



Voice Quality: Deploying Tools for a Stable Infrastructure

Contents

Introduction	3
Voice Impairment Parameters	3
Requirements of a Voice Quality Measurement Product	4
Components of Cisco Unified Communications	4
Cisco Unified Service Monitor	5
Data Collection and Analysis	5
Cisco 1040 Sensor-Based Call Quality—R-Factor:	5
Cisco Voice Transmission Quality-Based Call Quality—K-Factor	11
Scalability	12
Preparing the Server for Service Monitor	12
Operating System and Server	12
Hostname	13
Verify Locale Settings	13
Verify ODBC Driver Manager	13
Browser Version and Flash Plug-in	13
NTP	13
Connectivity	14
Terminal Server Services	14
Antivirus and Platform Agents	14
Check Routing and Firewalls	14
Port Availability	14
TFTP Server	15
Cisco Unified Communications Manager Configuration	15
Service Monitor Installation	15
Preinstallation Checks	15
Installation Procedures	16
Licensing and Registration the Software	16
Uninstallation	16
Failover and Redundancy	17
Cisco 1040 Failover Mechanism	17
Cisco 1040 Failover Deployment	18
Preparing for Failover	18
Setting Up Failover	18
Setting Up Failover in Default Configuration	18
Setting Up Failover for a Specific Sensor	19
Viewing Failover Status	20
Cisco Voice Transmission Quality Redundancy	20
Backup and Restore	20
Configuring Low-Volume Schedule and Database Purging	21
Troubleshooting	22
Useful URLs Cisco.com URLs for Customers and Partners	24
Product Information	24
Evaluation Downloads	24
Patch Download	25
Miercom Review of Operations Manager and Service Monitor	25

Introduction

This white paper provides insight into voice quality and the different methods to measure voice quality.

Voice over IP (VoIP) has passed its infancy stage and is a mature technology that has been widely adopted by customers hoping to take advantage of the cost savings offered by VoIP in addition to a range of advanced features that improve efficiency. Voice quality is the qualitative and quantitative measure of sound and conversation quality on an IP phone call. Voice quality measurement describes and evaluates the clarity and intelligibility of voice conversation.

The shift from the traditional Time-division Multiplexing (TDM) world to a packet-based IP telephony solution poses challenges for voice quality. Unlike data, which is bursty in nature and tolerant to delay and packet loss, voice and video are extremely sensitive to jitter, packet loss, and delay. In a converged network with voice, video, and data residing on the same network, there is a huge demand for the network infrastructure to be reliable and scalable and to offer different levels of service for advanced technologies such as voice, video, wireless, and data.

Voice Impairment Parameters

The real-time nature of voice drives strict service-level agreements (SLAs) to be implemented in the network. The primary voice impairment parameters are jitter, packet loss, and delay.

Packet Loss

In data networks, even if a few packets are lost during transmission, TCP ensures the retransmission and assembly of the packets, and the user will not notice any difference. But when transmitting voice packets across the IP backbone, the missing packets cause distortion in voice quality on the receiving end, and retransmission of missing voice packets is useless. It is tolerable to have occasional packet loss, but consecutive loss of voice packets can affect the overall quality of the transmitted voice.

Jitter

Delay variation, or jitter, occurs when voice packets arrive at the destination at different time intervals. This can happen because of the connectionless nature of IP. Depending on the congestion and load on the network, the arrival rate of these packets at the destination may vary. The devices on the receiving end should be capable of buffering these packets and playing them back to the user at a consistent interframe interval. These types of devices are called dejitter buffers. A dejitter buffer usually adds a forced delay (default 60 milliseconds [ms]) to every VoIP packet received. Typically, this delay is in the 20 to 60 ms range. This delay is commonly called the play out delay.

Delay

Delay is the finite amount of time it takes a packet to reach the receiving endpoint after being transmitted from the sending endpoint. In the case of voice, this is the amount of time it takes for a sound to travel from the speaker's mouth to the listener's ear. Delay (or latency) does not affect voice fidelity. Extended network delay is perceived as echo in the conversation. Even though network delay is not a direct cause of echo, it does amplify the perception of any echo present in the media path. Extremely long delays can lead to "collisions" in the conversation, when both parties seem to be speaking simultaneously.

The network infrastructure must meet the following requirements:

- Packet loss must not be more than 1 percent
- Average one-way jitter must not be more than 30 ms
- One-way delay must be under 150 ms

To help ensure good voice quality, it is imperative to keep jitter and packet loss under control by paying close attention to voice impairment factors.

Requirements of a Voice Quality Measurement Product

By now, you have a basic understanding of the voice impairment parameters and their importance in voice quality measurement. Some of the key requirements for a voice quality measurement product follow. It must:

- Be able to calculate voice quality for actual calls. Although it is possible to generate synthetic voice traffic and calculate voice quality for these calls, the voice quality generated from synthetic calls does not represent end-user experience. The capability to simulate voice traffic must be made use of to verify voice quality when a real-time alert is received.
- Provide details of voice impairment parameters such as jitter, packet loss, and so on. It is important to understand the cause of voice quality degradation; for example, knowing whether jitter or packet loss is causing the problem will help in fine-tuning the network, if required.
- Report voice quality in real time. For example, in a call that lasts for 5 minutes, if the user encounters voice quality problems in the second minute, it is important for the administrator to get an alert at the second minute rather than at the end of the conversation.
- Provide details about the endpoints involved in a conversation, the type of codec used, and the IP addresses and phone numbers of the endpoints. Such detailed information is important in troubleshooting a voice quality problem.
- Be scalable and easily deployable.
- Be nonintrusive.

Components of Cisco Unified Communications

The Cisco[®] Unified Communications Family of products can be divided into two parts:

- Cisco Unified Communications infrastructure
- Application infrastructure

The Cisco Unified Communications infrastructure consists primarily of call-processing devices such as Cisco Unified Communications Manager, Cisco Unified Communications Manager Express, Cisco Unity[®], and Cisco Unity Express. The Cisco Unified Communications infrastructure layer is the brain of the Cisco Unified Communications Family of products, and it performs the critical functions of call setup, call routing and call tear-down.

The application infrastructure consists of software applications such as Cisco Emergency Responder, MeetingPlace[®], MeetingPlace Express, IP Contact Center, IP Contact Center Express, and Personal Assistant. These applications provide added functionality to the Cisco Unified Communications family of products.

It is imperative to have a rock-solid network infrastructure. The network infrastructure offers the platform for advanced technologies such as voice, video, wireless, and so on.

Cisco Unified Service Monitor

Given the dynamic nature of IP networks and the strong dependency of the Cisco Unified solution on network infrastructure, it is imperative for network administrators to have voice quality information on real calls (not simulated calls) at their fingertips to help enable them to resolve voice quality problems. Cisco Unified Service Monitor meets user requirements in reporting voice quality issues.

There are two ways to measure call quality using Service Monitor:

- Call quality measurement and reporting as the call progresses: Use the Cisco 1040 Sensor hardware.
- Call quality measurement and reporting at the end of the call: Use the Cisco Unified Communications Manager Cluster (4.2.x and 5.x and above provide Cisco Voice Transmission Quality–based mean opinion scores [MOSs]; for earlier call manager versions, Service Monitor reports jitter and packet loss for the call).

Service Monitor analyzes the data that it receives and sends traps when a MOS falls below a threshold. Service Monitor provides a set of default global thresholds, one per supported codec. It helps enable you to change the default global thresholds and to override them by creating sensor threshold groups and cluster threshold groups.

