings. Also keep in mind that modems must be rebooted, for the changes to take effect.

Single Modem Speeds

If the goal is to concatenate large frames and achieve the best possible per-modem speeds, you can change

the minislot to 32 bytes, to allow a maximum burst of 8160. The pitfall to this is that it means the smallest packet ever sent will be 32 bytes. This is not very efficient for small US packets, such as Requests, which are only 16 bytes in length. Because a Request is in the contention region, if it is made bigger, there is a higher potential for collisions. It also adds more minislot round-off error, when it is slicing the packets into minislots.

The DOCSIS configuration file for this modem will need to have a Max Traffic Burst and Max Concat Burst setting of around 6100. This would allow four 1518-byte frames to be concatenated. The modem would also need to support fragmentation, to break it apart into more manageable pieces. Because the next request is usually piggybacked and will be in the first fragment, the modem might get even better PPS rates than expected. Each fragment will take less time to serialize than if the CM tried to send one long concatenated packet.

A few settings that can affect per-modem speeds must be explained. Max Traffic Burst is used for 1.0 CMs, and it should be set for 1522. Some CMs need this to be greater than 1600, because they included other overhead that was not supposed to be included. Max Concat Burst affects 1.1 modems that can also fragment, so they can concatenate many frames with one request but still fragment into 2000-byte packets for VoIP considerations. You might need to set the Max Traffic Burst and Max Concat Burst equal to each other, because some CMs will not come online otherwise.

One command in the CMTS that could have an effect is the **cable upstream n rate-limit token-bucket shaping** command. This command helps police CMs that will not police themselves as instructed in their configuration file settings. Policing could delay packets, so turn this off if you suspect that it is throttling the throughput. This might have something to do with setting the Max Traffic Burst the same as the Max Concat Burst, so more testing might be warranted.

Toshiba did well without concatenation or fragmentation because it did not use a Broadcom chipset in the CM. It used Libit and now uses TI, in CMs higher than the PCX2200. Toshiba also sends the next Request in front of a grant, to achieve a higher PPS. This works well, except for the fact that the Request is not piggybacked and it will be in a contention slot; it could be dropped when many CMs are on the same US.

The **cable default-phy-burst** command allows a CMTS to be upgraded from DOCSIS 1.0 IOS software to 1.1 code, without CM registration failures. Typically, the DOCSIS configuration file has a default of 0 or blank for the Max Traffic Burst, which would cause modems to fail with reject(c) when they register. This is a reject CoS because 0 means unlimited maximum burst, which is not allowed with 1.1 code (because of VoIP services and maximum delay, latency, and jitter). The **cable default-phy-burst** command overrides the DOCSIS configuration file setting of 0, and the lower of the two numbers takes precedence. The default setting is 2000 and the maximum is now 8000, which will allow five 1518-byte frames to be concatenated. It can be set to 0 to turn it off:

```
cable default-phy-burst 0
```

Some Recommendations for Per-Modem Speed Testing

1. Use Advanced Time-Division Multiple Access (A-TDMA) on the US for 64-QAM at 6.4 MHz channel.
2. Use a minislot size of 2. The DOCSIS limit is 255 minislots per burst, so 255×48 bytes per minislot = 12240 maximum burst \times 90 percent = ~11,000 bytes.
3. Use a CM that can fragment and concatenate and that has a full-duplex, FastEthernet connection.
4. Set the DOCSIS configuration file for no minimum, but with a maximum of 20 MB up and down.
5. Turn off US rate-limit token-bucket shaping.
6. Issue the **cable upstream n data-backoff 3 5** command.
7. Set the Max Traffic Burst and Max Concat Burst to 11000 bytes.
8. Use 256-QAM and 16 interleave on the DS (try 8 also). This gives less delay for MAPs.

9. Issue the **cable map–advance dynamic 300 1000** command.
10. Use a IOS Software Release 15(BC2) image that fragments correctly, and issue the **cable upstream n fragment–force 2000 5** command.
11. Push UDP traffic into the CM and increment it until you find a maximum.
12. If you are pushing TCP traffic, use multiple PCs through one CM.

