

How to Increase Return Path Availability and Throughput

Introduction

Networks currently being designed and implemented are broadband in nature and are shared among many people for economic/cost reasons. It is important to understand the factors that affect the throughput per user in broadband, shared networks and how to increase this throughput.

The availability of the service offered to customers can be increased in a number of ways, including developing techniques to narrow the modem transmit window, designing return path segmentation, increasing the modem bandwidth or modulation complexity, decreasing ingress, and monitoring the plant for quality assurance (QA). Some recommendations regarding the cable modem termination system (CMTS) configuration will also be explained in this paper.

Eighty percent of shared network problems are usually associated with the physical plant and how it is designed, tested, and maintained. This paper will begin by addressing Layer 1 of the Open System Interconnection (OSI) model—the physical layer.

Balancing the Reverse Path for Increased Availability

Modems that transmit at reasonably high levels have a better carrier-to-ingress ratio (CIR) and better availability than modems transmitting at lower levels. However, all modems funnel back or combine at one point, the CMTS upstream port, so each modem will have the same carrier-to-noise ratio (CNR). To increase the overall CNR, the noise must be padded down. One way to achieve a high modem output and at the same time pad down the noise is to have consistent balancing techniques and the same losses from the house to the next reverse active device.

There are two balancing techniques that can be used—balancing constant inputs to the port or balancing constant inputs to the hybrid. Because different amplifiers have different internal losses, the same loss from the modem to the reverse hybrid is not achievable unless reverse input pads are installed in the line extenders, or special taps or inline pads are used. Another way to keep the modems that are located off the line extenders running high is to balance for constant inputs to the port and not the hybrid.

SA, C-COR, GI (Motorola), and other amplifier manufacturers all support constant inputs to the port. By reverse-balancing constant inputs to the port, higher RF levels hit the hybrids in the line extenders and force a higher modem level out of the houses that are attached to taps off of line extenders.

45-52 dBmV is an optimal modem output, but is often difficult to achieve because of logistics and economics/cost. A tight operating range would provide a good CIR and still provide some headroom (Data over Cable Service Interface Specifications [DOCSIS] = 8-58 dBmV) for temperature effects, differing bandwidths, customer-induced losses, age effects, and modulation changes. With a network interface device (NID), installed a year later for telephony service, an insertion loss of almost 3 dB can be expected.

When balancing constant inputs of roughly 17 dBmV to the amplifier ports, the line extender chip now yields a better CNR. There are more line extenders in today's architectures than any other style of amplifier. Distortions and compression aren't an issue in RF amplifier upstream hybrids. Hybrids have a relatively high compression point and a low noise figure of about 5 dB. When considering thermal noise from the active devices in the system, the overall CNR doesn't change much because the laser is usually the limiting factor. The CIR is better, because the ingress from the lower-value taps off of the line extenders will be padded more at the line extender.

Any line extender that feeds a node or bridger amp will be padded more when doing constant inputs to the port instead of the hybrid. This helps to force a higher level from the modems and a better CIR. It also drops the noise and ingress when exiting the higher output pad. Some systems may opt to use the amplifier input pad to achieve this scenario as well. If amplifiers were designed with the same internal losses, balancing levels would be the same for all the different devices whether balanced to the port or hybrid.

Increasing Access Speed

There are a few ways to increase access speed. The most obvious is to decrease the number of homes per CMTS port. One could also increase the initial speed by using different symbol rates (different bandwidth), different modulation techniques, or more channels. Quality of service (QoS) provisioning is another option. This could entail more time slices per user, code-division multiple access (CDMA) with dedicated codes, or DOCSIS 1.1 QoS service flows.

DOCSIS 1.0 has specified two modulation schemes-quadrature phase shift keying (QPSK) and 16-quadrature amplitude modulation (QAM)-and five different bandwidths-.2, .4, .8, 1.6, and 3.2 MHz-to be used in the reverse path. This allows agility in selecting the throughput required versus the robustness needed for the return system being used. The only physical layer difference between DOCSIS 1.0 and 1.1 is that DOCSIS 1.1 has added eight taps of pre-equalization to the cable modems for added robustness to in-channel tilt and group delay issues on the upstream.

There is a give-and-take relationship between robustness and throughput. Higher network speed is usually related to the bandwidth used, resources allocated, or the robustness to interference-not to mention cost.

DOCSIS 2.0 Benefits

The advanced physical layer spec. in DOCSIS 2.0 has added 8-, 32-, and 64-QAM modulation schemes; 6.4 MHz channel width; and up to 16 Tbytes of Forward Error Correction (FEC). It allows 24 taps of pre-equalization in the modems and upstream interleaving. This adds robustness to reflections, in-channel tilt, group delay, and upstream burst noise. DOCSIS 2.0 also adds the use of CDMA in addition to time-division multiple Access (TDMA).

Some physical layer problems are listed in Table 1, along with the DOCSIS 2.0 solution.

Table 1

RF Plant Problems and DOCSIS 2.0 Solutions

Problem	Solution
Additive white gaussian noise (AWGN)	Higher FEC (T=16), lower implementation loss, Trellis Coded Modulation.
Impulse/burst noise	Higher FEC, byte interleaving (A-TDM mode), time spreading (lower symbol rate, S-CDMA mode). Mitigation of 100-micro seconds (usec) burst noise demonstrated in both modes at CableLabs.
Ingress, Common Path Distortion (CPD)	All silicon implementations employ ingress cancellation, although not in the DOCSIS 2.0 spec. Cancellation of four modulated narrowband carriers demonstrated at CableLabs. Cancellation of wideband ingress also demonstrated.
Micro reflections	Transmit equalizer increased to 24 T-spaced taps.
Hum modulation	Receiver tracking loops.

Some features of DOCSIS 2.0 are listed in Table 2, along with their associated benefits.

Table 2

DOCSIS 2.0 Features and Advantages

Feature	Advantage
Higher-order modulation schemes	More bits transmitted per RF bandwidth. Better coding and implementation loss to permit higher-order modulation transport. Best tactic for non-compressible large packets.
Higher symbol rate	Larger RF channel size means better statistical multiplexing and more efficient use of RF bandwidth.
Synchronous operation	Reduced preamble length for more efficient small packets.