Data Collection and Analysis

Service Monitor receives and analyzes MOSs from the following sources when they are installed and configured properly in your voice network:

- **Sensors:** Cisco 1040 Sensors compute MOSs for each Real-Time Transport Protocol (RTP) stream and send syslog messages to Service Monitor every 60 seconds.
- **Cisco Voice Transmission Quality:** Cisco Unified Communications Manager collects data from Cisco H323 Voice Gateways and Cisco IP phones; MOSs are calculated on the gateways and phones using the Cisco Voice Transmission Quality algorithm. At the termination of a call, Cisco Unified Communications Manager stores the data in call detail records (CDRs) and call management records (CMRs).

The Cisco Unified Service Monitor 2.0 software runs on an Intel-based machine, with a server running Windows 2003 with Service Pack 1. The software must be procured by the customers.

Cisco 1040 Sensor-Based Call Quality—R-Factor:

The ITU E model, as defined in G.107 (03/2003), predicts the subjective quality that is experienced by an average listener by combining the impairment caused by transmission parameters (such as loss and delay) into a single rating: the transmission rating factor R (the R-factor). This rating, expressed on a scale of 0 (worst) to 100 (best) can be used to predict subjective user reactions, such as the MOS. The MOS can be obtained from the R-factor with a converting formula. Thus the R rating is an estimate of the quality that can be expected if the network is realized the way it is planned.

Cisco 1040 Sensor Deployment

The Cisco 1040 Sensor is the hardware component that will be deployed on the Switched Port Analyzer (SPAN) port of a switch, as close to the IP phones and other problem areas (gateways, and so on) as possible,

The sensor has two Fast Ethernet interfaces:

- Management port
- SPAN port

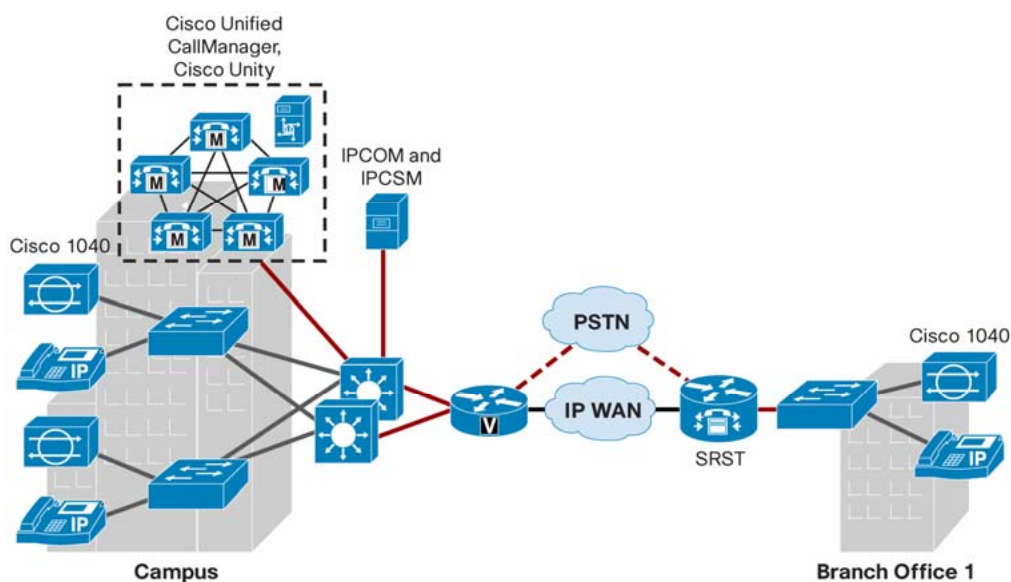
The sensor uses IEEE 802.3af standard Power over Ethernet (PoE) from the switch to which it connects. When the Cisco 1040 Sensor boots up, it uses Dynamic Host Configuration Protocol (DHCP) option 150 to retrieve its configuration and image files on a TFTP server. Similar to the way that an IP phone registers with call manager a Cisco 1040 Sensor registers (using Skinny Client Control Protocol [SCCP]) to the Service Monitor application. On the TFTP server, the Cisco 1040 first looks for its configuration file, named QOV[Cisco 1040 MAC address].CNF. If that file does not exist, the Cisco 1040 looks for a file named QOVDefault.CNF. These CNF files provide the image filename for the Cisco 1040 to download in addition to the Service Monitor IP addresses. The Cisco 1040 then downloads this image and registers to the Service Monitor, just like a phone registers to a Cisco CallManager, using SCCP.

Then the Cisco 1040 Sensor utilizes the SPAN port on a switch to monitor the actual voice calls. It collects voice impairment parameters such as jitter and packet loss from the RTP stream and computes MOS values. The sensor works in passive mode to collect voice impairment statistics.

The sensor reports voice quality details and MOS values every 60 seconds, providing near real-time voice quality measurement. Each Cisco 1040 can monitor 100 RTP streams with optimal SPAN port configuration.

As shown in Figure 1, multiple sensors can be deployed in the network and configured to register to the Service Monitor software component. Each instance of a Service Monitor software component can manage up to 50 sensors and report call quality for 30,000 IP phones.

Figure 1. Multiple Sensors Register to Cisco Unified Service Monitor



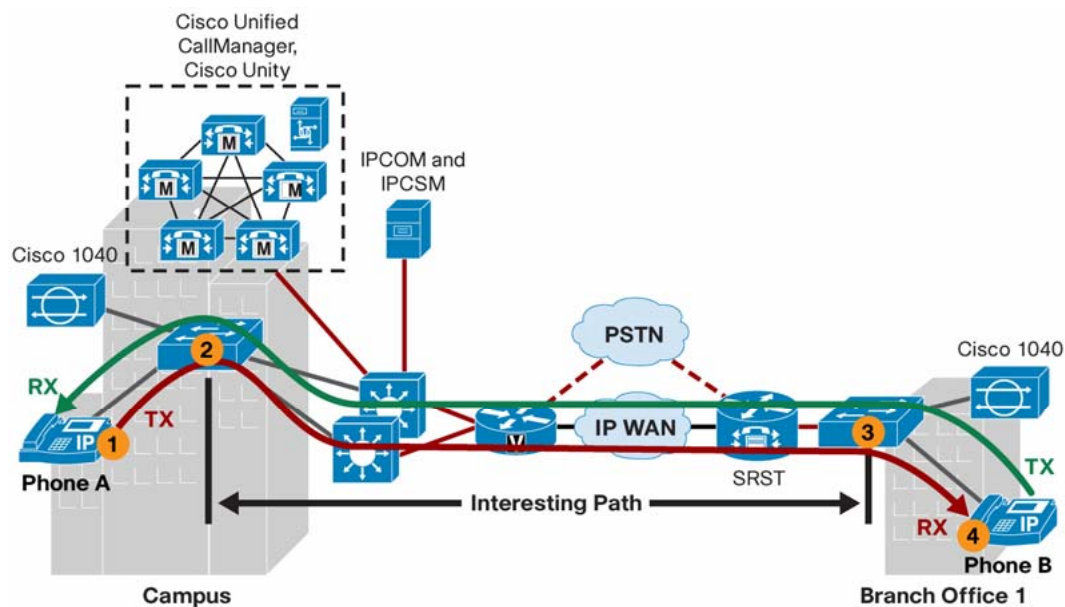
Note: In Figure 1, the Cisco Unified Service Monitor and Cisco Unified Operations Manager software instances co-reside on a single machine. This type of deployment is supported for small networks (fewer than 5,000 phones). For networks with more than 5,000 phones, the Service Monitor and Operations Manager software must run on separate machines.