Results

- Terayon TJ735 gave 15.7 Mbps. This is possibly a good speed because of less bytes per concatenated frame and a better CPU. It seems to have a 13–byte concatenation header for the first frame and 6–byte headers after, with 16–byte fragment headers and an internal 8200–byte maximum burst.
 - Motorola SB5100 gave 18 Mbps. It also gave 19.7 Mbps with 1418–byte packets and 8 interleave on the DS.
 - Toshiba PCX2500 gave 8 Mbps, because it seems to have a 4000–byte internal maximum burst limit.
 - Ambit gave the same results as Motorola: 18 Mbps.
-
- Some of these rates can drop when in contention with with other CM traffic.
 - Make sure 1.0 CMs (which can not fragment) have a maximum burst less than 2000.
 - 27.2 Mbps at 98 percent US utilization was achieved with the Motorola and Ambit CMs.

New Fragment Command

cable upstream n fragment–force *fragment–threshold number–of–fragments*

Parameter	Description
<i>n</i>	Specifies the upstream port number. Valid values start with 0, for the first upstream port on the cable interface line card.
<i>fragment–threshold</i>	The number of bytes that will trigger fragmentation. The valid range is 0 to 4000, with a default of 2000 bytes.
<i>number–of–fragments</i>	The number of equal size fragments into which each fragmented frame is split. The valid range is 1 to 10, with a default of 3 fragments.

DOCSIS 2.0 Benefits

DOCSIS 2.0 has not added any changes to the DS, but it has added many to the US. The advanced physical layer specification in DOCSIS 2.0 has these additions:

- 8–QAM, 32–QAM, and 64–QAM modulation schemes
- 6.4 MHz channel width
- Up to 16 T bytes of FEC

It also allows 24 taps of pre–equalization in the modems and US interleaving. This adds robustness to reflections, in–channel tilt, group delay, and US burst noise. Also, 24–tap equalization in the CMTS will help older, DOCSIS 1.0 modems. DOCSIS 2.0 also adds the use of S–CDMA in addition to A–TDMA.

Greater spectral efficiency with 64–QAM creates better use of existing channels and more capacity. This provides higher throughput in the US direction and slightly better per–modem speeds with better PPS. The

Issue the **show controllers cable slot/port upstream port** command to see the current US settings of the modem.

```
ubr7246-2# show controllers cable 6/0 upstream 0

Cable6/0 Upstream 0 is up
Frequency 33.000 MHz, Channel Width 6.400 MHz, 64-QAM Sym Rate 5.120 Msps
This upstream is mapped to physical port 0
Spectrum Group is overridden
US phy SNR_estimate for good packets - 36.1280 dB
Nominal Input Power Level 0 dBmV, Tx Timing Offset 2066
Ranging Backoff Start 2, Ranging Backoff End 6
Ranging Insertion Interval automatic (312 ms)
Tx Backoff Start 3, Tx Backoff End 5
Modulation Profile Group 243
Concatenation is enabled
Fragmentation is enabled
part_id=0x3138, rev_id=0x02, rev2_id=0x00
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 = 2
Minislot Size in Symbols = 64
Bandwidth Requests = 0x7D52A
Piggyback Requests = 0x11B568AF
Invalid BW Requests= 0xB5D
Minislots Requested= 0xAD46CE03
Minislots Granted = 0x30DE2BAA
Minislot Size in Bytes = 48
Map Advance (Dynamic) : 1031 usecs
UCD Count = 729621
ATDMA mode enabled
```

Issue the **show interface cable slot/port service-flow** command to see the service flows for the modem.

```
ubr7246-2# show interface cable 6/0 service-flow
```

Sfid	Sid	Mac Address	QoS	Param	Index	Type	Dir	Curr	Active
			Prov	Adm	Act			State	Time
18	N/A	00e0.6fle.3246	4	4	4	prim	DS	act	12d20h
17	8	00e0.6fle.3246	3	3	3	prim	US	act	12d20h
20	N/A	0002.8a8c.6462	4	4	4	prim	DS	act	12d20h
19	9	0002.8a8c.6462	3	3	3	prim	US	act	12d20h
22	N/A	000b.06a0.7116	4	4	4	prim	DS	act	12d20h
21	10	000b.06a0.7116	3	3	3	prim	US	act	12d20h

Issue the **show interface cable slot/port service-flow sfid verbose** command to see the specific service flow for that particular modem. This will display the current throughput for the US or DS flow and the configuration file settings of the modem.

```
ubr7246-2# show interface cable 6/0 service-flow 21 verbose
```

```
sfid : 21
Mac Address : 000b.06a0.7116
Type : Primary
Direction : Upstream
Current State : Active
Current QoS Indexes [Prov, Adm, Act] : [3, 3, 3]
Active Time : 12d20h
Sid : 10
Traffic Priority : 0
Maximum Sustained rate : 21000000 bits/sec
```

```

Maximum Burst                : 11000 bytes
Minimum Reserved Rate        : 0 bits/sec
Admitted QoS Timeout         : 200 seconds
Active QoS Timeout           : 0 seconds
Packets                      : 1212466072
Bytes                       : 1262539004
Rate Limit Delayed Grants    : 0
Rate Limit Dropped Grants    : 0
Current Throughput           : 12296000 bits/sec, 1084 packets/sec
Classifiers                   : NONE

```

Ensure that no delayed or dropped packets are present.