Tables 3 and 4 report some of the recent testing results from CableLabs.

Table 3

Continuous Wave (CW) Ingress Performance

Modulation	FEC	No INCA™ Technology Minimum C/I dB	INCA Minimum C/I dB	INCA Advantage	INCA C/I Benefits
QPSK	K=16, T=10	5.65	-1.65	7.3 dB	5.3X
QPSK	K=253, T=1	5.65	-0.35	6.0 dB	4X
16-QAM	K=128, T=10	12.15	-0.5	12.7 dB	18X
16-QAM	K=253, T=1	12.15	0.75	11.4 dB	13X

I believe these measurements were made with a spectrum analyzer with an RBW setting of 300 kHz, which would require a 9 dB correction factor for Table 3 if looking at a 3.2 MHz wide channel.

Table 4

INCA QAM Ingress Performance

Modulation	FEC	No INCA Minimum C/I dB	INCA Minimum C/I dB	INCA Advantage	INCA C/I Benefits
QPSK	K=16, T=10	10.5	1.7	8.5 dB	5.3X
QPSK	K=253, T=1	12.2	3.2	9.0 dB	7X
16-QAM	K=128, T=10	15.1	7.7	7.3 dB	5.5X
16-QAM	K=253, T=1	16.1	8.8	7.3 dB	5.3X

INCA is a trademark of Texas Instruments Incorporated

K= FEC codeword (CW) size and T = Tbytes of FEC, which is (2*T) bytes of FEC coverage per codeword. INCA = Ingress Noise Cancellation Algorithm.

Mitigating the Variables

One method to increase availability includes functionality built into the customer premises equipment (CPE). A good place to start is using set-top boxes and cable modems with variable or high outputs, frequency agility, and different modulation schemes.

Variable outputs will allow tighter compliance with the specifications. Higher outputs will allow a greater CIR and frequency agility will allow the modems to be used on an unoccupied frequency. This will allow the modems to possibly scan for a reliable frequency, which may allow communication free of bit errors. Flexibility to grow and change with the demand is always a must.

Different modulation schemes allow flexibility in robustness and throughput. QPSK is approximately 7 dB more robust than 16-QAM, but requires twice as much bandwidth for equal throughput. High QAM uses less bandwidth, but is more susceptible to bit errors from noise and interference, as well as group delay and reflection susceptibility. CDMA and high QAM have a very high peak-to-average ratio and may cause laser clipping if not taken into consideration.

Another method involves using return-path automatic gain control (AGC) and interdiction units. Return-path AGC is ideal for compensating for temperature-induced variances, but the cost may not be justified. Return-path, thermal equalizers can compensate for some temperature fluctuations that occur in cable, hybrids, and passive amplifiers, which may have an adverse effect on levels-as much as +/- 3 dB. If return-path AGC were used, the pilot can be placed below 15 MHz in the noisy band, not wasting any prime bandwidth. This pilot could be used for referencing, maintenance, monitoring, troubleshooting, and constant laser modulation.

Fabry-Perot (FP) lasers can have a high noise floor when there is nothing to transmit, so node manufacturers will insert a continuous wave carrier and call it a noise suppressor or quieting oscillator. Harmonic uses one at 40 MHz; SA uses one at 500 kHz or 6 MHz. This continuous wave carrier can sometimes cause issues, especially if the node and fiber receivers in the headend are from different manufacturers. A small jumper may need to be removed in the node to eliminate the carrier.

An interdiction unit is a programmable tap controlled from the headend. It is like a set-top box on the strand-it can be controlled to shut down the return or allow certain pay-per-view movies to be seen. It eliminates truck rolls (and may be the only way to really offer plug-and-play cable modems through the retail market without a truck roll). Most multiple system operators (MSOs) don't like the idea of interdiction units because of higher cost, more power consumption, and lost revenues from converter-box and remote-control rentals. Some vendors are providing return-path switching gear for noise mitigation and troubleshooting (such as Electroline's ClearPath system).

Noise Mitigation

For a better CIR, the CPE should transmit as high as possible. Because most of the ingress is from the house and drop, padding down both the ingress and the signal at the tap input will keep a good CIR. Another way to get a better CIR is to get rid of the ingress, but customers are hard to control.

The house must be grounded and bonded correctly for safety and elimination of ground loops, which can cause hum-related problems. Coils, beads, and filters may be used to mitigate noise and ingress. Coiling the drop cable prior to insertion into the house appliance may have a debilitating affect on reverse ingress, which is caused by common mode currents entering through a breach in the cable sheath. SignalVision makes a clamp-on ferrite bead that acts as a choke just like the cable coil. More testing needs to be done to find the best location for this device-the tap, inside or outside the ground block, or in the house. The tap location may not be the best; it could be too late to choke off ingress if a breach is found in the sheath and the ingress is already on the center conductor.

If a house amplifier is used for forward-path-level requirements, it can also act as either a highpass filter, a passive reverse with reverse padding, or an active reverse gain. An amplifier is one-way and offers great isolation. If reverse gain is required, go back to the design stage. The problem with reverse gain at the house is it also amplifies the ingress, which is mostly from the house in the first place.

Other possibilities include using highpass filters on houses without modem service, using only high-value taps in the design, or standardizing in-house losses and components.

Highpass filters attenuate the ingress from all the non-modem houses by at least 40 dB. Cost and logistics are the biggest problem with this. A highpass, bandpass, or dual-bandpass filter is used to allow only the return frequencies of the modem to pass. This prevents extra ingress from affecting the other services. Some systems will protect their investment and buy windowed filters. This way it can be moved from the tap to the splitter feeding the TV set, which may need a reverse window for video on demand (VOD). These systems may also use a dual-windowed filter to the modem to eliminate theft of cable service. Only the modem's forward and reverse frequencies are distributed.

High-value taps will create a more equal loss, but will require more actives and associated cost. Standardized in-house losses and components will limit some of the variances encountered as well as mitigate some potential problems. There are a few manufacturers now offering better ferrites and voltage blocking capacitors to eliminate potential problems of isolation, transient hum, harmonic generation, and magnetism.

