



# Monitoring VoIP with Cisco Network Analysis Module

White Paper

March 2009

## Contents

<a href="#">Introduction</a> .....	3
<a href="#">Target Audience</a> .....	3
<a href="#">Overview</a> .....	3
<a href="#">Deploying NAM for VoIP Monitoring</a> .....	4
<a href="#">How Voice Monitoring Works</a> .....	4
<a href="#">Using the NAM GUI for Voice Monitoring</a> .....	5
<a href="#">Use Case: Troubleshooting Voice Quality Degradation</a> .....	6
<a href="#">Use Case: Linking Voice Quality to Quality of Service (QoS) Issues</a> .....	8
<a href="#">Use Case: Voice Quality Thresholds for Proactive Troubleshooting</a> .....	10
<a href="#">Cisco NAM Integration with Cisco Unified Communications Management Suite</a> .....	12
<a href="#">Summary</a> .....	12
<a href="#">Glossary</a> .....	13
<a href="#">References</a> .....	13

## Introduction

The convergence of voice, video and data networks has resulted in increased network management complexity, and demands a new set of tools, techniques and best practices. Real-time applications such as voice lay more stringent requirements on the network than the typical applications in a business enterprise. For example, VoIP is very sensitive to the latency and jitter properties of the network; users quickly perceive and raise trouble tickets when such problems are experienced. On the other hand, web server transactions and email traffic have a much higher tolerance for latency and jitter. Effectively managing the additional complexity of voice over IP (VoIP) deployments offers significant benefits to IT personnel by reducing trouble tickets and maintaining a high level of user satisfaction.

A comprehensive management strategy requires alignment with organizational processes as well as consistency in managing the entire application delivery lifecycle. Examples of the latter include detailed plans for deployment, troubleshooting, network maintenance and upgrades. An important part of implementing the process in each part of the life-cycle is the use of the management tools that aid in planning, managing and troubleshooting enterprise networks. In this paper, you will be introduced to the capabilities within the Cisco® Network Analysis Module (NAM) that will proactively detect VoIP quality issues and help troubleshoot them. VoIP analysis capabilities in NAM 4.0 will be described along with use cases that detail the use of these tools to solve day to day operational problems. With the NAM 4.0 release, Cisco Unified Service Monitor, a component of Cisco Unified Communication Management Suite<sup>1</sup>, is able to aggregate voice performance metrics from NAMs deployed across the network for Enterprise-wide voice quality reporting. This suite which includes Cisco Unified Service Monitor, Cisco Unified Operations Manager and Cisco Unified Service Statistics Manager provides enterprise wide reporting and management for VoIP deployments.

## Target Audience

This paper should be helpful to anyone whose job role involves monitoring, managing, planning and responding to network performance and quality issues in VoIP networks. Examples of job roles include network architects, network engineers, IT architects/engineers, Performance Management and Optimization personnel and IT managers.

## Overview

The goal of a voice quality monitoring system is to support a set of capabilities that provide accurate and timely indications of voice quality and notify the user when needed. Mean Opinion Scores (MOS) offer a standard method to monitor the quality of a VoIP call. This method takes into account latency, jitter and other network properties that affect voice quality. In addition to computing call quality, the system must be cognizant of various other pieces of information, such as the identity of phones making the calls, call history, call volume and relevant historical trends such as peak and low periods. Finally, the ability to notify the IT administrator of a problem that has occurred is very important. Cisco NAM provides these capabilities: NAM detects and computes MOS for VoIP calls transported through Real Time Protocol (RTP) streams. MOS is computed periodically and reported via NAM's GUI which provides real-time and historical reporting on voice quality. Administrators can get access to MOS scores in real-time, even as calls are in active state. It is important to note that NAM is an open device, which means that its data is

---

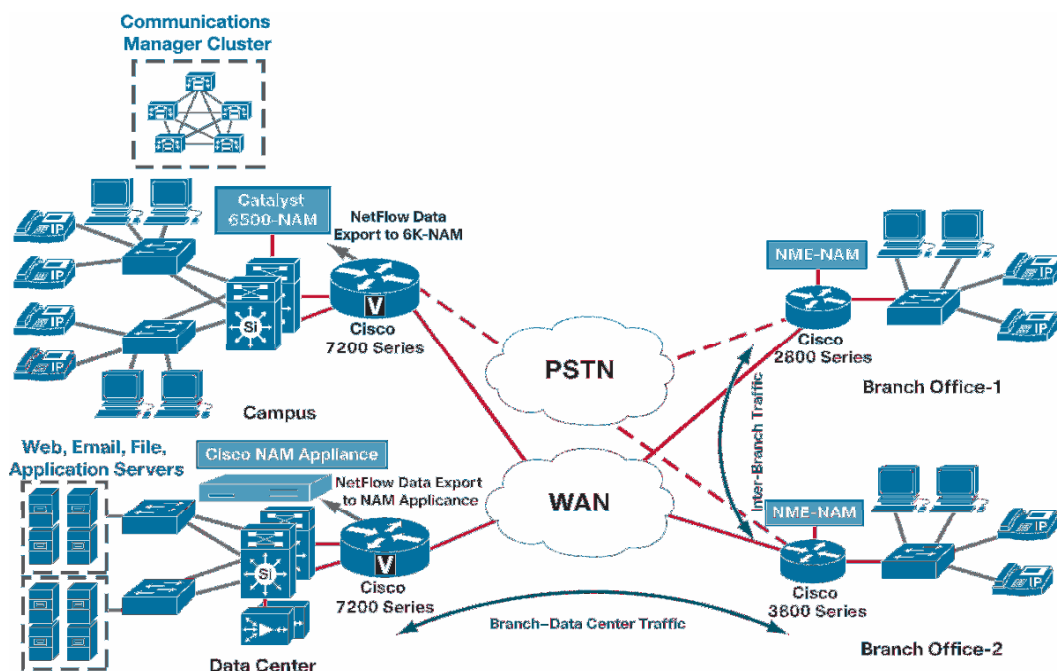
<sup>1</sup> Cisco Unified Communications Management release supporting the integration is expected to be available in the second half of 2009.

available for consumption by any application that can poll the NAM via its interfaces. This openness allows for easy integration into other applications being used.

## Deploying NAM for VoIP Monitoring

Consider the following diagram which depicts an enterprise network, comprising a campus, a data center and two branch offices.

**Figure 1.** Cisco NAM Deployment



NAM can be deployed in various locations depending on the use case being addressed. For example, when deployed in a branch, NAM can monitor the quality of all calls entering and exiting that branch office and monitor any dips in quality levels in that location. Alternatively, a NAM in the data center can monitor signaling messages sent from a phone to the Cisco Communications Manager cluster, and collect detailed information about the calling and called parties. A NAM located at the edge of the main campus can raise alarms about poor call quality from the main campus to a particular remote office location.

## How Voice Monitoring Works

VoIP phone calls are setup using a signaling protocol, such as SIP, using which the phone endpoints exchange information. Once setup is complete, voice traffic is sent through a streaming protocol such as RTP. NAM detects both signaling messages and voice traffic over RTP streams. It also has the ability to link a set of call setup messages to the RTP streams associated with them.

NAM monitors signaling messages for Skinny Call Control Protocol (SCCP), SIP, H.323, and Media Gateway Control Protocol (MGCP). The information collected from signaling messages includes currently active calls, call statistics, call history, detailed information about the calling and called parties, codecs used, port numbers used and other relevant information. These statistics can be collected both on a real-time and historical basis.

In order to measure the quality of the call, NAM detects and monitors RTP streams. First, NAM examines the packet header and identifies whether it is an RTP packet. If so, it checks whether the packet belongs to a new or existing RTP stream. Once the RTP packet is detected and associated with a stream, it is sent to the MOS process for quality analysis. The MOS process performs real time computations to measure voice quality metrics such as jitter, actual packet loss, adjusted packet loss, seconds of concealment, and severe seconds of concealment. Using the aforementioned metrics, the NAM computes the R-Factor MOS based on the ITU-T recommendation G.107. The best, worst and average values for these metrics are reported every minute through GUI.

An important aspect of this real-time reporting is that voice quality metrics are available to users even when the call is active. There is no need to wait for the end of the call before such statistics are collected. This real-time visibility is a critical part of the solution, and facilitates rapid responses to problems.


**Note:** To monitor signaling messages, the NAM has to be in the path of call signaling messages from the VoIP end point (IP Phone) to the call management server. In a Cisco Unified Communications System, the server would be a Unified Communications Manager.