Issue the **show cable hop** command to verify that there are no uncorrectable FEC errors.

```

ubr7246-2# show cable hop cable 6/0

```

Upstream Port	Port Status	Poll Rate (ms)	Missed Poll Count	Min Poll Sample	Missed Poll Pcnt	Hop Thres	Hop Period (sec)	Corr FEC Errors	Uncorr FEC Errors
Cable6/0/U0	33.000 Mhz	1000	* * *	*set to fixed frequency	* * *	* * *	0	0	
Cable6/0/U1	admindown	1000	* * *	frequency not set	* * *	* * *	0	0	
Cable6/0/U2	10.000 Mhz	1000	* * *	*set to fixed frequency	* * *	* * *	0	0	
Cable6/0/U3	admindown	1000	* * *	frequency not set	* * *	* * *	0	0	

If the modem is dropping packets, then the physical plant is affecting the throughput and must be fixed.

Summary

The previous sections of this document highlight the shortcomings when you take performance numbers out of context with no understanding of the impact on other functions. While you can fine tune a system to achieve a specific performance metric or overcome a network problem, it will be at the expense of another variable. To change the MAPs/s and interleave values might get better US rates, but at the expense of DS rate or robustness. To decrease the MAP interval does not make much difference in a real network and just increases CPU and bandwidth overhead on both the CMTS and the CM. To incorporate more US FEC increases US overhead. There is always a trade-off and compromise relationship between throughput, complexity, robustness, and cost.

If admission control is used on the US, it will make some modems not register when the total allocation is used up. For example, if the US total is 2.56 Mbps to use and the minimum guarantee is set to 128k, only 20 modems would be allowed to register on that US if admission control is set to 100 percent.

Conclusion

You must know what throughput to expect, to determine what subscribers data speed and performance will be. Once you determine what is theoretically possible, a network can then be designed and managed to meet the dynamically changing requirements of a cable system. Then you must monitor the actual traffic load, to determine what is being transported and when additional capacity is necessary to alleviate bottlenecks.

Service and the perception of availability can be key differentiating opportunities for the cable industry, if networks are deployed and managed properly. As cable companies make the transition to multiple services, subscriber expectations for service integrity move closer to the model that has been established by legacy voice services. With this change, cable companies need to adopt new approaches and strategies which ensure that networks align with this new paradigm. There are higher expectations and requirements now that we are a telecommunications industry and not just entertainment providers.

While DOCSIS 1.1 contains the specifications that assure levels of quality for advanced services such as VoIP, the ability to deploy services compliant with this specification will be challenging. Because of this, it is imperative that cable operators have a thorough understanding of the issues. A comprehensive approach to choose system components and network strategies must be devised, to ensure successful deployment of true service integrity.

The goal is to get more subscribers signed up but not jeopardize service to current subscribers. If Service Level Agreements (SLAs) to guarantee a minimum amount of throughput per subscriber are offered, the infrastructure to support this guarantee must be in place. The industry is also looking to serve commercial customers and add voice services. As these new markets are addressed and networks are built, it will require new approaches: denser CMTSs with more ports, a distributed CMTS farther out in the field, or something in between (like adding a 10baseF to your house).

Whatever the future has in store, it is assured that networks will get more complex and the technical challenges will increase. The cable industry will only be able to meet these challenges if it adopts architectures and support programs that can deliver the highest level of service integrity in a timely manner.

Related Information

- [Troubleshooting Slow Performance in Cable Modem Networks](#)
- [Troubleshooting uBR Cable Modems Not Coming Online](#)
- [Configuring Cable Modulation Profiles on Cisco's CMTS](#)
- [Troubleshooting Slow Performance in Cable Modem Networks](#)
- [What is the Maximum Number of Users per CMTS?](#)
- [Cisco Cable/Broadband Software Center \(registered customers only\)](#)
- [Broadband Cable](#)
- [Technical Support – Cisco Systems](#)

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

© 2007 – 2008 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Aug 25, 2008

Document ID: 19220