Tap Flat Loss and Cable Non-linear Loss

If coaxial cable had flat loss, the forward and reverse paths could be balanced for flat inputs and outputs, with no need for equalizers or cable simulators, and modems would all transmit the same regardless of which tap they were located off of. Another way to look at it would be to make the tap loss non-linear opposite that of cable. This is the advantage of fiber optics. Temperature and electro-magnetic interference have no affect and there is no tilt problems with fiber as compared to coaxial cable. It also has much less attenuation per unit length.

This equalized tap will assure equal forward outputs and reverse inputs at all frequencies by attenuating variably at different frequencies instead of flat loss. This will compensate for cable loss when comparing forward and reverse frequencies, making the modem outputs equal no matter where the location-and ingress will be relatively the same no matter which tap it originates from. This is already done for mid-split and high-split LAN systems. This tap could have a plug-in accessory to change the value of the equalizer, or cable simulator or a diplex filter after the directional coupler. This would allow separation of the forward and return frequencies and only affect the tap spigots, not the through leg of the tap. This will allow us to insert one pad for all the spigots and affect only the forward or return paths, which ever is desirable.

Since feeder equalizers (FEQs) are used for long cable spans and this is what causes the biggest variation in levels from the modems, the FEQ could possibly have variable attenuation. For example; instead of setting the attenuation from 54-750, the largest attenuation could start at 5 MHz and decrease in attenuation until the high frequency (750 MHz) is reached.

Eagle Comtronics produces an equalizer and cable simulator in a small filter design that installs directly on the F ports of the tap. A cable simulator can be used at the first tap out of an amplifier with the equalizers at the end of lines. The cable simulator is ideal for new passive optical network (PON) designs. With only one active node, you can run higher levels for better reach. Distortions are limited because there are no amplifier cascades. The taps used can be higher-value taps-but then the reverse from the modems may not work because of the excessive loss. The special cable-simulated tap looks like a 32 dB tap at the high frequency and a 23 dB tap on the low frequency. The special equalizer taps start at 5 MHz and go to 860 MHz, eliminating the need for a feeder equalizer and forcing the modems off of low-value taps to transmit more. This also drops the noise and ingress from within the house to achieve a better CIR for everyone. Price is the only disadvantage to this method; it is less expensive to buy a tap that does this instead of using four or eight filters per tap. Several vendors (SA, Motorola, Maspro, Lindsay, and Antronix) are making split band taps with plug-in reverse pads. Some of those companies are producing equalizer and cable simulator taps as well. With an equalizer, there's no worry about the frequency split if you ever increase upstream from 40 to 65 MHz.

If a "hot tap" is required, it is advantageous to install a pad that only affects the reverse band. Arcom offers a high-pass filter with 3 dB increments of attenuation on the return passband. These are known as step attenuators. Most of these filters will pass at least .5 amperes of current, which should be more than enough for a NID, if one is being used. This allows more attenuation on the reverse band without affecting the tilt or levels on the forward band. Now the amplifiers and the taps need to be balanced!

How to determine what filter, pad, equalizer, or cable simulator to install follows. Some users rely entirely on the design for placement of these devices, or actually test it. Attach a sweep system to the tap and observe the forward sweep response with no reference. This is known as a raw sweep. For example, depending on the system requirements, if you want 4 dB of positive tilt from every tap, look at a scan or sweep and balance the forward output for that tilt by changing the equalizer or cable simulator. The cable simulators in the first tap off an active equalizer will probably be design-oriented, but the use of equalizers may be subject to the conditions.

For hot taps, a step attenuator or plug-in reverse pad may be warranted. A reverse sweep can indicate the value needed. Document the telemetry level when sweeping from the first tap off of an active equalizer-this is the "magic number" to attempt to hit when balancing the other taps. Example: Suppose the first tap off the active is a 23-dB tap and the reverse sweep indicated 1 dBmV of telemetry. If an end-of-line, eight-dB tap is reverse-swept and the telemetry says 10 dBmV, install a 9-dB pad. Make sure the telemetry is at an upstream frequency not near the roll-off of diplex filters. There are test units on the market that "talk" to the CMTS (Sunrise CM1000 and Acterna DSAM 2500). These can be used to see what level they produce to hit the CMTS at the required input. If the level isn't within your design constraints, insert the appropriate accessory.

Eagle Comtronics in conjunction with Com21 was making a highpass filter that automatically disabled when a certain dual tone, multiple frequency (DTMF) signal was detected from the modem. The level was high enough to power this switch, but it would have to be extremely fast. I'm not sure if they're still pursuing this. Another company, Spacenet out of Sweden, has developed the Dynamic Ingress Blocking (DIB) device that only opens the return channel during the short period a connected house actually transfers data. More info is available at <http://www.spacenet.se>.

Upstream Frequency Hopping

Frequency hopping can increase your upstream availability. Modems dynamically hop to a clear area of the spectrum. Modems can run 3 dB hotter if the levels are based off of power/Hz-some of the spectrum is allocated just for redundancy. Frequencies previously avoided like 27 MHz can now be used. In addition, the signal can hop based on time and/or day, and can assign a different power based on the hop frequency or range of frequencies. Availability is increased even more by changing the modulation and/or bandwidth in addition to frequency hopping. Dynamic modulation selection with user-configurable criteria allows you to maintain the highest channel width and modulation rate, which equates to the highest bit rate. Dynamic modulation also will revert to the original channel when available.

There are some disadvantages to frequency hopping. Frequency hopping was never intended to replace good quality craftsmanship and monitoring equipment-it is intended to help with the intermittent problems that arise after you have qualified your plant. Remember that ingress at 15 MHz will still affect the modem at 30 MHz because of laser clipping or harmonics. Another pitfall to frequency hopping is that more bandwidth must be allocated for redundancy and hopping. If it's a blind hop, more time could be spent hopping than actually transmitting useful data.

The benefits to frequency hopping are great. It improves response time to ingress and noise impairments. If the CMTS incorporates a spectrum analyzer on-board, or cross-references with a spectrum analyzer, it eliminates "blind" and unnecessary frequency hops. This saves time and effort when troubleshooting cable plant impairments. It also increases cable plant availability, improves the percentage of modems online, and maximizes spectrum use.