### Using the NAM GUI for Voice Monitoring

The GUI onboard NAM allows for easy access to real-time and historical VoIP data collected from the network. It also provides the ability to set and monitor thresholds proactively, so that there is minimal impact on end users of the system. This section will explore various aspects of the voice related screens in the GUI.

The voice monitoring GUI is divided into different categories. The Active Call Monitoring section provides visibility on various metrics for currently active call. The MOS Quality and Alarm Threshold charts provide a 1 hour window of the quality of active calls. Note that since quality metrics are computed for RTP streams, the NAM must be placed in the path of the call. For example, a NAM placed at the edge of a branch office can provide quality metrics for all calls entering or leaving that branch.

The Active Calls Table provides information gathered by associating call signaling messages with the RTP streams used for the call. To provide correlation between signaling and RTP streams, NAM must ideally be placed at locations where it has visibility into a call's signaling messages and RTP streams.

**Figure 2.** Active Calls - MOS Quality

The Terminated Calls Table provides analysis of aspects such as which calls suffered from quality issues, where those calls originated, what codecs were used and other relevant details. For example, the Worst N Calls menu provides detailed visibility into the Worst N Calls. This will show the caller and called party for the calls and the start and end times along with the quality. This helps isolate the problem to specific network locations or to transient network conditions during specific times in the day. The known phones and RTP stream sections provide details from the perspective of an individual phone and from the perspective of raw RTP streams, respectively. During troubleshooting, it is useful to navigate from the Active Calls menu to the RTP stream screen, for example to get additional details on the codecs used in the calls.

The use cases that follow will utilize the GUI screens described in this section to solve real-world problems. Workflows will be provided along with the relevant screenshots, to show how problems commonly experienced in VoIP networks can be tackled.

### Use Case: Troubleshooting Voice Quality Degradation

Consider a situation in which the network administrator learns about problems with VoIP quality by monitoring the NAM GUI. What steps could the administrator follow to isolate the problem's root cause?

As illustrated in the Figure 2, NAM classifies the voice calls by quality into poor, fair, good and excellent categories. This rating is based on MOS scores and can be configured by the user to suit their network's sensitivity levels. NAM uses preset default values for the MOS ranges. The chart indicates that there were a few calls with poor quality a few minutes ago.

**Figure 3.** Individual RTP Streams and their Associated MOS

**RTP Stream Traffic**

Current Data: as of Thu 09 Oct 2008, 09:31:21 UTC

Auto Refresh

Source Address

#	Source Addr : Port	Dest Addr : Port	Payload Type	SSRC	Pkt Loss /million	MOS	Adj Pkt Loss (%)	Jitter (ms)	SSC
1.	10.14.1.2 : 1280	10.14.1.20 : 1250	G711Ulaw_64k	34933	40.00	1.76	40.00	0.05	60.0
2.	10.14.1.2 : 1296	10.14.1.20 : 24614	G711Ulaw_64k	27379	40.00	1.76	40.00	0.06	60.0
3.	10.14.1.2 : 1494	10.14.1.20 : 6374	G711Ulaw_64k	54306	40.00	1.76	40.00	0.06	60.0
4.	10.14.1.2 : 1730	10.14.1.20 : 54846	G711Ulaw_64k	1750	40.00	1.76	40.00	0.06	60.0

In response to this problem, the next step in troubleshooting is to navigate to obtain more detail about poor calls. Figure 3 shows the individual RTP streams and the MOS scores associated with them. As indicated in the highlighted portion of the table below, the MOS score of the first several calls is very low (1.76) as per the ranges defined in the foregoing chart.

There are other interesting clues that can be gleaned from Figure 3. Note that the Packet Loss column indicates that VoIP streams are experiencing packet loss.

The next step is to get clues as to where packets are being dropped. The source address of the RTP stream should be examined. All calls have the source IP address 10.14.1.2 but with different port numbers. This is typical of a conferencing system that uses different port numbers for different streams. By looking up the network topology diagram, we learn that 10.14.1.2 is located in Building 3 of the main campus of the company. The topology also indicates that there is a NAM at the edge router for Building 3. We login to that NAM looking for clues on where packets might be getting dropped.

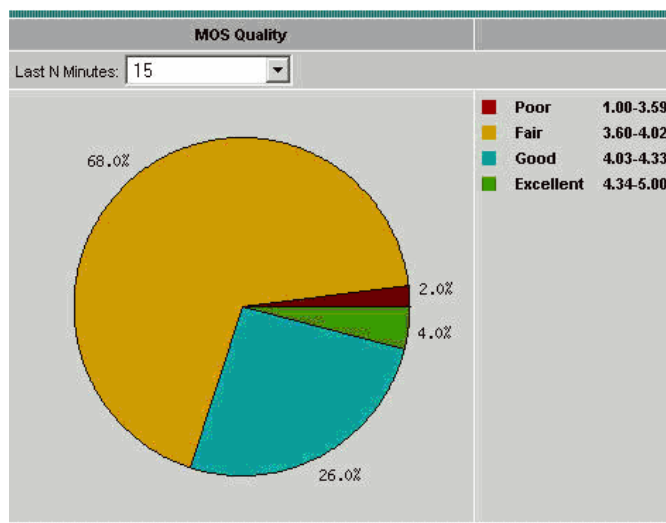
**Figure 4.** Interface Statistics from the NAM in Building 3

10.	VLAN-11	3 G	467 M	3 G	3%	2 G	0	0	0	0	0	0
11.	Vl11	3 G	467 M	3 G	3%	1 G	0	0	50	0	0	0
12.	VLAN-35	290 M	308 M	2 G	3%	1 G	0	0	0	0	0	0
13.	Gi1/22	4 G	4 G	2 G	3%	1 G	0	0	881 K	0	0	0
14.	Vl10	467 M	3 G	2 G	3%	2 G	0	0	0	0	0	0
15.	VLAN-10	467 M	3 G	2 G	3%	3 G	0	0	0	0	0	0
16.	Vl16	4 G	4 G	2 G	3%	2 G	0	0	0	0	0	0
17.	VLAN-16	4 G	4 G	2 G	3%	2 G	0	0	0	0	0	0
18.	Gi1/16	4 G	4 G	2 G	3%	2 G	0	0	935 K	0	1	0
19.	Gi1/34	143 M	41 M	2 G	3%	1 G	0	0	0	0	0	0

881K InDiscards found on Gi1/22 which connects to the enterprise core

By navigating to the Interface statistics screen that provides details about packet related statistics, we find that Gi1/22, the interface that connects to the core of the campus network, is experiencing serious packet loss.

**Figure 5.** MOS Quality Chart



As this interface serves all traffic going from and to Building 3 and the rest of the campus including voice traffic, this is most likely the root cause for packet drops on the RTP stream. The problem in this case was found to be a hardware defect on the line card that affected the interface. Replacing the card fixed the issue.


This troubleshooting workflow highlights some of the VoIP quality monitoring capabilities, and also shows how VoIP features can be used in combination with other traffic monitoring features on the NAM. In this particular case, we used interface statistics monitoring in the NAM in Building 3 to isolate the root cause of the problem.

**Use Case: Linking Voice Quality to Quality of Service (QoS) Issues**

The NAM GUI offers multiple reports, and depending on the use case, users have the option of using one report or the other to start their workflow. This use case approaches a problem that is a slight variation of the problem highlighted in the previous use case. While the previous use case describes a proactive response, this one shows how to respond rapidly to a problem.

The situation is that the administrator has received a series of complaints in the last few minutes. He realizes that there is a quality problem but is unsure where in the network it is. In this case, it is useful to narrow down the problem to a set of calls. Consider the MOS quality pie chart in Figure 5. This chart allows you to vary the time period being monitored so that you can narrow your analysis into a 5-minute or a 15-minute period. Figure 5 shows that the overall call quality is less than desirable. There are very few calls that were excellent. This points to a wide spread problem related to voice traffic.

**Figure 6.** Worst N Calls



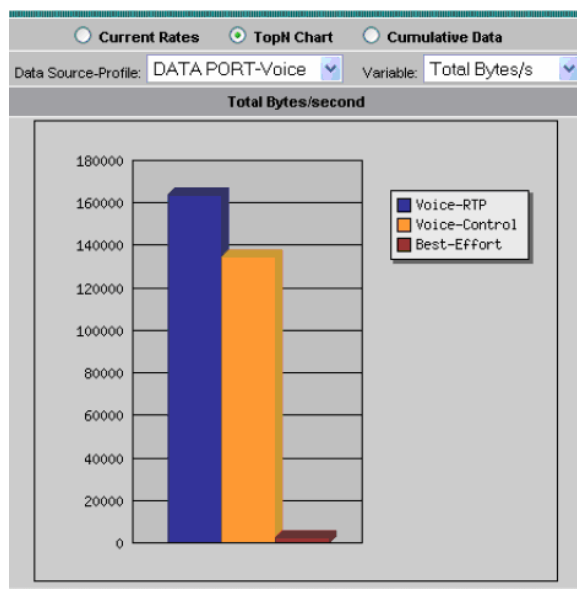
Further evidence of the issue is provided by the Worst N Calls display (Figure 6). A MOS of 2.36 is in the “Poor” range.

**Figure 7.** Additional Call Details

	Calling Party	Called Party
Phone Number:	sip:id00248@10.16.10.249	sip:id50248@10.15.10.249
Alias:	id00248	id50248
Signaling Protocol:	SIP	
Receiving Signaling Address/Port:	10.16.10.249	10.15.10.249
Call Identifier:	705dafc5@10.16.10.249	
Call Tag:	2B5391B	3961
<b>General Information</b>		
Start Time:	01-22-09 20:21:00 UTC	
End Time:	01-22-09 20:26:00 UTC	
<b>RTP Stream Information</b>		
Receiving Audio Address/Port:	10.16.10.249 / 11656	10.15.10.249 / 11656
Proposed Audio Payload Types:	PCMA	
Receiving Video Address/Port:	- / -	- / -
Proposed Video Payload Types:	-	

We navigate to get additional details on the call to track down the phone location, codec used, the related RTP stream, and any other useful information (Figure 7).

**Figure 8.** Traffic Distribution After QoS is Implemented



Since the problem appears to be wide spread and we are unable to isolate it to a specific location or time duration, we use the quality of service (QoS) monitoring feature of the NAM to check if voice traffic is getting the service level in the network. First, Differentiated Services (Diffserv) profiles are created to identify which applications are being associated with DSCP or Type of Service (ToS) values. The NAM allows the administrator to observe application traffic flow, the DSCP values associated with each application and total bandwidth utilization per DSCP value.

This analysis showed that voice traffic was being treated as best effort traffic, i.e. priority was not being given to voice streams. Even though the network was provisioned with large bandwidth connections, lack of QoS led to issues with VoIP traffic during peak hours. Applying an appropriate QoS scheme eliminated this problem.

**Figure 9.** Traffic Distribution After QoS is Implemented

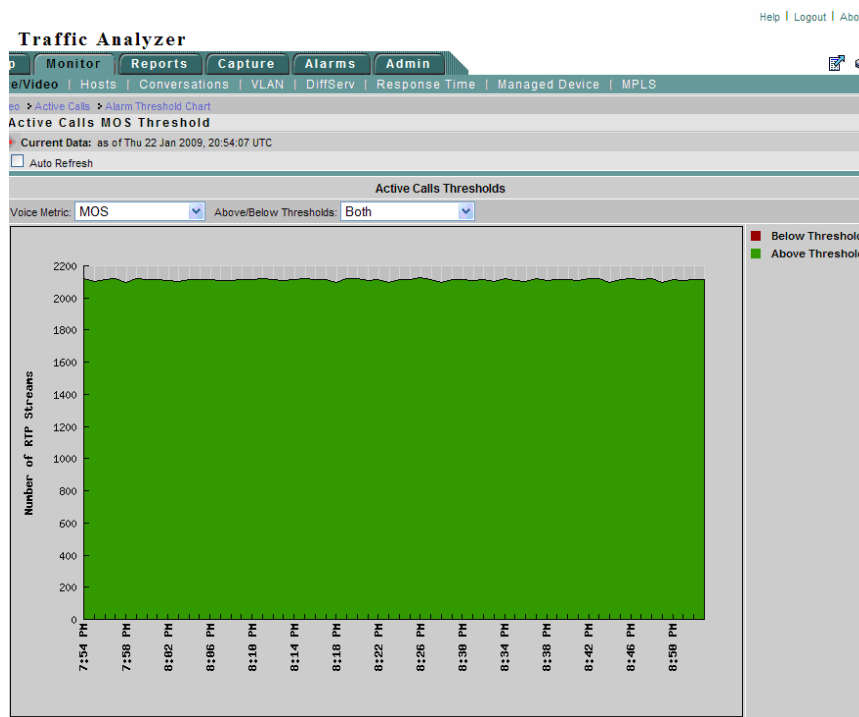
Voice/Video Stream Thresholds		
	Codec	Thresholds
<input checked="" type="checkbox"/> MOS:	G711 (1-4.4):	4.1
	G722 64k (1-4.5):	3.3
	G722 56k (1-3.3):	3.3
	G722 48k (1-4.5):	3.3
	G723.1 (1-3.9):	3.4
	G728 (1-4.0):	3.6
	G729 (1-4.11):	3.9
	GSM (1-4.3):	3.5
	G726 32k (1-4.2):	3.6
	G726 24k (1-4.2):	3.3
G726 16k (1-4.2):	3.1	
<input checked="" type="checkbox"/> Adjusted Pkt Loss (1-100%):		3.0
<input checked="" type="checkbox"/> Actual Pkt Loss (1-100%):		3.0
<input checked="" type="checkbox"/> Jitter (1-10000 ms):		5
<input checked="" type="checkbox"/> Seconds of Severe Concealment (1-60):		2
<input checked="" type="checkbox"/> Seconds of Concealment (1-60):		6

### Use Case: Voice Quality Thresholds for Proactive Troubleshooting

To manage voice quality proactively the system should automatically alert you to potential problems. This is achieved on the NAM through the use of thresholds and alerts. As shown in Figure 9, a number of metrics can be configured for alerts. MOS thresholds can be configured with different values for different codecs because the quality tolerances could be different depending on the codec. In addition, jitter, seconds of concealment, and packet loss can be configured with thresholds based on the requirements of the network. Preset default values are provided and can be modified as needed.

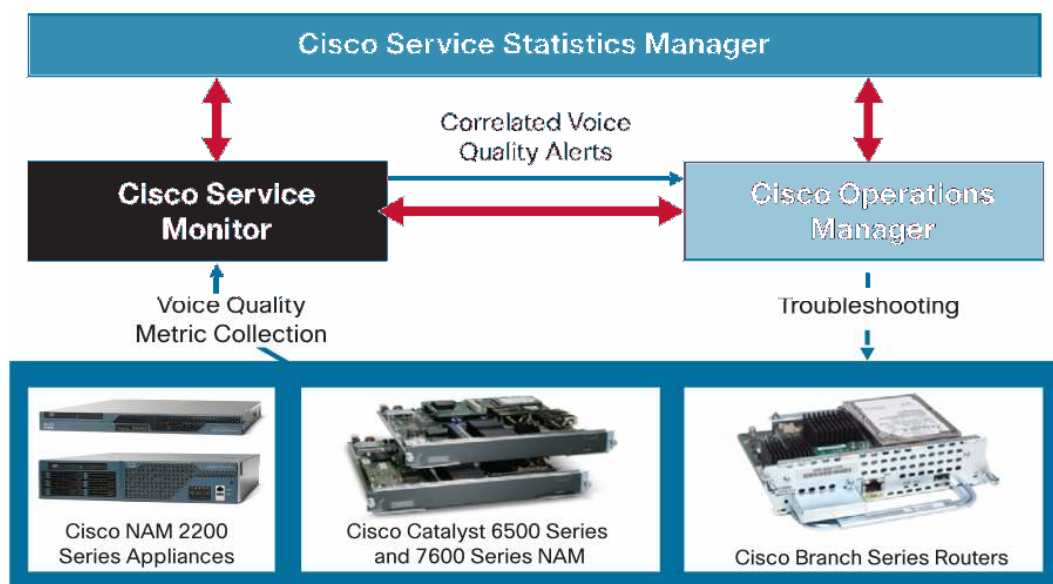
Alerts are sent as syslog messages, therefore, the NAM should be configured to export alerts to a specific syslog receiver. In a typical enterprise, a syslog receiver collects syslog messages from various devices, and feeds them into a network or performance management product that processes the message. The figure below shows the various thresholds that can be configured. This capability would be useful when different MOS tolerances are required for each codec.

**Figure 10.** MOS Threshold Chart



NAM provides a snapshot of current call quality vis-à-vis the thresholds that are configured (Figure 10). This report allows administrators to quickly determine whether or not any calls have exceeded the set thresholds. In the picture below, all calls are of excellent quality indicating that VoIP is operating without any problems.

**Figure 11.** Cisco NAM Metrics are Rolled Up into the Cisco Unified Communications Management Suite



### **Cisco NAM Integration with Cisco Unified Communications Management Suite**

The Cisco Unified Communications Management suite of products provides end to end visibility into all aspects of VoIP networks. To analyze voice quality, the Cisco Unified Communications management suite rolls up voice metrics from the NAM for network wide quality analysis. Such an integrated system provides a comprehensive view of voice quality across the network. Also, it allows troubleshooting of performance issues experienced in one part of the network but whose root cause may be located elsewhere. As noted earlier, the Cisco Unified Communications Management suite release supporting the integration is expected to be available in the second half of 2009 (Figure 11).

The products within this suite, that will collect and use information rolled up from the NAM, directly or indirectly, are defined briefly:

- Cisco Unified Operations Manager monitors & diagnoses problems as well as tests and tracks changes and inventory, providing network visibility and real-time operations management capabilities
- Cisco Unified Service Monitor tracks & reports on the user experience, providing automated diagnostics for high service quality assurance
- Cisco Unified Service Statistics Manager provides robust executive & operational reports as well as capacity planning reports

Figure 11 shows how NAM data is rolled up by Service Monitor. These metrics are processed and any resulting network alarms and alerts are exported to Cisco Operations Manager. Cisco Operations Manager notifies the administrator if an alarm is raised, and provides a direct link to the NAM module responsible for triggering the alarm. The administrator can login directly into the NAM to continue troubleshooting. This workflow between NAM, Cisco Service Monitor and Cisco Operations Manager allows administrators to manage quality across the network through early notification and useful information about the problem's root cause. Finally, Cisco Unified Service Statistics Manager collects NAM data indirectly through Cisco Service Monitor and uses it to provide historical trending and capacity planning reports that help IT personnel plan for future deployments.

### **Summary**

In today's complex multimedia networks, the ability to measure the quality experienced by end users is critical. By providing visibility into the quality of VoIP traffic, NAM 4.0 significantly enhances the ability of IT personnel to detect, isolate and troubleshoot VoIP problems. Its real-time monitoring and alerting capabilities facilitate a proactive approach to monitoring user experience, which helps increase end user satisfaction.

## Glossary

<b>MOS</b>	The computed Mean Opinion Score for the minute interval. MOS is computed according to ITU-T G.107 E-Model every three seconds. The reported MOS score is the average of all three second scores for the minute. The minimum stream duration to compute MOS is one second of media flow.
<b>Packets Lost</b>	The count of aggregate packets lost due to network transmission during the reporting period. This is computed based on observed RTP sequence number analysis.
<b>Jitter</b>	The RFC3550 jitter value in milliseconds. This value is a smoothed metric and may not be adequately indicative of problems given short and sudden spikes in jitter. This value should be a good description of the jitter given a uniform and constant distribution of jitter events.
<b>Percent network loss</b>	The percentage of packets dropped by the network on the way to the destination address.
<b>Adjusted packet loss</b>	The percentage of packets lost due to high jitter. This value is computed based on a reference jitter buffer with a fixed length play-out delay. It is not affected by network loss.
<b>SOC</b>	Seconds Of Concealment. The number of seconds during which any impairment was experienced. Impairment can be due to network loss or high jitter. If just one packet is lost during the entire reporting interval, this value should be 1. If each second of the reporting period experiences at least one lost packet then this value should be 60.
<b>SSOC</b>	Severe Seconds Of Concealment. The number of seconds during which severe impairment was experienced. Severe impairment is defined by packet loss greater than or equal to 5 percent, including both network loss and loss due to jitter buffer discards.
<b>Codec</b>	The codec used by the media stream. This value is derived from the RTP payload type and may also include information from media stream payload lengths and packetization properties.

## References

<b>Cisco Network Analysis Module</b>	<a href="http://www.cisco.com/go/nam">http://www.cisco.com/go/nam</a>
<b>Cisco Unified Service Monitor</b>	<a href="http://www.cisco.com/go/cusm">http://www.cisco.com/go/cusm</a>
<b>Cisco Unified Operations Manager</b>	<a href="http://www.cisco.com/go/cuom">http://www.cisco.com/go/cuom</a>



**Americas Headquarters**  
Cisco Systems, Inc.  
San Jose, CA

**Asia Pacific Headquarters**  
Cisco Systems (USA) Pte. Ltd.  
Singapore

**Europe Headquarters**  
Cisco Systems International BV  
Amsterdam, The Netherlands

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](http://www.cisco.com/go/offices).

CCDE, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo 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. (0812R)