Service Monitor satisfies most quality-monitoring requirements for enterprise IP telephony. Deployment strategies include:

- **Strategic monitoring:** The Cisco 1040 Sensor is installed to continuously monitor IP phones at some or all locations in the managed environment. Depending on monitoring goals, significant coverage of all or most sites could be included, or by using sampling techniques, representative sites would be selected for monitoring and would determine the location of the Cisco 1040 sensors. Service Monitor can scale to support up to 500 RTP streams a minute total from multiple Cisco 1040 sensors and can provide real-time alerting on call quality issues as well as providing information that can be used to evaluate general service levels and to validate performance of SLAs.
- **Tactical monitoring:** Cisco 1040 sensors can be inexpensively shipped overnight to a site (such as a branch office) that has voice quality concerns or problems. Once it is installed, it can immediately begin monitoring and assessing the quality of IP-based calls without elaborate setup or complicated installation. The Cisco 1040 Sensor is FCC Class B compliant and can easily be installed in any office environment.

Figure 2 shows centralized Cisco Unified Communications Manager deployment with one remote branch office connected through a WAN circuit. To monitor voice quality for calls across the WAN circuit, two Cisco 1040 Sensors must be deployed as shown. The key is to deploy the sensors switch as close to the IP phone as possible. In most cases, the sensors will sit on the access layer switch in the campus.

Figure 2. Centralized Cisco Unified Communications Manager Deployment



- For each phone, there are transmit and receive RTP streams.
- For the RTP stream originating from Phone A (TX RTP stream), the segment between 1 and 2 in Figure 2 experiences the least impairment, and the probability of voice quality degrading in this segment is slim to none. The RTP stream between segments 2 and 3 traverses several network devices and is prone to network conditions.
- The previous statement is also true for RTP streams originating from Phone B.
- For the sensor on the left in Figure 2, you can safely ignore the RTP stream from Phone A

covering segments 1 and 2. The RTP stream coming from Phone B covering segment 4, 3, and 2 is of importance; the switch on the left must be configured to span the incoming RTP stream, to span the destination port seen by the left sensor.

With optimal SPAN port configuration, each sensor can monitor 100 calls.

The number of sensors required depends on the busy hour call completion (BHCC) handled by the switch. For a 10,000-phone network, for example, the cluster could handle 1000 to 4000 call simultaneously. As the size of the network increases, it becomes more appropriate to take samples of the calls generated from the cluster and measure voice quality for a subset of these calls. On the other hand, if the network consists of only about 1000 phones, it is easier to monitor voice quality on all of the calls; however, the same sampling technique can also be applied to a 1000-phone network.

The Unified Communications deployment follows one of the following call processing models:

- Single site with centralized call processing
- Multiple-site WAN with centralized call processing
- Multiple-site WAN with distributed call processing

For a single site with centralized call processing, most often a Catalyst[®] 6500 is used in the access layer/wiring closet to connect the IP phones. It is common to expect about 200 IP phones (4 blades with 48 ports = approximately 200) on a single Catalyst 6500. If 4 out of 10 phones are active at any point, a single Cisco 1040 Sensor can be placed on the SPAN port of the Catalyst 6500 to monitor the active calls. It is possible to deploy multiple sensors on the same switch to address situations in which the switch is handling high call volume.

It is not necessary to measure voice quality for every call. The general practice is to measure voice quality for a subset of the calls on a continuous basis and use a tactical approach for troubleshooting voice quality problems.

Based on this analysis, for a 1000-phone deployment, if you were to sample 30 percent of the active calls and measure voice quality on a continuous basis, you would need three Cisco 1040 Sensors (30 percent of 1000 is 300, and each sensor can monitor 100 RTP streams; hence, you would need three sensors). Usually, the sensors are deployed in pairs (one each at the origination and termination endpoints).

This sampling could be quite aggressive for most common deployments. Based on the simultaneous calls that are active on a switch, the number of sensors required for voice quality measurement varies. As the network size increases, the sampling policy can be reduced.

The previous scenario had a dense population of phones on a single switch; another scenario could have phones scattered among multiple smaller switches. Typically, the number of phones in a branch is less, and you will see smaller-density switches used in the branch. For this scenario, it is excessive to have one sensor per switch, especially if the number of RTP streams on these switches is small (20–40 RTP streams). To address this issue, you can use Remote SPAN (RSPAN) and combine RTP streams on multiple switches. Another alternative is to use an active hub and connect the SPAN destination port from multiple switches to the same active hub. With this alternative, there is potential for Layer 2 loops, and you must evaluate the best option before embarking on any approach.

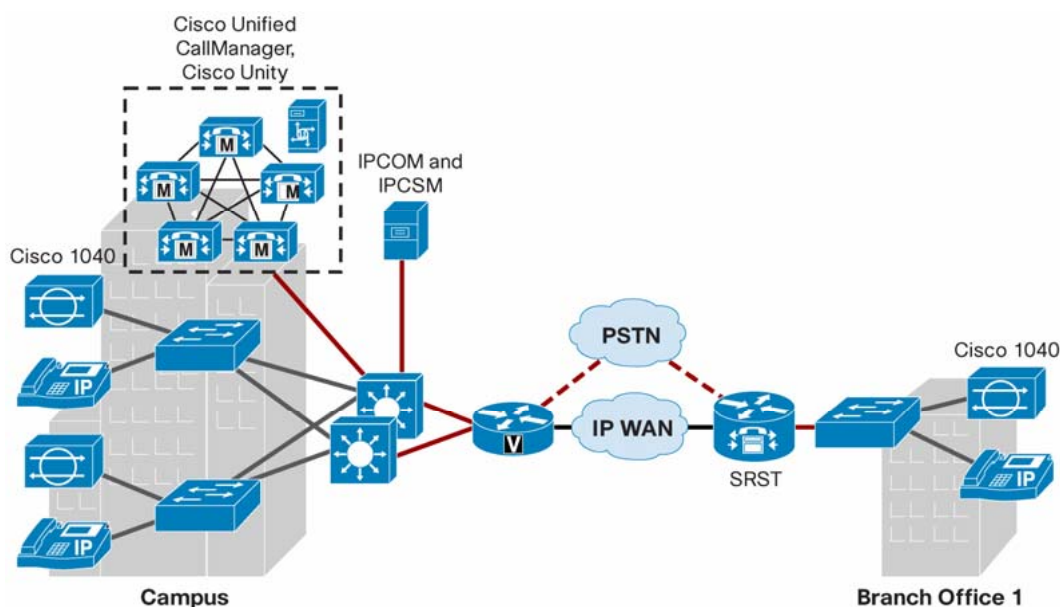
In summary, Service Monitor can monitor actual voice calls in real time and provide details of the parameters that cause voice quality degradation. The combination of Service Monitor and

Operations Manager provides a powerful product for monitoring and troubleshooting voice quality problems.

Placement of Cisco 1040 Sensors

The Cisco 1040 Sensor is FCC Class B certified and can be deployed in a wiring closet or on a desk. The Cisco 1040 uses the SPAN port on the switch to monitor the RTP stream. As shown in the following diagram, the sensors are deployed as close to the IP phone as possible so that the voice quality measurement will be close to what the user experiences. The sensor reports voice quality measurements every 60 seconds. For each conversation, there are four RTP streams: two each from originating and terminating phones. In the default SPAN port configuration, Service Monitor receives four MOS values every 60 seconds for each conversation. As discussed in the previous section, of the four RTP streams, two provide meaningful statistics; hence, the MOS value calculated for the interesting RTP stream is of importance and should be considered for further analysis.