Some emerging enhancements allow the hopping criteria to be based on the CNR instead of the signal-to-noise ratio (SNR) from the demodulator chip. You may also be able to select a few different hop criteria, such as modem performance of Correctable/Uncorrectable FEC and user-configurable CNR threshold criteria. Some devices allow an interface for real-time spectrum display, so you can monitor the return-path frequency from the comforts of your home.

Additional Plant and Public Relations Suggestions

Some additional suggestions include self-terminating taps and outlets, placement of different carriers at different frequencies, tight cumulative leakage index compliance, better manufacturer testing, specifications and equipment, different frequency splits, monitoring with performance history, and consumer awareness and accountability. By making consumers aware of the problems they can cause, they may feel accountable and help to limit the in-house ingress. Some training and public relations work can go a long way toward educating subscribers and changing their perceptions of the cable provider.

As far as carrier placement goes, wider-frequency carriers should be located near the middle of the range and narrow carriers near the upper end, because of group delay problems. More robust carriers should be located between 5-15 MHz; this is usually the problem area. Stay away from 27 MHz because of CB and radio-controlled cars.

CMTS Configuration and Optimization

The following CMTS settings should be optimized for better upstream throughput: Upstream Power Level, Upstream Power-Adjust Threshold, Upstream Power-Adjust Continue, Upstream Power-Adjust Noise, Modulation Profiles, and Upstream Minislot Size. Some other suggestions for optimized upstream throughput include using Upstream Rate Limit Token Bucket Shaping, proper settings for the DOCSIS configuration files, the use of upstream concatenation, proper use of downstream Dynamic Map Advance, and the use of variable downstream interleaving.

Downstream Dynamic Map Advance, delay from downstream interleaving, and the use of upstream concatenation affect the upstream packet-per-second (PPS) rate for each individual modem.

Upstream Rate Limit Token Bucket Shaping helps control modems that have exceeded the “cap” built into their DOCSIS configuration file. Instead of dropping packets and forcing resends and more requests that could collide with other modem requests, the modem is throttled by the CMTS.

The Minimum Guaranteed and Max Upstream rates can also be set in the configuration file to affect individual modem speeds. The Max Rate can be thought of as a “throttle,” while the Min Rate is a guaranteed minimum.

For upstream concatenation to operate with larger packets, the upstream Max Burst field in the DOCSIS configuration file should be of adequate size to be efficient. To concatenate three 1024-byte packets, the Max Burst should be at least 3072 bytes.

Note: The upstream Max Burst field should be set for at least 1522 bytes and no more than 2000 for DOCSIS 1.0 modems-these modems cannot fragment, and this could affect voice-over-IP (VoIP) traffic from DOCSIS 1.1 modems.

Cable Modem Upstream Ranging, Levels, and SNR

Ranging

A modem is polled at least every 25 seconds during station maintenance. This is when the CMTS will command the modem to adjust its level if it is outside the “power-adjust threshold” (default +/-1 dB) from the nominal receive level (default of 0 dBmV). The threshold range can be changed from 0 to +/- 10 dB. A big threshold eliminates many adjustments, but will produce modems with different CNRs. The flap list also has a range of +/-1 to +/- 10 dB for reporting purposes.

Note: Never set the power-adjust threshold range to zero-modems will not set up unless they hit the CMTS exactly at 0 dBmV, and ranging opportunities will be taken up by modems continually changing levels. The default range of +/- 1 dB may be sufficient, but a range of +/- 2 may be justified for temperature swings you do not want to track.

The modem will be commanded to change its power in units of .25 dB by the CMTS, but the cable modems only have one dB of granularity. The CMTS may instruct the modem to change by 14 (14*.25 = 3.5 dB), but the modem will only change by three or four dB.

The nominal level can be set between -10 and +25 dBmV, but is easier to manage if left at 0 dBmV, and there is less chance of total power overload from other ingress. The DOCSIS RFI Specification 4.2.12 states that the total input power to the upstream demodulator must not exceed 35 dBmV for the entire 5-42 MHz band.

Note: Station maintenance on Cisco universal broadband routers is one second per modem up to 25 modems. For example, if there are only four modems on a particular Media Access Control (MAC) domain (one downstream and all its associated upstreams), each modem will be polled every four seconds. Station maintenance is a maximum of every 15 seconds when Hot-standby Communication-to-Communication Protocol (HCCP) is configured for N+1 availability.

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

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

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

If a modem does not receive a station maintenance message within its T4 timer (30 seconds minimum to 35 seconds maximum), it will go offline and restart scanning downstream.

Tip: A novel method of forcing the cable modem and the CMTS to exchange station maintenance “keepalives” is to use the Cisco CMTS command “ping docsis {modem-mac-address | modem-ip-address} [count] [verbose]”. This command will order the cable modem to send station-maintenance ranging requests to the CMTS in quick succession in order to verify DOCSIS connectivity to the cable modem. By using the “verbose” option, the power, timing, and frequency adjustments sent by the CMTS can be viewed for each ranging response message. Time adjustments between +/- 3 ticks are not tracked.

```
uBR7246-VXR# ping docsis 10.2.2.36 verbose
```

```
Queuing 5 MAC-layer station maintenance intervals, timeout is 25 msec:
```

```
Reply from 0001.9659.43fd: 4 ms, tadj=0, padj=0.25, fadj=0
```

```
Reply from 0001.9659.43fd: 4 ms, tadj=-1, padj=0.25, fadj=0
```

```
Reply from 0001.9659.43fd: 4 ms, tadj=-1, padj=0, fadj=0
```

```
Reply from 0001.9659.43fd: 4 ms, tadj=-1, padj=0, fadj=0
```

```
Reply from 0001.9659.43fd: 4 ms, tadj=-1, padj=0, fadj=0
```

```
Success rate is 100 percent (5/5)
```

If regular “ping” packets generate errors, but “ping docsis” packets do not, it could be because of the modulation profile. If using a Mix profile, the Short and Long data packets for ping will use 16-QAM, while the station maintenance packets for ping docsis will be QPSK.

A cable modem will adjust properly when a CMTS commands it to, but if the CMTS receives a level outside the threshold again, it will poll the cable modem-within one second. If the level of noise on the upstream is enough to distort the upstream level being received by the CMTS, then the modem and CMTS will go into “power-adjust noise” averaging mode. An asterisk (*) will be displayed next to the receive level in the “Show Cable Modem” command. When this occurs, modems are polled using a one-second interval. By default, the percentage of “noisy” ranging responses that cause a modem to enter “power-adjust noise” mode is 30 percent. This percentage may be increased to alleviate excessive power level adjustments in the presence of noise.

If the modem has reached its maximum transmit power (DOCSIS spec of 58 dBmV for QPSK and 55 dBmV for 16-QAM; could be more), an explanation point (!) will be displayed next to the receive level in the “Show Cable Modem” command. The modem will be permitted to stay online if it is within the “power-adjust continue” range (the default is two dB from nominal). This can be increased up to 25 dB, but six dB works well. The modem will be commanded to change level every 25 seconds during its station maintenance, but only once, not continuously.

There is no special message sent by the modem when it reaches max transmit power. If the CMTS keeps telling the modem to increase its transmit power several times in a row, the CMTS concludes that the cable modem cannot increase its transmit power any higher and flags it with “!”.

There are some corner conditions where the flag may be erroneously set. For example, if the CMTS commands the modem to adjust its power by one dB, but the modem only adjusts its transmit power if the adjustment request is above one dB, the CMTS may erroneously conclude that the cable modem is out of transmit power. Thus, besides looking at the “!”, the customer should also look at the actual reported power level and see if it is outside the target power level. The real purpose of the “!” is troubleshooting and identification.

By increasing the “Continue” command to six dB, the modem will be permitted to stay online if the CMTS receive level is between -6 dBmV and 0 dBmV. If the level is above -1, you will not see a “!” because the level is within the “threshold” window. If the level is below -6 dBmV, the modem will go offline. For systems that still have high-value taps (29 and 26 dB), this helps keep the modems online, but will produce modems with different CNRs.

The HFC plant should be designed so that in normal DOCSIS 1.x operation, modems transmit between 42 and 52 dBmV. This gives enough headroom for customer-induced losses, modulation changes, temperature-related issues, and age factors. Designing a lower bound of 42 dBmV helps alleviate noise induced from low-value taps (4, 8, 11 dB). Special equalized taps and feeder equalizers with return frequency pads are available to achieve this.

Levels

The CMTS upstream channel uses the equivalent of a bandpass filter that is tuned to the operator-configured frequency for the upstream port; it measures occupied channel power of the 3.2 MHz signal. If the upstream receiver is tuned to 30 MHz in the CMTS configuration file with a 3.2 MHz channel-width, then it is only looking at 28.4 MHz through 31.6 MHz. The adjacent channel's energy is not a factor in the measurement.

Another issue concerning the upstream is that if you change the bandwidth from 1.6 to 3.2 MHz, the CNR will degrade by 3 dB—you will see double the amount of noise floor with the same average power. The CPE is not told to increase power, but to keep a constant average power, which does not follow power/Hz calculations.

Reasons Why Cisco Chose this Route

If doubling the bandwidth automatically increased the expected return level at the CMTS by three dB, the modem could reach maximum transmit power in some situations.

1. If you change the channel width from 1.6 MHz to 3.2 MHz, you may want the carrier to be three dB hotter. You can change the input level on the CMTS from 0 to +3 dBmV, which will make the modems transmit three dB more. You will have the same CNR you had to begin with, because of double the noise bandwidth.
2. If you install three dB more padding in front of the CMTS upstream port, the modems will transmit three dB higher, but the total noise and interference on the front-end of the CMTS upstream port will be three dB lower. This is the preferred method, and carries a side benefit—configurations do not change. It is advantageous to use return-path combining networks (made by PCI technologies, QRF, ADC, and SA, among others). These include plug-in pads that can be changed without disconnecting cables.
3. Installing three dB padding in front of the reverse laser in the node provides less chance of laser clipping, less noise in the node, and the same overall affect at the CMTS. The disadvantage here is that you must adjust the node-receive levels for that one service. Set-top signals and telephony also require more work.

Remember to observe the transmit levels from the modems to see if they are reaching their maximum transmit power, especially the ones off of high-value taps. You can perform a “show cable modem remote-query” (if configured) to display the modem's outputs and a “show cable modem” to see if there is a “!” displayed next to the receive column, indicating the modem is at its maximum output power.

SNR or CNR

SNR is signal-to-noise ratio from the demodulator chip after the carrier has been eliminated to give a pure baseband, signal-to-noise ratio. CNR is carrier-to-noise ratio, before it reaches the chip (RF domain), and is measured with a spectrum analyzer.

People may question why they see an SNR reading of zero when modems are just coming online. The SNR is not accurate within +/-10 percent unless it has averaged over 10,000 valid symbols, or within +/-1 percent after 20,000 valid symbols. After that, the SNR reflects the moving average value.

As a result, the accuracy of the long-term channel SNR is within +/-1.5 dB for QPSK and +/-2 dB for 16-QAM when the SNR is in the range of 13 to 27 dB. By convention, a channel with 30 dB or higher SNR is considered a “clean channel.” Because of the way SNR is calculated, impulse noise and transient sources of noise cannot be detected. The formula used is better associated with white gaussian noise.

Note: In Cisco IOS[®] Software Release 12.1(10)EC1 and Cisco IOS Software Release 12.2(4)BC1b, the algorithm for calculating the SNR estimate was refined for a more accurate value. Depending on the RF plant characteristics, the new SNR estimate could be up to six dB lower than the values shown in earlier software releases. This value is only an estimate. For the most accurate value, use a spectrum analyzer or use the “show controllers” command with the Cisco uBR-MC16S Cable Line Card.

Another factor in upstream SNR readings and error generation could be total power overload from ingress and other services. Section 4.2.12 of the DOCSIS 1.0 specification states: The maximum total input power to the upstream demodulator MUST NOT exceed 35 dBmV in the 5-42 MHz frequency range of operation. Some systems have experimented with bandpass filters to eliminate extra noise below 20 MHz, CB at 27 MHz, or Ham radio at 28 MHz.

Upstream Modulation Profiles

Modulation profiles define how information will be transmitted upstream from a cable modem to the CMTS. Many upstream modulation profile variables can be changed, such as guard time of the burst, preamble, modulation (QPSK or 16-QAM), and FEC protection. Cisco has created three default profiles-qpsk, 16qam, and mix-to eliminate confusion, but changes may be necessary depending on the application.