Figure 3. Cisco 1040 Placement



The number of sensors per switch depends on the following:

- Type of switch
- Type of customer
- Number of simultaneous calls

The type of switch determines the number of SPAN destination ports that can be configured on the switch. Modular switches such as the 65xx support two SPAN destination ports with different source ports. Most of the fixed configuration switches support a single destination port. On the modular configuration switches, you can have two sensors deployed on the same switch, and on the fixed configuration switch, you can have a single sensor deployed.

A typical enterprise customer has various organizations, such as engineering, human resources, marketing, sales, and support. The traffic patterns are different in these organizations. You can expect more calls for support, sales, and marketing groups, and the switches that house these users will have a higher busy hour call attempt (BHCA) value. In a call center environment, you can expect high call volume and, therefore, high BHCA, which will be the deciding factor for the number of sensors required.

The number of simultaneous calls is tied to the previous argument about the BHCA value. It is important, therefore, to understand the BHCA value and average call hold (ACH) time to determine the number of sensors required. Keep in mind the limit on the number of RTP streams supported by each sensor (100) and the number of simultaneous calls generated by the phones connected to the switch where the sensor is deployed.

You could have a situation in which the number of RTP streams exceeds 100, in which case you can add sensors or configure the SPAN source port in such a way that you monitor only selected phones on the switch. If the SPAN destination port sees more than 100 RTP streams, the sensor goes into sampling mode and the MOS value reported is diluted.

Note: 1040 bandwidth requirements are very little. Given that syslog is around 60 bytes per stream, so max syslog each minute would be 6000 bytes when reporting 100 streams each minute. Skinny keep alives are every 30 seconds and a heartbeat message only has about 12 bytes of TCP payload for each KA, and KA-ACK.

Cisco 1040 Sensor in Sampling Mode

A Cisco 1040 Sensor can monitor 100 RTP streams. If it is deployed on a switch that has more than 100 RTP streams, it performs sampling, in which case some of the RTP streams are not considered for MOS value generation. You must avoid this situation at all times. In sampling mode, the MOS value reported is diluted because some of the RTP streams are not considered. The sensor monitors the RTP stream and collects the information necessary to compute the MOS value. This information is stored in a buffer, from where the computation process picks data to compute the MOS value. If the packets arrive at a faster rate than the rate at which the buffer is emptied, part of the RTP stream is dropped before the sensor starts the collection process. Keep in mind that CPU resources are constantly utilized; hence, it is not just the buffer that becomes a bottleneck when the sensor is overwhelmed with more RTP streams: the CPU also falls short in serving the different processes. The MOS value reported by the sensor gets worse as the number of simultaneous RTP streams increases beyond 100. It is important to plan ahead and optimize the SPAN port configuration in these scenarios.

Cisco 1040 Sensors in the Branch

In a branch office, the density of IP phones is less when compared with the density seen in a campus. Typically, the branch office contains fixed-configuration switches, and the number of simultaneous calls is lower.

In a fairly large branch, it is common to see multiple fixed-configuration switches stacked to provide more density and avoid the need to run gigabit uplink to aggregate switches and routers. The Cisco 1040 Sensor fits into this model the same as with any other switch. It utilizes the SPAN port to monitor the RTP stream.

In a scenario in which the switches are not stacked but have gigabit uplink to an aggregate switch, if the number of RTP streams is below 100, then one sensor per switch is excessive. In this situation, RSPAN is useful. The configuration done on the switch with SPAN, RSPAN, or ESPAN is transparent to the sensor; the sensor functions normally as long as it sees the RTP stream.

For cases in which RSPAN is not a desirable configuration or not an approved configuration, a simple active hub can be used to connect the individual SPAN port from the various switches, and the sensor can be deployed on the hub. It is very important to keep spanning tree loops in mind when such a configuration is attempted. The use of a hub must be selected as a last resort.

SPAN Port Limitation

SPAN ports are widely used to connect packet sniffers for troubleshooting. In the contact center world, the SPAN port is used to record the voice conversation. In the service monitor world, the SPAN port is used to monitor voice quality. If the SPAN port is required for a packet sniffer, contact center, and Service Monitor at the same time, the SPAN port does not allow configuration of the same source port tied to multiple SPAN destination ports. This is a limitation of SPAN port configuration. The only alternative is to use an active splitter that offers one-to-many streams. The simplest splitter must be an active hub that offers a one-to-many stream. In this model, the packet sniffer, contact center application, and sensor connect to the hub, and the hub connects to the SPAN destination port on the switch.

It is also recommended that all fax devices be aggregated in one voice gateway, and the 1040s not span the switch to which the gateway is connected. Currently the 1040 is known not to correctly handle certain cases of calls going to fax machines.

Cisco Voice Transmission Quality-Based Call Quality—K-Factor

K-factor (Klirrfaktor) is a mean opinion score (MOS) estimator of the endpoint type defined in ITU standard P.564. This standard relates to the testing and performance requirements of such a device. K-factor predates the standard. A P.564-compliant version will follow. K-factor is trained using thousands of speech samples and impairment scenarios, along with target P.862.1 MOS scores for each scenario. The trained K-factor device in the IP phone or gateway can then recognize the current impairment and produce a running MOS score prediction.

R-factor is based on three dimensions: loss, delay, and echo. K-factor and other P.564 MOS estimators measure only packet loss, which is a network effect. They are packet loss metrics projected onto a psychological scale. In general, primary statistics (packet loss, jitter, and concealment ratio) show visible degradation well before MOSs start to degrade. Hence, MOSs are a secondary indication of network problems, because it is essentially a packet loss metric. Packet loss counts, jitter, concealment ratio, and concealment second counters are primary statistics, based on direct observation. MOSs are a secondary statistic. Hence, you should use MOSs as a flag, but then use primary statistics to investigate or qualify the alarm. Use primary metrics in SLAs rather than MOSs.

Cisco Voice Transmission Quality-based call quality reports can be obtained using Service Monitor in conjunction with Cisco Unified Communications Manager and the latest Cisco Unified IP Phones and Cisco Gateways.

When to Use Cisco Voice Transmission Quality-Based Call Quality Reporting

Some key points to keep in mind when choosing to go with Cisco Voice Transmission Quality-based call quality reporting are:

- Cisco Voice Transmission Quality is supported from Cisco Unified Communications Manager 4.2 or later versions.
- Cisco 7906, 7911, 7931, 7921, 7962-G, 7962-G/GE, 7942-G, 7942-G/GE, 7972-G/GE, 7940, 7960, 7941, 7961, 7970, and 7971 IP Phones support Cisco Voice Transmission Quality in SCCP and Session Initiation Protocol (SIP) mode. (You must have new firmware; the firmware can be downloaded from Cisco Unified Communications Manager 4.2 or 5.x.)
- No Cisco IP Phones, including 7985, support Cisco Voice Transmission Quality
- Sampling rate is every 8 seconds
- Score is sent at the end of the call by using CMRs

Unlike Cisco 1040-based call quality reporting, Cisco Voice Transmission Quality–based call quality is reported at the end of the call. You can use Cisco Voice Transmission Quality–based reporting if you prefer not to have call quality reporting as the call progresses. Also, the Cisco Voice Transmission Quality feature is inherent to Cisco Unified Communications Manager 4.2 and later. Therefore, if you do not want to invest in a Cisco 1040, Cisco Voice Transmission Quality–based call quality still provides MOSs to estimate user experience.

Scalability

Service Monitor scalability is dependent on the estimated call rates generated in the network. The call volume supported by a fully loaded Service Monitor 201 system with various scenarios is given in Table 1.

Table 1. Call Volume Supported by a Fully Loaded Service Monitor 201 system

Scenario	Cisco Voice Transmission Quality Calls/ Minute	Sensor Segments/Minute
Cisco Voice Transmission Quality + Sensor	250	1500
Cisco Voice Transmission Quality Only	500	N/A

Preparing the Server for Service Monitor

This section describes how to prepare your server for Service Monitor installation

Operating System and Server

Service Monitor is supported on Windows 2003 Server SP1/ SP2—Standard Edition and Windows 2003 Server SP1/ SP2—Enterprise Edition. No other operating systems are supported. It is recommended that software other than the operating system and antivirus software not be installed on this computer system.

Server suggestions:

- For a small network (under 5000 phones), Serial ATA (SATA) disks are required.
- For a medium network (5000 to 15,000 phones), SCSI disks are required (suggestion: MCS 7845-I1 or MCS 7845-H1 comes with SCSI).
- For networks with more than 15,000 phones, Serial Attached SCSI (SAS) disks are required (suggestion: MCS 7845-I2 or MCS 7845-H2 comes with SAS).

Hostname

It is recommended that you configure the hostname for the Service Monitor server before you start installing Service Monitor. Specify the hostname when you are installing the operating system or subsequently, using My Computer > Right Click and select Properties > Computer Name.

Once Service Monitor is installed, changing the hostname is a very laborious process involving file manipulation and the execution of scripts. The User Guide for Cisco Unified Service Monitor documents all the steps involved in changing the hostname of the Service Monitor server.

Verify Locale Settings

Service Monitor supports only the U.S. English and Japanese locales. Using other locales means that you are running on an unsupported configuration. Further, Service Monitor may display erratic behavior, such as JRunProxyServer services not starting automatically. However, non-U.S. English keyboard layouts should work.

Verify ODBC Driver Manager

Some components of Service Monitor require the presence of the correct version of Open Database Connectivity (ODBC) on the Service Monitor server.

To verify the ODBC Driver Manager version, do the following:

- Step 1. On the Service Monitor server, select Start > Settings > Control Panel > Administrative Tools > Data Sources (ODBC).
- Step 2. Click the About tab.
- Step 3. Make sure that all ODBC core components have the same version number (3.5xx or later). ODBC is not available from Microsoft as a standalone installation but is packaged along with Microsoft Data Access Component (MDAC).

Note: If the necessary ODBC is not listed, install MDAC 2.5 or later by referring to the Microsoft Website.

Browser Version and Flash Plug-in

The recommended browser is Microsoft Internet Explorer 6.0/7.0

When using Service Monitor, disable any software on your desktop that you use to prevent popup windows from opening. Service Monitor must be able to open multiple windows to display information.

NTP

The clocks on Service Monitor and call manager servers must be synchronized for Service Monitor reports to include complete and up-to-date information and accurately reflect activity during a given time period. These notes offer a starting point and do not provide complete instructions for configuring Network Time Protocol (NTP).

To get started:

1. Talk with your call manager administrators to determine the time server with which Service Monitor should synchronize. You might find Cisco IP Telephony Clock Synchronization: Best Practices, a white paper on Cisco.com, useful; you can read it at http://cisco.com/en/US/products/sw/voicesw/ps556/prod_white_papers_list.html.

2. Use your system documentation to configure NTP on the Windows Server 2003 system where Service Monitor will be installed. Configure NTP with the time server being used by call manager in your network. You might find How to configure an authoritative time server in Windows Server 2003 at <http://support.microsoft.com/kb/816042> useful.

Connectivity

Ensure that the Service Monitor server can reach the TFTP server and the phones or devices in the IP address range where the Cisco 1040 would be deployed.

Terminal Server Services

Remote Desktop Service and/or Virtual Network Computing (VNC) Services are recommended to remotely manage the Service Monitor server. VNC Services and Remote Desktop can be used to remotely install the Operations Manager and Service Monitor software.

Antivirus and Platform Agents

You should enable virus protection on the Service Monitor server, using antivirus software. Active scanning of drives and memory should be performed during off-peak hours. Please exclude from scanning the “CSCOPx” folder. You may experience delays, and performance may be degraded, when the virus protection software is scanning all files. Service Monitor has undergone interoperability testing with the following:

- Third-party virus protection software:
 - Symantec Antivirus Corporate Edition Version 9.0
 - McAfee VirusScan Enterprise 8.0
- Platform Agents:
 - (Optional) Cisco Security Agent 4.0.3 (build 736)

Check Routing and Firewalls

Make sure that any firewalls between the Service Monitor server and the call manager, TFTP server, and Cisco 1040s are configured to allow management traffic through. See the “Port Availability” section below for information on which ports should be opened.

Also make sure that there is connectivity between devices and the Service Monitor server. Even if a route exists to a network behind a device, it does not mean that one exists to (and from) the device itself.

Port Availability

Table 2 lists the ports used by Service Monitor. These ports should not be scanned.

Table 2. Service Monitor Port Usage

Protocol	Port Number	Service Name
UDP	22	SSH
UDP	53	DNS
UDP	67 and 68	DHCP
UDP	5666	Syslog—Service Monitor receives syslog messages from Cisco 1040
TCP	2000	SCCP—Service Monitor uses SCCP to communicate with Cisco 1040s
TCP	43459	Database.
TCP	5665-5680	Interprocess communication between the user interface and back-end processes

TFTP Server

Cisco Unified Communications Manager 5.x or 4.2 can be used as the TFTP server.

If you use Cisco Unified Communications Manager as a TFTP server, Service Monitor cannot copy configuration files to Cisco Unified Communications Manager due to security settings on the latter. You will need to manually upload the configuration file. After uploading the configuration file, reset the TFTP server on Cisco Unified Communications Manager. For more information, see Cisco Unified Communications Manager documentation.

Cisco Unified Communications Manager Configuration

Service Monitor can collect and analyze data from Cisco Unified Communications Manager. Follow the configuration steps described in the user guide in the “Cisco Unified Communications Manager Configuration” section or go to

http://bi/push_targets1/ucdit/cc/td/doc/product/rtrmgmt/srvmon/srvmon20/usrgd/confccm.htm.

These are some important points that you should be aware of when you are configuring Cisco Unified Communications Manager:

- If you don't see call records in the Cisco Voice Transmission Quality reports, make sure CDR and CMR are enabled on all call manager nodes of a cluster.
- Make sure that the cluster ID is unique in the system; System Monitor will generate an error when a duplicate cluster is entered.
- Make sure that on the call manager Enterprise Parameter Configuration page the CDR File Time Interval is set to 1 (minute). This determines how frequently call manager generates CDRs. (Typically this is not changed.)
- Check the owner of the sm_record_create_table in the CDR database after adding call manager to the Service Monitor server. Make sure the table owner is dbo. Caution: call manager cannot write CDRs to the database if the sm_record_create_table owner is not dbo

Service Monitor Installation

This section describes the tasks that should be performed during Service Monitor installation.

Preinstallation Checks

Perform the following checks before installing Service Monitor:

- Dual homing (dual network interface cards [NICs]), using two different IP addresses, is not supported on Service Monitor. If, during installation, you receive a warning message to edit a file named gatekeeper.cfg, then your server is dual homed, and you must disable one of the network interface card (NIC) interfaces before adding any devices to Service Monitor. Using two NICs with a single IP address (a failover configuration, in case one of the NIC cards fails) is supported.
- Make sure that you change the default call manager cluster ID setting (located at call manager Administration > Enterprise parameters). The default setting is Standalone Cluster. Unless you change this entry, all of the clusters will have the same cluster ID. This causes problems in Service Monitor. Changing the cluster ID requires a restart of RIS Collector service, Windows Simple Network Management Protocol (SNMP) service, and the CCMadmin service. Perform these restarts on the publisher and then on the subscribers.

Note: If Service Monitor is already using Cisco Unified Communications Manager and you are changing the cluster name, then you have to delete and re-add the Cisco Unified Communications Managers in Service Monitor for it to reflect the new cluster name.

If Operations Manager is already managing Cisco Unified Communications Manager and you are changing the cluster name, then the cluster names in the service-level view will not reflect the new cluster name. You have to delete and re-add the Cisco Unified Communications Manager in Operations Manager for it reflect the new cluster name.

- Make sure that the Service Monitor server's hostname is resolvable using Domain Name System (DNS). If DNS is not being used, edit the Windows hosts file and enter the Service Monitor hostname and IP address. The hosts file is located at `C:\Windows\system32\drivers\etc`.
 - # This file contains the mappings of IP addresses to hostnames. Each
 - # entry should be kept on an individual line. The IP address should
 - # be placed in the first column followed by the corresponding hostname.
 - # The IP address and the hostname should be separated by at least one
 - # space.
 - #
 - # Additionally, comments (such as these) may be inserted on individual
 - # lines or following the machine name denoted by a “#” symbol.
 - #
 - # For example:
 - #
 - # 102.54.94.97 rhino.acme.com # source server
 - 10.1.1.15 Service Monitor server.cisco.com
 - # add the Service Monitor server IP address and hostname entry into the hosts file if DNS is
 - not being used on the network

Installation Procedures

If you do not have a license key, then during the installation, select the evaluation version. The evaluation version can manage up to 1000 phones for up to 90 days. When the Service Monitor license has been acquired, simply upload the license into the Service Monitor. Operations Manager and Service Monitor require separate licenses.

Licensing and Registration the Software

Licensing grants you permission to manage a certain number of phones. You can enter licenses for Service Monitor during installation or add them later. There is a separate license for Service Monitor and Operation Managers.

Uninstallation

The uninstallation process may cause a warning message similar to the following to appear:

```
The uninstallation is waiting for a process to stop, do you wish to
continue to wait?
```

If you see this message, click Yes and continue to wait.

It is a good practice to delete the `C:\Program Files\CSCOpX` folder and then reboot the server after the Service Monitor application has been uninstalled from any server. Remember to save any Cisco 1040-related call metrics, performance, or node-to-node archived files that you might want to

keep from the C:\Program Files\CSCOPx\data folder.

Failover and Redundancy

Service Monitor supports failover for only sensor functionality. A primary and secondary Service Monitor server can be configured to provide redundancy and failover support to Cisco 1040 Sensor. In case of a primary Service Monitor server going down, the probe would automatically switch over to the secondary Service Monitor server. However there is no synchronization between primary the Service Monitor and secondary Service monitor servers.

Cisco 1040 Failover Mechanism

The 1040 establishes a connection with the Service Monitor and periodically sends an SCCP Keep Alive message. The Service Monitor acknowledges the Keep Alive message to maintain the connection.

The following scenario describes when a 1040 will fail over to the secondary server:

1. The Cisco 1040 stops receiving Keep Alive acknowledgement messages from the primary Service Monitor server.
2. After sending three Keep Alive messages without any acknowledgement, the Cisco 1040 sends a Keep Alive message to the secondary Service Monitor server.
3. The secondary Service Monitor server sends a Keep Alive acknowledgement message.
4. The Cisco 1040 sends a StationRegister message with the station user ID set to the Cisco 1040's ID.
5. Secondary Service Monitor goes to the TFTP server to get the latest configuration file for this Cisco 1040.
6. Secondary Service Monitor sends a StationRegister acknowledgement message.
7. Now Cisco 1040 will start sending syslog messages to the secondary Service Monitor server while still sending Keep Alive messages to the primary Service Monitor server to see whether it's back up again.

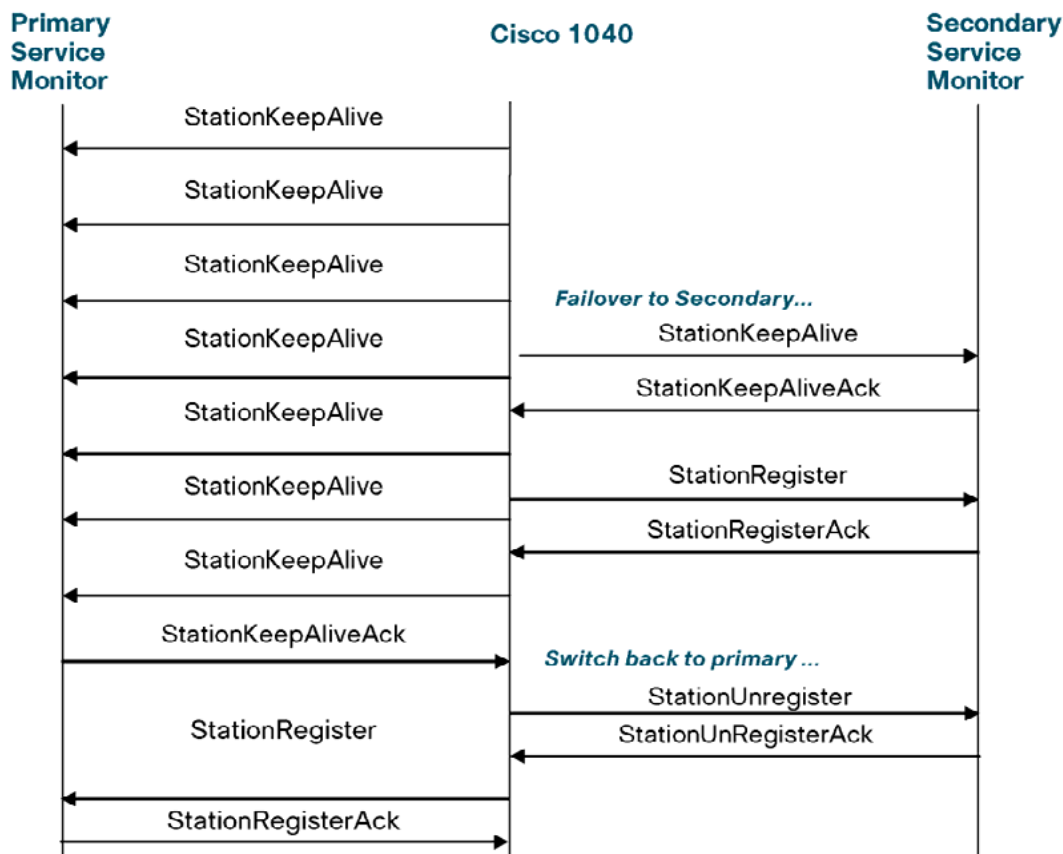
Note: Users cannot set the time of a failover Cisco 1040. The only way to make any configuration changes to a failover Cisco 1040 is to first make this Service Monitor server its primary Service Monitor server.

The following scenario describes how the Cisco 1040 will revert back to the primary server:

1. Cisco 1040 begins to receive a Keep Alive acknowledgement from its primary Service Monitor server once it comes back up again.
2. Cisco 1040 sends a StationUnregister message to the secondary Service Monitor server.
3. Secondary Service Monitor server sends a StationUnregister acknowledgement message to the Cisco 1040.
4. Cisco 1040 sends a StationRegister message with the station user ID set to the Cisco 1040's ID to the primary Service Monitor server.
5. Primary Service Monitor server sends back a StationRegister acknowledgement message.
6. Cisco 1040 starts to send syslog messages to this Service Monitor server now.

See Figure 4.

Figure 4. Cisco Failover Scenario



Cisco 1040 Failover Deployment

Preparing for Failover

The first step is to have two Service Monitor servers available for configuration. One server acts as the primary server and the second as a secondary server. Please refer to the Service Monitor installation guide for the hardware specification of these servers.

It is recommended that these servers connect to the network through redundant paths. This helps ensure that a failure in one part of the network that affects the primary server does not also affect the connectivity of the secondary server.

Setting Up Failover

Failover can be setup globally or for specific sensors. If a Service Monitor server acts as primary or secondary Service Monitor server for any sensors across all locations, failover can be setup globally. A Service Monitor server can act as the primary Service Monitor server for one domain or location and at the same time as the secondary Service Monitor server for another domain or location. In this case, failover needs to be set up for specific sensors.

Setting Up Failover in Default Configuration

Go to the primary Service Monitor server, select the Configuration tab, the Sensors option, and Setup from the TOC. The Setup dialog box is displayed (Figure 5). Enter the IP address or DNS name of the primary Service Monitor server and the IP address or DNS name of the secondary Service Monitor server to the default configuration for failover operations of any sensors.

Figure 5. The Setup Dialog Box

Setup

Call Metrics Archiving: Enable Disable

Data File Directory: C:/PROGRA~1/CSC0px/DataDir

Image File Directory: C:/PROGRA~1/CSC0px/ImageDir

Send traps every 10 minutes per endpoint

Default Configuration to TFTP Server

Image Filename: SvcMonAA2_34.img

Primary Service Monitor: 192.168.137.98

Secondary Service Monitor: 192.168.137.99

OK

Setting Up Failover for a Specific Sensor

Go to the primary Service Monitor server, select the Configuration tab, the Sensors option, and Management from the TOC.

The Cisco 1040 Details dialog box opens showing a list of any previously defined or registered Cisco 1040s. Select Add to create a specific configuration for a Cisco 1040.

The Add a Cisco 1040 dialog opens (Figure 6). Enter the IP address or DNS name of the primary Service Monitor server and the IP address or DNS name of the secondary Service Monitor server to the configuration for failover operations of a specific sensor.

Figure 6. Figure 6 Dialog Box for Adding a Cisco 1040 Sensor

Add a Cisco 1040 Sensor

Sensor Name: Cisco 1040

Image File Name: SvcMonAA2_34.img

MAC Address: 001120FFD012

Primary Service Monitor: 192.168.137.98

Secondary Service Monitor: 192.168.137.99

Description:

	Name	MAC	IP	Primary	Secondary	Registered With	Reset Time
1.	Cisco 1040	001120FFD012	192.168.137.6	192.168.137.98	192.168.137.99	192.168.137.98	30-May-2007 15:40:32 PDT

Add Edit Delete Reset Refresh

Viewing Failover Status

Enter `http://<IP address>` in your browser where IP address is the address of your Cisco 1040. The Current Service Monitor field will show the Service Monitor to which the Cisco 1040 is sending data; this could be a primary or secondary Service Monitor server.

Go to the primary Service Monitor server or secondary Service Monitor server if failover took place, select the Configuration tab, the Sensors option, and Management from the TOC. The sensor details page will show the primary Service monitor server, secondary Service Monitor server, and the Service Monitor server to which the Cisco 1040 is registering (Figure 7).

Figure 7. Figure 7 The Service Monitor Sensor Details Page



Cisco Voice Transmission Quality Redundancy

There is no failover capability for Cisco Voice Transmission Quality capability. However call manager can be added to multiple Service Monitor servers. Up to three Service Monitor servers can be configured as billing application servers in Cisco Unified Communication Manager 5.x or above to receive Cisco Voice Transmission Quality data. So when one Service Monitor server is down, the other Service Monitor server will still be able to obtain the Cisco Voice Transmission Quality data from CCM.

Note: Call manager publisher server is responsible for transferring Cisco Voice Transmission Quality data to the Service Monitor server. If the publisher server is unavailable, there is no mechanism for Cisco Unified Service Manager to obtain the Cisco Voice Transmission Quality data in the cluster.

Backup and Restore

The Service Monitor 2.x database normally is very big. Backup can take hours, so the backup process is not included when the Service Monitor software is reinstalled or upgraded.

On servers on which Operations Manager and Service Monitor are coresident, the database is not very big so it is OK to use the Common Services backup and restore process. The backup UI might time out, but the backup will be complete after some time (this might take a few hours for big database). Restore can be done only after backup has completed successfully.

In Service Monitor standalone servers, it is recommended to use the Common Services backup and restore process only when the database is less than 6 GB. For large Service Monitor databases, it is recommended to do backup manually (saving database password and copy database files) as documented in the user guide.

Configuring Low-Volume Schedule and Database Purging

Service Monitor needs 8 hours of low-volume time during a day. During a low-volume schedule, Service Monitor handles roughly 20 percent of the number records that are processed during a peak period and performs database maintenance. For sensor data, during regular call volume the maximum segment rate allowed is 1500 per minute. Anything over this rate will be discarded. During low call volume, the maximum segment rate allowed is 25 percent of the maximum regular call volume. Throttling of sensor data is based on the total amount of data received in any 5-minute interval; it is not per minute. This allows accommodation of temporary spikes in traffic while blocking continuous high-rate traffic that is over the supported limit. There is no throttling of Cisco Voice Transmission Quality data. Service Monitor standalone server has been tested to support a maximum rate of 500 Cisco Voice Transmission Quality calls per minute.

The default low-volume schedule is 10 p.m. through 6 a.m. To change the schedule, on the Service Monitor server, change the values of these properties in the NMSROOT\qovr\qovrconfig.properties file:

```
lowcallvolume-Mon=0-6,22-24
lowcallvolume-Tue=0-6,22-24
lowcallvolume-Wed=0-6,22-24
lowcallvolume-Thu=0-6,22-24
lowcallvolume-Fri=0-6,22-24
lowcallvolume-Sat=0-20,22-24
lowcallvolume-Sun=0-20,22-24
```

You can configure more than one low-volume period as long as the total time adds up to 8 hours and it covers midnight to 1 a.m. Here are some examples:

```
lowcallvolume-Mon=1-7,21-23
lowcallvolume-Tue=0-6,21-22,23-24
```

To put changes into effect after you edit qovrconfig.properties, you must stop and start the QOVR process. While logged on to the server where Service Monitor is installed, from the command line, enter these commands:

```
pdterm QOVR
pdexec QOVR
```

Service Monitor needs 4 hours data purge time. Data purging must occur during the low-volume schedule and must not run from midnight to 2 a.m. The default schedule is 2 a.m. to 6 a.m. To change the schedule on the Service Monitor server, change the values of these properties in the NMSROOT\qovr\qovrconfig.properties file:

```
datapurge-Mon=2-6;
datapurge-Tue=2-6;
datapurge-Wed=2-6;
datapurge-Thu=2-6;
datapurge-Fri=2-6;
datapurge-Sat=2-6;
datapurge-Sun=2-6;
```

Data purge need not run continuously for 4 hours. You can configure more than one data purge period as long as:

- The total time adds up to 4 hours
- Data purging occurs during low-volume schedule
- No data purging occurs from midnight through 2 a.m.

Here are some examples:

```
datapurge-Mon=2-5;22-23;  
datapurge-Tue=2-3;4-6;23-24
```

Note: Do not edit the properties files using Wordpad as it introduces a carriage return. Use Notepad to edit the file instead.

The data retention period determines the number of days that data is retained in the Service Monitor database before being purged. The default value depends on the deployment scenario:

- Service Monitor alone on a server: 7 days
- Any coresident server: 3 days.

Troubleshooting

This section provides a few troubleshooting tips.

Q. I don't see any service quality alerts in the Operations Manager dashboard. What could be the problem?

A. Before debugging this problem, do the following:

- Go to the Service Monitor application and verify that the probes are registered and visible in the Service Monitor GUI.
- Go to the Service Monitor setup page and verify that Operations Manager is entered as a trap recipient, even if Service Monitor is on the same machine.
- In Operations Manager, go to Administration > Service Quality and add Service Monitor, even if it is on the same machine.
- If all of the previous are correct, then only calls that fall below the Service Monitor threshold will be shown in the Alerts and Activities page.
 - When all of the previous have been done, debug the problem by doing the following:
- Check syslog.log under:
 - NMSROOT\CSCOPx\syslog.log in IPCOM 1.0
 - NMSROOT\CSCOPx\log\qovr\syslog.log in CUOM 1.X/ CUSM 2.X

Check if you see recent syslogs. Check the D=<value>. This is the MOS score multiplied by a factor of 10. Check to see whether this value is below the threshold multiplied by 10.

- If this is true, check NMSROOT\log\qovr\trapgen.log. This file should contain the traps that are being generated by the system. If traps are available in this file, then it means that the Service Monitor portion is functional and ready. If not, then check for exceptions in probemanager.log and datahandler.log.

Q. A few rows in the sensor diagnostic reports do not show directory numbers. What could be the issue?

A. Make sure that the call is over. CDRs and CMRs are generated at the end of the call only, whereas sensors send data even while the call is in progress.

- Make sure that the corresponding call in the Cisco Voice Transmission Quality diagnostic report shows directory number for both the caller and callee. Some gateway ports do not have directory numbers. Sensor rows for those calls will not have directory number information.

Q. Why is the directory number missing in sensor reports?

A. Please check the following:

- The Cisco 1040 Sensors monitor RTP traffic; they do not report directory numbers.
- Directory numbers can be seen in the diagnostic sensor reports only if the call whose streams are being reported by sensors to Service Monitor is also reported by call manager to Service Monitor meaning that the call manager server must be added to the Service Monitor server.
- Service Monitor and call managers have to be time synched in order to display the directory numbers.

Q. General sensor troubleshooting tips:

A. Please check the following:

- If the sensor is not receiving the IP address, check the DHCP configuration on the DHCP server.
- Start with `http://<ip-addr>/Communication` and see what the sensor communication debug page says. If there are any startup issues on the sensor (such as the TFTP server not being reachable or inability to download the image file), this page should tell you.
- If you are suspecting a sensor/Service Monitor issue, then use the sniffer to check the communication from the sensor to the rest of the world. You should use the sniffer on the management port of the sensor to begin with.
- If Cisco Unified Communication Manager 5.x is used for TFTP server, please restart the TFTP service after the configuration files and image files are manually copied over for changes to take effect.
- The spanning port doesn't play a role in booting the sensor. But do make sure SPAN is configured on the switch port instead of the routed port.
- Service Monitor installs the TFTP server by default; If the customer installs another TFTP server, make sure the TFTP service in Common Service (CS) is shut down to avoid any potential issue.
- Do not edit configuration files using Wordpad (where the TFTP server is not editable), as it introduces a carriage return.

Q. What type of credentials do we configure?

A. Please check the following:

- For call manager 5.x and above: HTTP/HTTPS credentials

These are the same as those used by the browser for accessing the call manager Web-based administration page. The IP address provided should be the publisher's IP address.

- For call manager 4.x : Both HTTP/HTTPS and CDR database credentials

CDR database: Is the Microsoft SQL Server CDR database credentials on 4.x versions. HTTP is for device name correlation.

- For call manager 3.3 : Both the device database and CDR database credentials

The device DB is the Microsoft SQL Server CCM030X database source credentials on 3.3 versions. This is for device name correlation

- HTTPS gets the device type for an endpoint involved in a call from the call manager.. HTTPS is also used to discover the IP address of the CDR source in 5.x systems. This is shown as CDRM in the UI. The CDR database gets the CDRs and CMRs from 3.x and 4.x systems. The device database gets the device type for an endpoint from a call manager 3.x.
- For all systems, failure to contact the call manager is immediately indicated by 'Failure'. Verify network connectivity and that credentials have not changed on call manager
- If the call manager server whose credentials are being entered does not have a DNS entry (and vice versa):
 - You need also to provide the hostname, which is the same as that displayed on the System->Server page in the call manager Web administration site.
 - Otherwise the hostname field is optional.

Detailed troubleshooting tips are documented at

http://cisco.com/en/US/products/ps6536/prod_troubleshooting_guides_list.html.

Useful URLs

Cisco.com URLs for Customers and Partners

Product Information

Product Page

http://www.cisco.com/en/US/products/ps6536/prod_literature.html

Release Notes

http://www.cisco.com/en/US/products/ps6536/prod_release_note09186a00807ee746.html

Data Sheet

http://www.cisco.com/en/US/products/ps6536/products_data_sheet0900aecd80580026.html

User Guide

http://www.cisco.com/en/US/products/ps6536/products_user_guide_book09186a00807ee7fb.html

Training

Instructor-led Cisco Unified Operations Manager and Service Monitor (CUOM-SM) 2-day training:

Customer and partners can send an email to aeskt_registration@cisco.com.

Evaluation Downloads

For Partners and Customers

Step 1. Go to the Marketplace site at <http://www.cisco.com/go/marketplace>. Note that you must log in with a Cisco employee (CEC) or authorized Cisco Partner login and password:

Step 2. Select the Collateral & Subscription Store link.

-
- Step 3. Read the notice to Cisco employees and click Continue.
- Step 4. From the navigation menu at the top-left corner of the page (above the Subscriptions link), select the Marketing Collateral link. From the Marketing Collateral navigation menu, select Network Management Evaluation Kits, and then select the desired evaluation kit.
- Step 5. Use Add to cart and Checkout to place the order for the desired kit, using your Access Visa or personal credit card.

For further questions on Cisco Unified Communications Operations Manager or Cisco Unified Communications Service Monitor, or for any other Cisco Unified Management-related questions, send an e-mail to ask-ipc-management@cisco.com.

Patch Download

<http://www.cisco.com/cgi-bin/tablebuild.pl/servmon>

Miercom Review of Operations Manager and Service Monitor

http://www.cisco.com/en/US/products/ps6535/prod_white_papers_list.html

**Americas 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: 408 527-0883

Asia Pacific Headquarters

Cisco Systems (USA) Pte. Ltd.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCVP, the Cisco logo, and Welcome to the Human Network are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PIX, ProConnect, ScriptShare, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0711R)