Note: There may be slight differences between Cisco IOS Software trains and versions. DOCSIS 1.1-based code (BC train) uses shortened last codeword as the default setting. DOCSIS 1.0 code (EC train) uses fixed last CW as the default setting.

The default modulation profiles can be inefficient, depending on the DOCSIS Extended Header being used. These modulation profiles were optimized for five-byte Extended Headers. An inefficiency occurs when Cisco modems add one extra null byte to the Extended Header (Cisco modems do this for even alignment on a word boundary). This can have drastic affects. It is not apparent if this only affects Cisco modems-for example; Toshiba modems use five-byte Extended Headers. More testing with multiple vendors is required.

Note: Piggybacking bandwidth requests requires an Extended Header and an Extended Header is also required if using BPI+ security.

Tip: If not explicitly assigned with a modulation profile, each upstream port on a Cisco CMTS is assigned modulation profile 1 (qpsk) by default. Up to eight profiles can be configured. It is recommended to not change modulation profile 1. If more profiles are needed, start with number 2.

Modulation Profile Tutorial

Looking at the default settings of modulation profile 1, here is what would be needed:

1. Under global configuration, type: "cable modulation-profile 1 qpsk"
2. Under the appropriate interface (cable 3/0), type: "cable upstream 0 modulation profile 1" (or leave it blank, as the default is modulation profile 1)
3. The actual profile when entered and viewed in the "show run" command looks like this, but only the Short and Long interval usage codes (IUCs) for profile 1 may be displayed

Table 5

Current "Inefficient" Modulation Profile 1

IUC	FEC T bytes	FEC CW	Max B	Guard Time	Mod Type	Scramble	Scramble Seed	Diff Enc	Preamble Length	Last CW	UW
cable modulation-profile 1 short	5	75	6	8	qpsk	scrambler	152	no-diff	72	fixed	uw8
cable modulation-profile 1 long	8	220	0	8	qpsk	scrambler	152	no-diff	80	fixed	uw8

The "Show Cable Modulation-Profile" command produces the following output:

Table 6

Current "Inefficient" Modulation Profile 1

Mod IUC	Type	Preamble Length	Diff Enco	FEC T bytes	FEC CW	Scrambl Seed	Max B	Guard Time	Last CW	Scrambler	Preamble Offset
1 Request	qpsk	64	no	0x0	0x10	0x152	0	8	no	yes	952
1 Initial	qpsk	128	no	0x5	0x22	0x152	0	48	no	yes	896
1 Station	qpsk	128	no	0x5	0x22	0x152	0	48	no	yes	896
1 Short	qpsk	72	no	0x5	0x4B	0x152	6	8	no	yes	944
1 Long	qpsk	80	no	0x8	0xDC	0x152	0	8	no	yes	936

As you can see, the fields are not in the same places. The Unique Word (UW) setting is not visible. You can see the Preamble Offset, which is not set, but calculated, based on what is set for the UW.

Each column is described below:

- IUCs such as short, long, req, init, station, etc. These are also known as information elements. The first three IUCs are for maintaining modem connectivity, while Short and Long IUCs are for actual data traffic.
- Modulation Type such as 16-QAM or QPSK. This is expanded for DOCSIS 2.0.
- Preamble Length in bits <2-512>. 16-QAM is usually double the Preamble Length over QPSK.
- Diff enco = diff encoding enabled; no-diff = diff encoding disabled. Always use no-diff encoding.
- FEC T bytes are entered as decimal <0-10>, but shown in hex. $2 * \text{FEC T bytes size} = \text{bytes of FEC in each FEC CW (codeword)}$. Zero indicates no FEC. You can also disable FEC on the interface of each individual upstream port. This has been expanded to 16 for DOCSIS 2.0.
- FEC CW = codeword length information bytes (k) entered in decimal <16-253>, but shown in hex.

Note: When using a shortened last codeword, the last codeword must be greater than or equal to 16 bytes. If less than 16 bytes, filler bytes are added to make it 16. A full codeword is $k + 2 * T$, and must be less than or equal to 255 bytes total.

- If no FEC is used, “codeword” has no meaning
- Scrambl seed is listed in hex <0-7FFF>. Do not change this.
- Max B is the maximum burst size in minislots <0-255>. Zero means no limit. Any burst less than or equal to the amount of bytes represented by the maximum burst will use this IUC.
- Guard Time is listed in symbols <0-255>. DOCSIS says this needs to be at least five symbols. QPSK has two bits per symbol and 16-QAM has four bits per symbol.
- Last CW of fixed = fixed last codeword; shortened = shortened last codeword and will say “yes” in the column. Shortened eliminates extra “stuffing.”
- Scrambler = scrambler-enabled and no-scrambler = scrambler-disabled. Always keep scrambler enabled.
- Preamble Offset is not entered into the configuration. It is calculated when you enter in the UW value of eight or 16. Preamble Offset plus Preamble Length will equal 1024, 768, 512, or 256 bits for UW16; if not, you can assume UW8 is being used.
 - The UW is entered in the configuration of a profile, but doesn’t show up in the “Show” command. UW16 = Detect on a 16-bit unique word; UW8 = Detect on an eight-bit unique word.

Caution: Be sure to use UW16 when using 16-QAM for Short or Long IUCs. Using UW8 with 16-QAM can cause uncorrectable FEC errors to increment. Observe the “Show Cable Hop” command to verify.

Modulation Profile 3 (mix) Example

1. Under global configuration, type: “cable modulation profile 3 mix”
2. Under the appropriate interface (cable 3/0), type: “cable up 0 modulation profile 3”
3. The actual profile when entered and displayed on a “show run” command looks like this:

Table 7

Current "Inefficient" Modulation Profile 3

IUC	FEC T bytes	FEC CW	Max B	Guard Time	Mod Type	Scramble	Scramble Seed	Diff Enc	Preamble Length	Last CW	UW
cable modulation-profile 3 request	0	16	0	8	qpsk	scrambler	152	no-diff	64	fixed	uw16
cable modulation-profile 3 initial	5	34	0	48	qpsk	scrambler	152	no-diff	128	fixed	uw16
cable modulation-profile 3 station	5	34	0	48	qpsk	scrambler	152	no-diff	128	fixed	uw16
cable modulation-profile 3 short	6	75	6	8	16qam	scrambler	152	no-diff	144	fixed	uw8
cable modulation-profile 3 long	0	220	0	8	16qam	scrambler	152	no-diff	160	fixed	uw8

The "show cable modulation-profile 3" command will look like this:

Table 8

Current "Inefficient" Modulation Profile 3

Mod IUC	Type	Preamble Length	Diff Enco	FEC T bytes	FEC CW	Scrambl Seed	Max B	Guard Time	Last CW	Scrambler	Preamble Offset
3 Request	qpsk	64	no	0x0	0x10	0x152	0	8	no	yes	0
3 Initial	qpsk	128	no	0x5	0x22	0x152	0	48	no	yes	0
3 Station	qpsk	128	no	0x5	0x22	0x152	0	48	no	yes	0
3 Short	qam	144	no	0x6	0x4B	0x152	6	8	no	yes	0
3 Long	qam	160	no	0x8	0xDC	0x152	0	8	no	yes	0

Note: Notice in the display above that the Preamble Offset indicates 0. The Preamble Offset will not show up until you assign this modulation profile to an upstream port.

Tip: Decrease the minislot size from eight ticks to four. This will keep the number of bytes in a minislot closer to 16 when you use the more complex modulation scheme. If the minislot size is left at eight ticks, the minimum burst sent will be at least 32 bytes. This is inefficient when sending upstream requests, which only require 16 bytes total. Refer to Appendix B for minislot configuration.

DOCSIS 1.0-Based Code (EC and Earlier Cisco IOS® Software Trains)

Consider Cisco modems with six-byte Extended Headers and using all current Cisco CMTS defaults in the EC code, such as 1.6 MHz channel width, minislot size of eight ticks (16 bytes), and the modulation profile shown below:

```
cable modulation-profile 1 short 5 75 6 8 qpsk scrambler 152 no-diff 72 fixed uw8
```

If sending 64-byte Ethernet frames (46-byte packet data unit [PDU] + 18-byte Ethernet Header) on the upstream, the modem uses a Long burst and the total packet size becomes 256 bytes. This will be 16 minislots. See Appendix A for the calculations. This is inefficient for a 46-byte PDU. The PPS rate for 64-byte packets will drop because of this. Concatenation can help with the upstream throughput when sending 64-byte packets, but sending extra bytes wastes time.

This inefficiency could affect downstream TCP flows, because this will also be true for a TCP "acknowledgment" on the upstream. Even though an acknowledgment is less than 46 bytes, it will be padded to make it at least 46. Upstream concatenation can help, but it is still inefficient to send 256 bytes when only 96 are typically needed.

If the Extended Header is only five bytes, the modem will use a short grant at six minislots, for a total of 96 bytes. This is a difference of 160 bytes (256-96).

The fix to this dilemma with modulation profile 1 (qpsk):

1. Increase the FEC codeword size from 75 to 76 for the Short IUC.
2. Decrease the FEC T bytes from five to four for the Short IUC.
 - a. If the minislot size is changed from the default of eight ticks to four, make sure the Max Burst Size field for the Short IUC is changed from six to 12.
3. Shortened last codeword is recommended for the Short and Long IUCs.
 - a. Modems with older code may have to be upgraded since they may not register when using shortened last codeword in the IUCs.
4. If you want FEC to be high, increase it to ten and change the Max Burst field from six to seven.
 - a. If the minislot size is changed from the default of eight ticks to four, use eight T bytes of FEC and make sure the Max Burst field for the Short IUC is changed to 13.

Recommended profiles assuming eight-tick minislots at 1.6 MHz, or four ticks at 3.2 MHz:

Table 9

Recommended Modulation Profile 1

IUC	FEC T bytes	FEC CW	Max B	Guard Time	Mod Type	Scramble	Scramble Seed	Diff Enc	Preamble Length	Last CW	UW
cable modulation-prof 1 short	4	76	6	8	qpsk	scrambler	152	no-diff	72	short	uw8
cable modulation-prof 1 long	8	220	0	8	qpsk	scrambler	152	no-diff	80	short	uw8

Looking at the “Mix” profile defaults and the same situation as above, 46-byte PDUs will use 288 bytes total. This is even worse than the QPSK example because of more Preamble and Guard Time.

The fix to this dilemma with modulation profiles 2 (qam-16) and 3 (mix):

1. Increase the FEC codeword size from 75 to 76 for the Short IUC.
2. Increase the FEC T bytes from six to seven for the Short IUC.
3. Increase the Max Burst field from six to seven.
4. Be sure to use UW16 when using 16-QAM for Short or Long IUCs.
5. Shortened last codeword is recommended for the Short and Long IUCs.
 - a. If you have old code on some modems and you enable shortened last codeword in the modulation profile, it may not register. You will need to upgrade the modem code.
6. The FEC T bytes can be increased on a Long IUC from eight to nine when using 16-QAM.

Recommended profiles assuming four-tick minislots at 1.6 MHz, or two ticks at 3.2 MHz:

Table 10

Recommended Modulation Profile 3

IUC	FEC T bytes	FEC CW	Max B	Guard Time	Mod Type	Scramble	Scramble Seed	Diff Enc	Preamble Length	Last CW	UW
cab modulation-prof 3 short	7	76	7	8	16qam	scrambler	152	no-diff	144	short	uw16
cab modulation-prof 3 long	9	220	0	8	16qam	scrambler	152	no-diff	160	short	uw16

Conclusion

It is imperative to understand how all the variables such as minislot size, channel width, modulation, and Max Burst size work together. Keeping the minislot size to a minimum adds better resolution between minislot usage. The current default settings from the factory may not be optimized for all situations.

Manufacturers are finally specifying the performance of certain components with digital loading. These parameters include group delay at various frequencies, dynamic range and network bit error ratio (BER) performance for a particular modulation scheme, and component configuration.

Discrete amplifiers have the advantage of less power consumption than hybrids. On the other hand, hybrid-return amplifiers usually exhibit better distortion performance and temperature stability when compared with a discrete amplifier of the same gain. Hybrid amplifiers have a better dynamic range than discrete/single-ended amplifiers, which can overload and cause harmonics. Choose your equipment wisely.

A recent idea is assigning more power to less robust carriers and less power to the more robust carriers. For example, instead of just allocating power based on the bandwidth it uses, allocate power based on its robustness to noise. If you have a 1.6-MHz wide carrier using QPSK, it would get less power assigned than a 1.6-MHz wide carrier using 16-QAM. (This will also depend on which service is more critical to your bottom line.)

Whatever options you chose, time is always critical. Implement some type of proactive maintenance and monitoring plan to stay abreast of issues that can arise—even on new cable systems. About 90 percent of DOCSIS problems are attributed to physical layer and/or Cisco IOS Software configuration issues. Physical layer issues such as upstream levels, ingress, and noise must be understood and controlled. Knowledge of cable modem ranging processes and Cisco IOS Software configuration features can help mitigate these potential problems as well.

Additional References

Multimedia Traffic Engineering for Cable Networks

Combined Wisdom on the Upstream

Successfully Deploying Data over Cable Networks: An Overview

Business Case for Two-Way Service Deployment over HFC Networks

Determining Radio Frequency or Configuration Issues on the CMTS

<http://www.cisco.com/warp/public/109/rf.html>

Appendix A

Total Packet Size Calculations for a 46-Byte PDU

QPSK, 1.6 MHz, eight-tick minislots example:

$$(8 \text{ ticks/minislot} * 6.25 \text{ usec/tick} * 1.28 \text{ Msym/s} * 2 \text{ bits/sym}) / (8 \text{ bits/byte}) = 16 \text{ bytes/minislot}$$

Using the default settings for modulation profile 1 as shown below:

```
cable modulation-profile 1 short 5 75 6 8 qpsk scrambler 152 no-diff 72 fixed uw8
cable modulation-profile 1 long 8 220 0 8 qpsk scrambler 152 no-diff 80 fixed uw8
```

46-byte Ethernet frame + 18-byte Ethernet header + 6-byte DOCSIS header + 6-byte DOCSIS Extended Header = 76 bytes. An FEC codeword size of 4B in hex equals 75 bytes. $76/75 =$ one full codeword needed and one leftover byte. If using the default setting of fixed last codeword, this would require two full codewords. $2 * (75 + 2 * 5) = 170$ bytes + 9 bytes of Preamble + 2 bytes of Guard Time = 181 bytes. {The Preamble was (72 bits) / (8 bits/byte) = 9 bytes. The Guard Time of eight symbols would be $(8 \text{ sym} * 2 \text{ bits/sym}) / (8 \text{ bit/byte}) = 2$ bytes.}

$181 / (16 \text{ bytes/minislot}) = 11.3125$ minislots needed. Round this up to 12. Since the default setting for Max Burst size for the Short IUC is six, the Long IUC has to be used. $76 \text{ bytes} / 220 \text{ byte FEC CW} = 1$ full codeword needed + $2 * 8 = 236$ bytes + 10 bytes of Preamble + 2 bytes of Guard Time = $248 \text{ bytes} / 16 = 15.5$. Roundup to $16 * 16 \text{ bytes/minislot} = 256$ bytes.

Using the “modified” modulation profile 1 shown below:

```
cab modulation-prof 1 short 4 76 6 8 qpsk scrambler 152 no-diff 72 short uw8
```

46-byte Ethernet frame + 18-byte Ethernet header + 6-byte DOCSIS header + 6-byte DOCSIS Extended Header = 76 bytes. An FEC codeword size of 76 means exactly one codeword will be needed + 2*T. We have $76 + 2 * 4 = 84$ bytes + 9 bytes of Preamble + 2 bytes of Guard Time = 95 bytes. $95 / 16 \text{ bytes/minislot} = 5.9375$ minislots needed. Round up to $6 = 6 \text{ minislots} * 16 \text{ bytes/minislot} = 96$ bytes.

Appendix B

Minislot Configuration

Setting the Minislot Size

It is best to set the minislot size to a value that will make it eight or 16 bytes. This is sometimes not achievable because the DOCSIS limit says the minislot must be at least 32 symbols.

The following table lists the channel width vs. the number of ticks allowed for a minislot

Table 11 .

Channel Width vs. Minislot Selection

Channel Width	Ticks Allowed			
.2	32	64	128	
.4	16	32	64	128
.8	8	16	32	64
1.6	4	8	16	32
3.2	2	4	8	16

The number of ticks allowed will be affected by the symbol rate (channel width) used on the upstream. The modulation used and the number of ticks per minislot will affect the total amount of bytes in a minislot.

To configure the minislot size:

```
ubr7246vxr(config-if)#cable upstream 0 minislot-size 8
```

To verify the minislot size:

```
ubr7246vvr#show controllers c3/0 u0
```

Cable3/0 Upstream 0 is up

Frequency 24.848 MHz, Channel Width 1.600 MHz, QPSK Symbol Rate 1.280 Msps

Spectrum Group 1, Last Frequency Hop Data Error: NO(0)

MC16S CNR measurement: -26 dB

Nominal Input Power Level 0 dBmV, Tx Timing Offset 2952

Ranging Backoff automatic (Start 0, End 3)

Ranging Insertion Interval automatic (60 ms)

Tx Backoff Start 0, Tx Backoff End 4

Modulation Profile Group 2

Concatenation is disabled

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 = 8

Minislot size in symbols = 64

Bandwidth requests = 0xED97D0

Piggyback requests = 0x2DB623C

Invalid bandwidth requests = 0xE4B

Minislots requested = 0x12B17492

Minislots granted = 0x12B16E64

Minislot size in bytes = 16

Map Advance (Dynamic): 2468 usecs

UCD count = 3566700

DES Ctrl Reg#0 = C000C043, Reg#1 = 4016

DOCSIS 2.0

DOCSIS 2.0 actually relaxes this limitation somewhat:

6.1.2.4 Upstream Intervals, Minislots and 6.25-Microsecond Increments

Allowed single time-tick minislots for Advanced PHY-only channels.

Instead of the tick range of two to 128, it can have a range of one to 128. The symbols per minislot of 32 may be relaxed to 16, but only for the advanced modulation schemes in DOCSIS 2.0.



Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the
Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland
Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland
Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe