CHAPTER 8

# Network Management

The Cisco WAN and application optimization solution provides a powerful set of WAN optimization features. To translate these features into verifiable user benefits, effective network management tools are required.

This chapter describes important network management features of the WAN and application optimization solution.

Instead of a general discussion on network management covering a wide range of management functions (for example, network planning and security), the focus here is specifically on traffic and performance monitoring, with a brief overview of configuration management. These functions are immediately relevant to the WAN and application optimization solution deployment and operation.

## 8.1 Centralized Monitoring, Reporting, and Troubleshooting

Adequate monitoring, reporting, and troubleshooting are essential elements of successful WAN optimization initiatives. The value of these features transcends functional teams, providing technology and business managers with the information needed to enhance productivity, enabling IT and network architecture groups to improve infrastructure design; and reduce time spent solving problems.

For the WAN and application optimization solution, the core monitoring features help the user to:

1. **Profile** traffic patterns and resource bottleneck loads, enabling the user to prioritize links and protocols to be optimized
2. **Baseline** performance metrics for applications (for example, transaction times) and for resources (for example, link and server CPU/memory utilization) before WAN optimization
3. **Assess** the effectiveness of each successive WAN optimization initiative
4. **Troubleshoot** problems that occur during and after the successive WAN optimization initiatives

### 8.1.1 Monitoring Challenges and Solutions

WAN optimization devices present new challenges for developers of network monitoring products. Any WAN optimization device has at least some potential to disrupt network traffic monitoring to varying degrees.

For example, optimization device pairs can employ Layer 3 (L3) tunneling architectures and LAN-WAN segmentation. Most WAN optimization devices do not preserve original IP addresses and TCP port information for optimized user-to-server connections. Such changes can cause many conventional monitoring techniques to generate obscure or unintelligible traffic profile data.

Optimization devices can disrupt performance measurements. Leading WAN optimization devices rely on TCP proxy to manage connections locally at each appliance, as shown in Figure 8-1. Each WAN optimization device acknowledges received TCP transmissions before actually sending data across the network. This behavior can skew response time analysis.

For example, application response time monitors may report the near-immediate acknowledgment of data that has not yet been transmitted, received, and acknowledged over the WAN, causing incorrect reporting of response times. Numerous network monitoring techniques can fail to accurately report performance metrics in this environment.

*Figure 8-1 TCP Proxy Architecture Used in Typical WAN Optimization Devices*

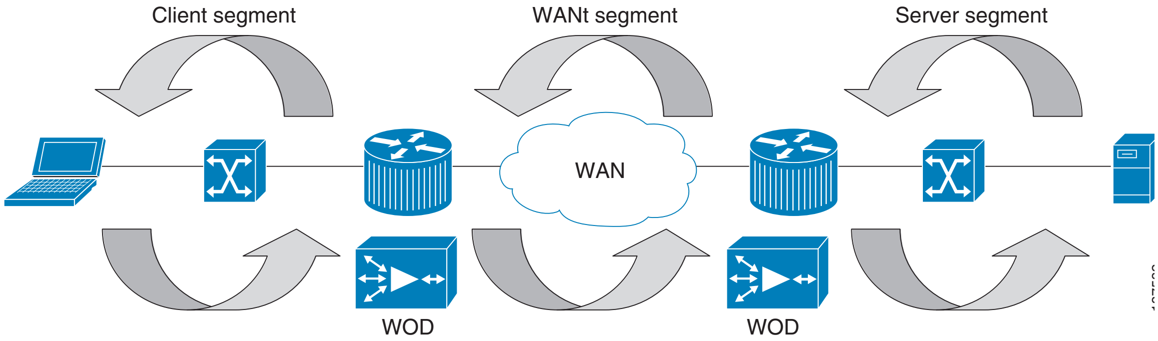

Fortunately, Cisco WAN optimization technologies are architected to help overcome the transparency issues found in competing products. For example, Cisco Wide Area Application Services (WAAS):

- Preserves client and server IP addresses and TCP port numbers
- Exposes performance metrics for optimized traffic using FlowAgent technology on the WAAS Wide-Area Application Engine (WAE) appliance and NME network modules
- Does not mask detailed NetFlow traffic records.

When deployed with NetQoS and Cisco monitoring products specifically designed for use in Cisco WAN optimized environments, users benefit from end-to-end and tier-to-tier visibility of application performance, more accurate reporting of traffic traversing the network, and detailed device performance metrics.

## 8.2 NetQoS Performance Center: Network-Wide Monitoring and Reporting

Today's best-of-breed network performance management products draw from multiple data sources to expedite problem resolution and improve infrastructure planning. The NetQoS Performance Center suite of management modules presents a rich set of analytical data using a single interface. NetQoS Performance Center derives this data from passive monitors and Cisco monitoring instrumentation that is already present in network and data center devices. See Chapter 4, “Cisco Monitoring Instrumentation” and Figure 8-2.

The Cisco WAN and application optimization solution incorporates NetQoS Performance Center to deliver powerful and integrated WAN traffic and performance monitoring. Refer to the following URL for more information about NetQos:

http://www.netqos.com/Cisco_wan_optimization/index.html

The NetQoS Performance Center provides flexible views and reporting to help effectively manage networks, applications, and devices. From this intuitive, Web-based network performance monitoring, the network operator can drill into detailed information provided by the following NetQoS product modules:

- NetQoS SuperAgent for end-to-end performance monitoring
- NetQoS ReporterAnalyzer for traffic analysis
- NetQoS NetVoyant for device performance management

The NetQoS Performance Center also facilitates integration with third-party products – making it possible, for example, to publish summary views of infrastructure performance in business portals such as Microsoft SharePoint.

*Figure 8-2 NetQoS Products*

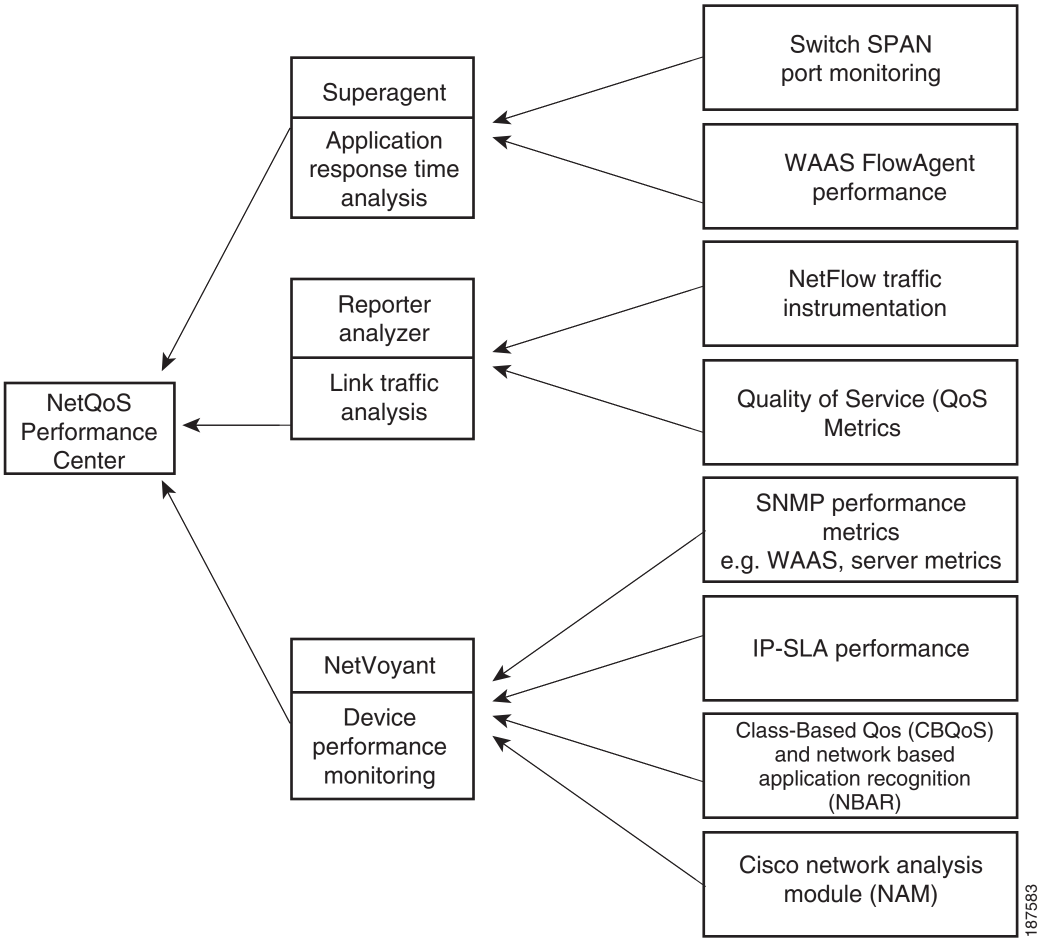
Switch SPAN port monitoring
WAAS FlowAgent performance
NetFlow traffic instrumentation
Quality of Service (QoS Metrics
SNMP performance metrics e.g. WAAS, server metrics
IP-SLA performance
Class-Based Qos (CBQoS) and network based application recognition (NBAR)
Cisco network analysis module (NAM)
Superagent
Application response time analysis
Reporter analyzer
Link traffic analysis
NetVoyant
Device performance monitoring
NetQoS Performance Center
187583

The NetQoS Performance Center enables users to view performance metrics from multiple data sources to facilitate troubleshooting, capacity planning, and management reporting. For example, correlating a slowdown in application access times with emerging congestion helps detect and correct the problem.

In this case, NetQoS SuperAgent reported the slowdown in application access times, which were detected through changes in application response times derived from FlowAgent instrumentation in WAE devices. This was correlated with emerging congestion caused by database replication traffic at peak office work times, which NetQoS ReporterAnalyzer reported. (ReporterAnalyzer takes metrics from IOS NetFlow-enabled routers and switches.)

NetQoS Performance Center provides *role-based views* of performance data. Users having different responsibilities, or working in different geographies, can view data in ways that are tailored to their needs. Figure 8-3 illustrates role-based views.

Consider the example of separate application teams where each team is responsible for monitoring the performance of a different set of applications. When appropriately configured, NetQoS Performance Center can present tailored reports that deliver performance information only for the applications relevant to each team.

***Figure 8-3 NetQoS Performance Center***

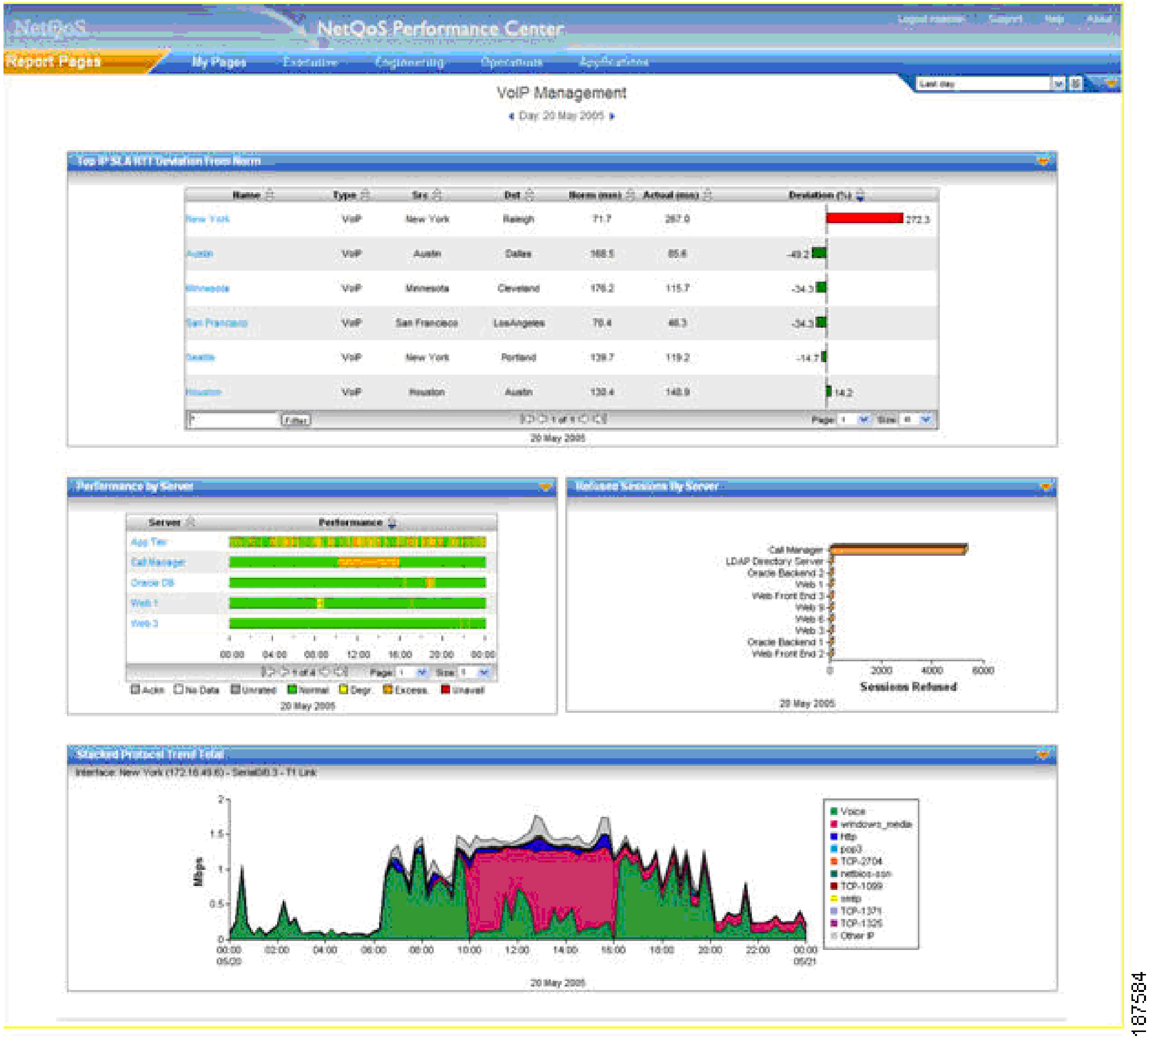

VoIP Management

Day 20 May 2005

Top IP SLA RTT Deviation From Norm

| Name | Type | Src | Dst | Norm (ms) | Actual (ms) | Deviation (%) |
|---|---|---|---|---|---|---|
| New York | VoIP | New York | Raleigh | 71.7 | 267.0 | 272.3 |
| Austin | VoIP | Austin | Dallas | 168.5 | 85.6 | -49.2 |
| Minnesota | VoIP | Minnesota | Cleveland | 176.2 | 115.7 | -34.3 |
| San Francisco | VoIP | San Francisco | LosAngeles | 70.4 | 46.3 | -34.3 |
| Seattle | VoIP | New York | Portland | 139.7 | 119.2 | -14.7 |
| Houston | VoIP | Houston | Austin | 130.4 | 148.9 | 14.2 |

## 8.2.1 NetQoS SuperAgent: Measuring Application Response Times

NetQoS SuperAgent is a passive application response time (ART) monitoring and reporting module that is installed in the data center. Through continuous performance analysis of TCP/IP applications, SuperAgent calculates performance baselines for applications, servers, and network links, and alerts the user when performance degrades beyond a preconfigured threshold or an automatically learned baseline. Operations staff can then isolate performance issues to the specific IT resource responsible for degraded performance.

The SuperAgent product monitors all the TCP application packets from the network into the data center and out again providing a way to measure Network round trip time (RTT), server response time (SRT), data transfer time (DTT), and much more. The SuperAgent product separates response time into application, network, and server delay components, enabling the customer to detect network performance bottlenecks rapidly to help protect application performance.

The SuperAgent product comprises at least one collection device and a management console. A collection device gathers relevant data from network devices and forwards the data to the SuperAgent Management Console. SuperAgent uses two types of collection devices: the SuperAgent Collector and the SuperAgent Aggregator. The SuperAgent Collector collects data from regular switches using a SPAN mechanism. The SuperAgent Aggregator collects data from an embedded agent (called a flow agent) residing in WAAS WAEs.

Integrating SuperAgent Aggregator and WAAS FlowAgent enables more accurate reporting of response times on optimized traffic where WAAS is deployed. Figure 8-4 shows a typical SuperAgent deployment comprising a SuperAgent Aggregator, a SuperAgent Collector, and a SuperAgent Management Console for long-term data storage and reporting. SuperAgent monitors end-to-end application delivery in optimized WAAS and nonoptimized environments. Users can quickly install and configure SuperAgent to monitor very large environments, making SuperAgent a practical choice for complex datacenters and networks.

***Figure 8-4 NetQoS SuperAgent Application Response Time Collection Architecture and WAAS***

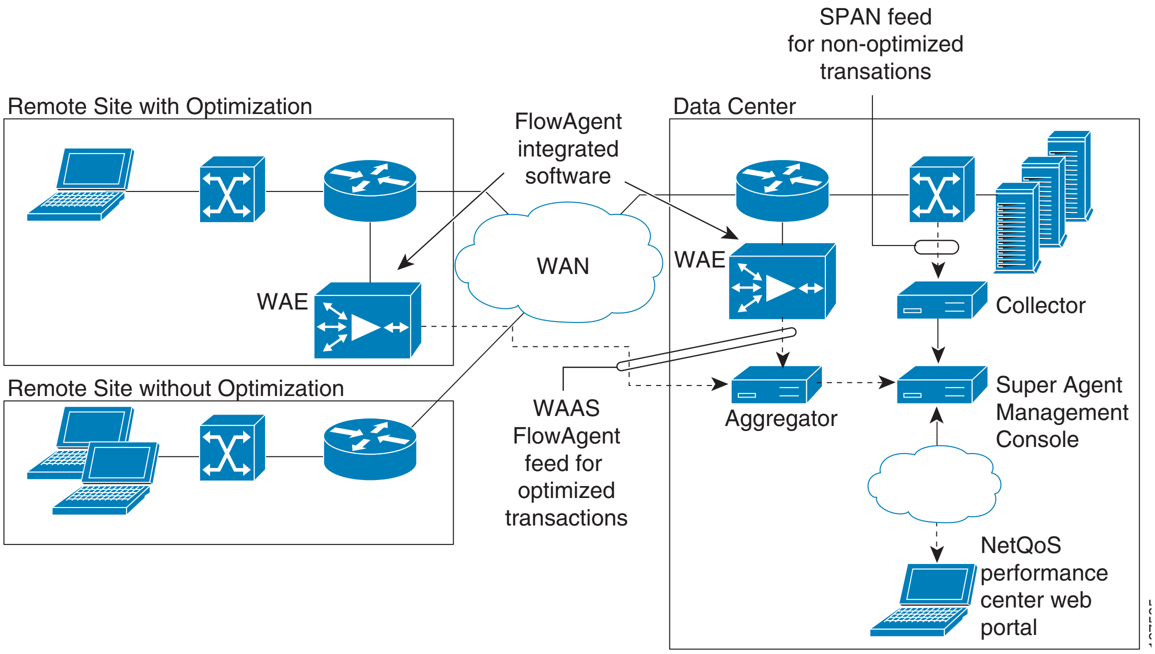

With a relatively small number of collection appliances (and without the installation of end user client or server based agents), SuperAgent can help the user to determine:

- Historically normal performance for individual servers, applications, and network links
- The specific origin of application, network, and server performance problems
- Precisely how planned and unplanned changes affect application delivery
- The impact of IT initiatives (for example, WAN and application optimization) on performance service level objectives

The SuperAgent reporting dashboard provides:

- Time-based baselines by application, server, and network
- Response-time graphs(see the example shown in Figure 8-5) showing how server, application, and network delays affect the delivery of any configured TCP/IP application, at particular sites or for user-groups at those sites.

***Figure 8-5 SuperAgent Response Time Composition Graphs***

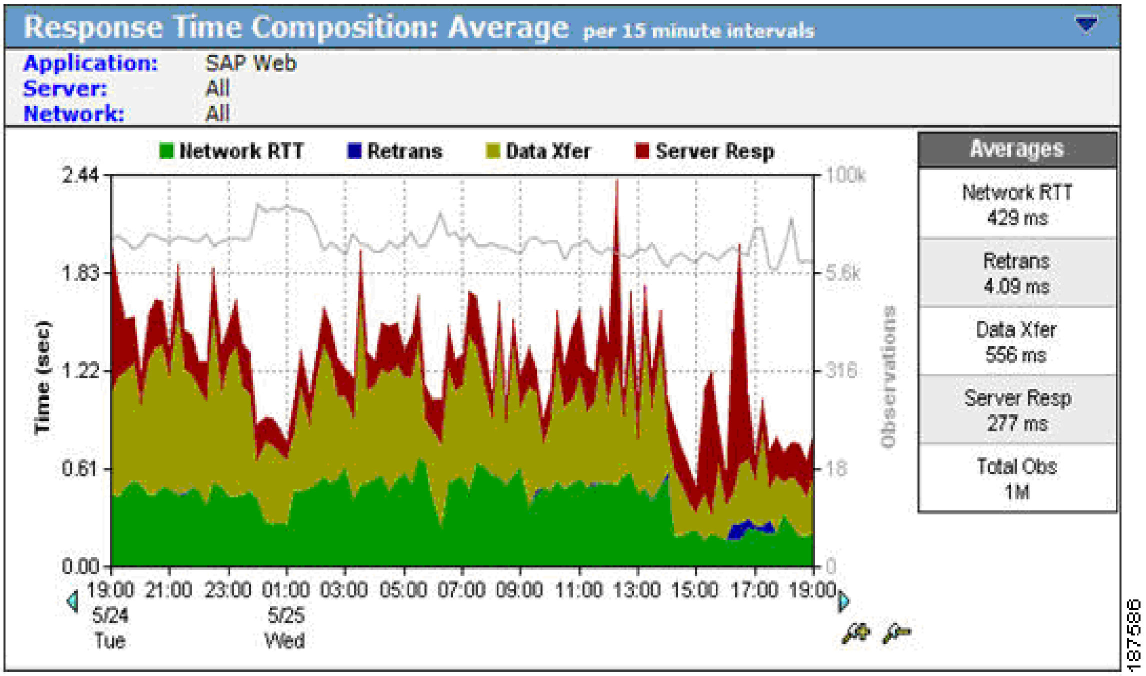

Response Time Composition: Average per 15 minute intervals

Application: SAP Web
Server: All
Network: All

Averages

Network RTT 429 ms

Retrans 4.09 ms

Data Xfer 556 ms

Server Resp 277 ms

Total Obs 1M

- Operations views (see the example shown in Figure 8-6) showing groups of worst performing networks, servers, and applications in time-based graphs.

*Figure 8-6 SuperAgent Operations View*

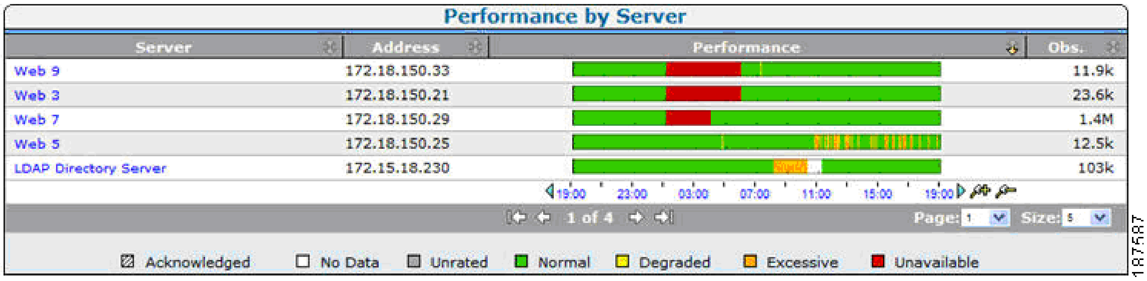

Performance by Server

| Server | Address | Performance | Obs. |
|---|---|---|---|
| Web 9 | 172.18.150.33 | | 11.9k |
| Web 3 | 172.18.150.21 | | 23.6k |
| Web 7 | 172.18.150.29 | | 1.4M |
| Web 5 | 172.18.150.25 | | 12.5k |
| LDAP Directory Server | 172.15.18.230 | | 103k |

- Performance maps (see the example shown in Figure 8-7) that help compare a wide range of metrics for similar IT resources, including servers, applications, and network links, to identify performance issues and candidates for improvement.

*Figure 8-7 SuperAgent Performance Maps*

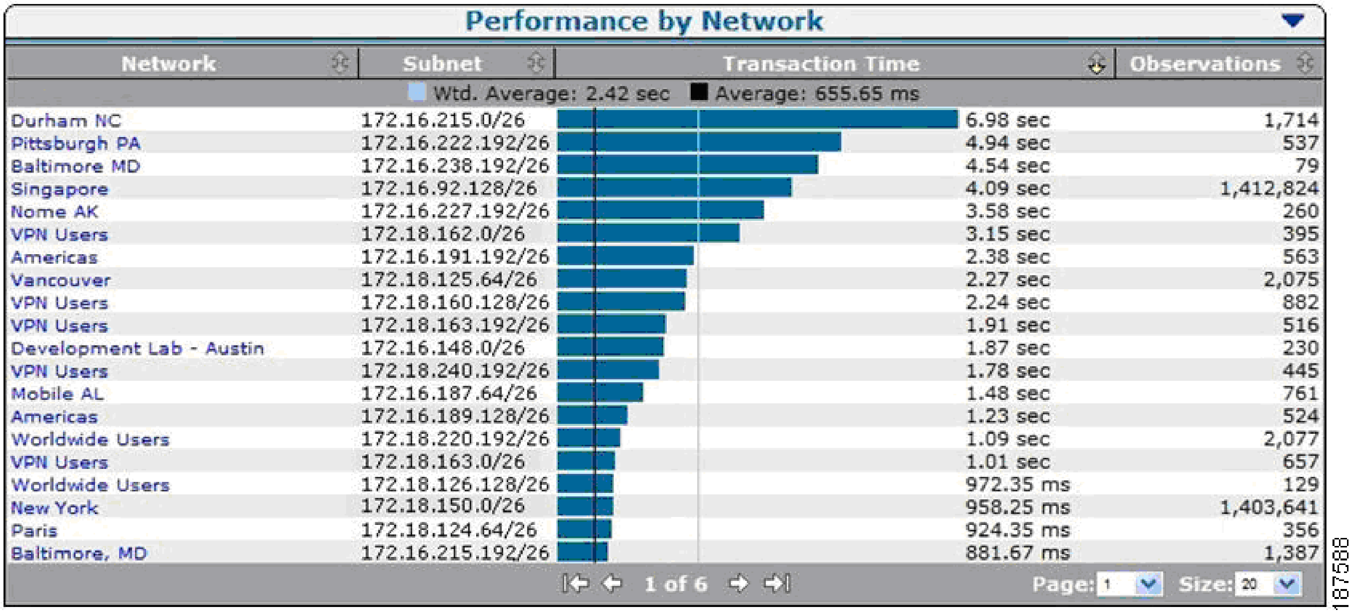

Performance by Network

Wtd. Average: 2.42 sec ■ Average: 655.65 ms

| Network | Subnet | Transaction Time | Observations |
|---|---|---|---|
| Durham NC | 172.16.215.0/26 | 6.98 sec | 1,714 |
| Pittsburgh PA | 172.16.222.192/26 | 4.94 sec | 537 |
| Baltimore MD | 172.16.238.192/26 | 4.54 sec | 79 |
| Singapore | 172.16.92.128/26 | 4.09 sec | 1,412,824 |
| Nome AK | 172.16.227.192/26 | 3.58 sec | 260 |
| VPN Users | 172.18.162.0/26 | 3.15 sec | 395 |
| Americas | 172.16.191.192/26 | 2.38 sec | 563 |
| Vancouver | 172.18.125.64/26 | 2.27 sec | 2,075 |
| VPN Users | 172.18.160.128/26 | 2.24 sec | 882 |
| VPN Users | 172.18.163.192/26 | 1.91 sec | 516 |
| Development Lab - Austin | 172.16.148.0/26 | 1.87 sec | 230 |
| VPN Users | 172.18.240.192/26 | 1.78 sec | 445 |
| Mobile AL | 172.16.187.64/26 | 1.48 sec | 761 |
| Americas | 172.16.189.128/26 | 1.23 sec | 524 |
| Worldwide Users | 172.18.220.192/26 | 1.09 sec | 2,077 |
| VPN Users | 172.18.163.0/26 | 1.01 sec | 657 |
| Worldwide Users | 172.18.126.128/26 | 972.35 ms | 129 |
| New York | 172.18.150.0/26 | 958.25 ms | 1,403,641 |
| Paris | 172.18.124.64/26 | 924.35 ms | 356 |
| Baltimore, MD | 172.16.215.192/26 | 881.67 ms | 1,387 |

- SLA reporting (see the example shown in Figure 8-8) reports whether SLAs are met, and the causes of any deficiencies, to support service level management (SLM).

*Figure 8-8 SuperAgent SLA Performance Detail*

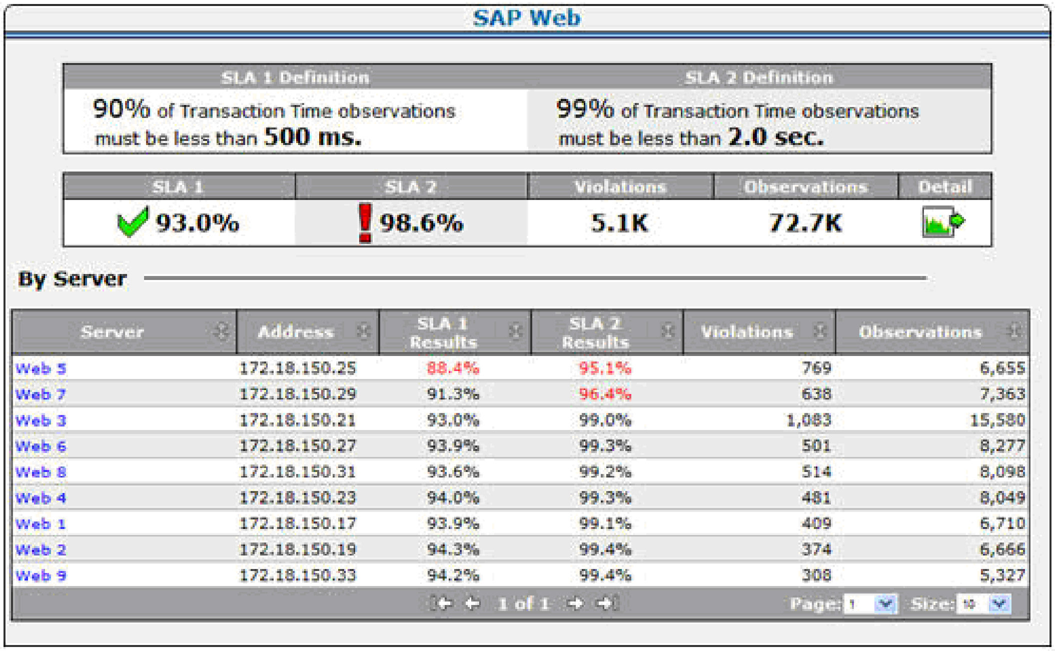

SAP Web

| SLA 1 Definition | SLA 2 Definition |
|---|---|
| 90% of Transaction Time observations must be less than 500 ms. | 99% of Transaction Time observations must be less than 2.0 sec. |

| SLA 1 | SLA 2 | Violations | Observations | Detail |
|---|---|---|---|---|
| ✓ 93.0% | ! 98.6% | 5.1K | 72.7K | |

By Server

| Server | Address | SLA 1 Results | SLA 2 Results | Violations | Observations |
|---|---|---|---|---|---|
| Web 5 | 172.18.150.25 | 88.4% | 95.1% | 769 | 6,655 |
| Web 7 | 172.18.150.29 | 91.3% | 96.4% | 638 | 7,363 |
| Web 3 | 172.18.150.21 | 93.0% | 99.0% | 1,083 | 15,580 |
| Web 6 | 172.18.150.27 | 93.9% | 99.3% | 501 | 8,277 |
| Web 8 | 172.18.150.31 | 93.6% | 99.2% | 514 | 8,098 |
| Web 4 | 172.18.150.23 | 94.0% | 99.3% | 481 | 8,049 |
| Web 1 | 172.18.150.17 | 93.9% | 99.1% | 409 | 6,710 |
| Web 2 | 172.18.150.19 | 94.3% | 99.4% | 374 | 6,666 |
| Web 9 | 172.18.150.33 | 94.2% | 99.4% | 308 | 5,327 |

## 8.3 NetQoS ReporterAnalyzer: Analyzing Link Traffic using NetFlow

NetQoS ReporterAnalyzer is a traffic analysis module that analyzes and reports how application traffic is affecting network performance by leveraging IOS NetFlow instrumentation present on Cisco routers and switches. It enables the user to see which applications are using bandwidth, who is using the bandwidth, and when. In a single reporting interface, ReporterAnalyzer can display trends that impact a global WAN infrastructure (for purposes of traffic policy monitoring, capacity planning and control) alongside reports of traffic anomalies (for example, malware, peer-to-peer, and unauthorized service protocols) detected from individual devices among many thousands on the network.

The ReporterAnalyzer product comprises the following components:

- Harvester, which passively collects and processes data from NetFlow enabled routers.
- Flow Manager, which aggregates data from multiple Harvesters.
- ReporterAnalyzer Console, which provides a web interface to display collected data.
- Data Storage Appliance (DSA), which stores data for as many as 500 interfaces, or Super DSA for up to 2500 interfaces.

In a typical ReporterAnalyzer implementation, as shown in Figure 8-9, a Harvester collects raw NetFlow data from configured routers, switches, and other NetFlow compliant devices, processes the data for collection by Flow Manager, and creates local archives for detailed flow analysis and reporting.

***Figure 8-9 ReporterAnalyzer Link Traffic Analysis Architecture***

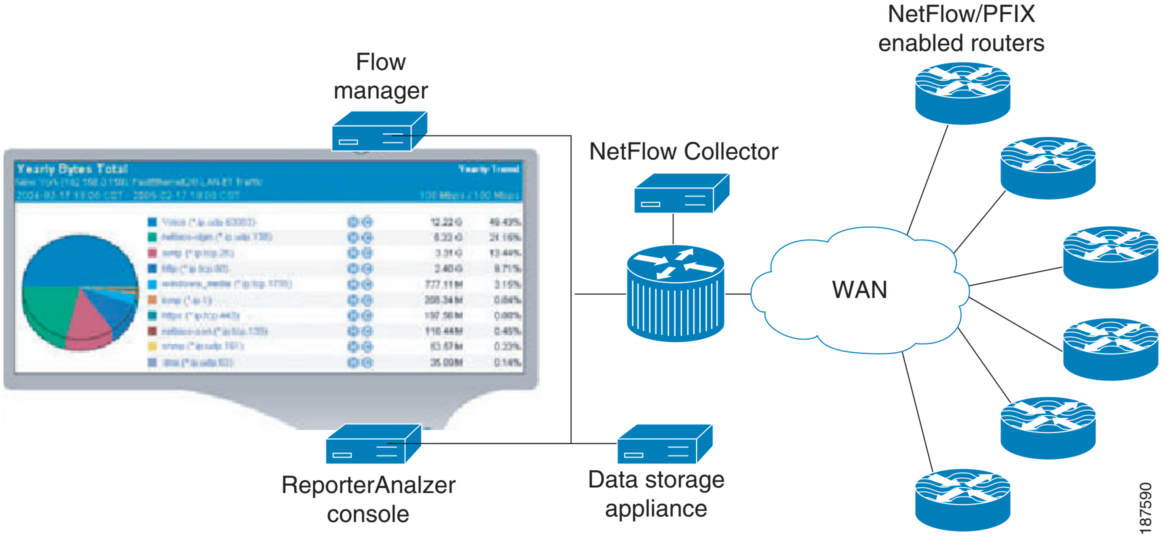

After enabling networking devices to export NetFlow data, and deploying a few Harvester NetFlow collection devices, ReporterAnalyzer users can:

- Identify interfaces, hosts, and applications that generate the most traffic or are most utilized
- View baselines for protocol and flow data
- Identify network traffic that exceeds specified thresholds
- View real-time alerts and reports
- Help pinpoint the cause of a network problem by reporting on, and drilling into, all traffic flows
- Identify bandwidth requirements for network applications and users
- Design and run reports based on user-selected criteria

The ReporterAnalyzer user interface provides views and analytics to support network troubleshooting, forensics, policy monitoring, capacity planning, and management reporting.

ReporterAnalyzer provides the following views.

- An enterprise overview page displays a summary of interfaces that exceed configured utilization thresholds and the top interfaces, protocols, and hosts for the entire network during the preceding 24 hours.
- Interface views show summary information, protocols, hosts, conversations, type of service (ToS), as shown in Figure 8-10, growth reports, baselines, and other data.

*Figure 8-10* ***ReporterAnalyzer Stacked Trend Plot Showing ToS Distribution on a Link***

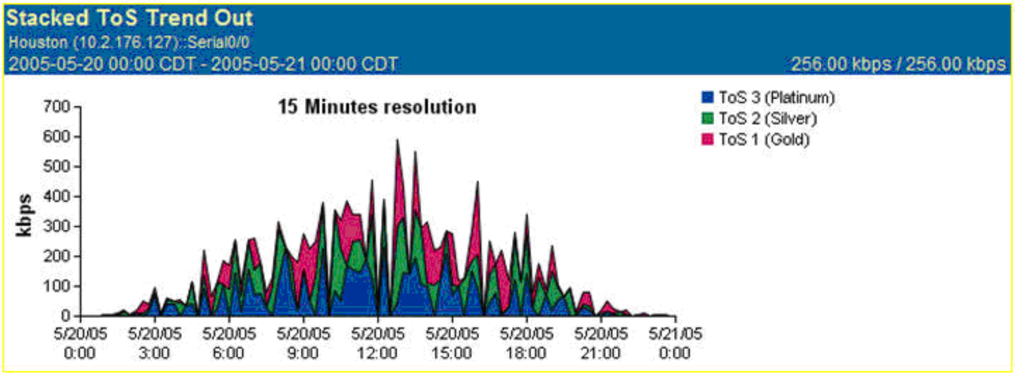

- Custom reports (see the example shown in Figure 8-11), which use a wizard to create interface, protocol, ToS, host, conversation, and combination reports for any timeframe, and that can be run at any designated time.

*Figure 8-11* ***ReporterAnalyzer Custom Report***

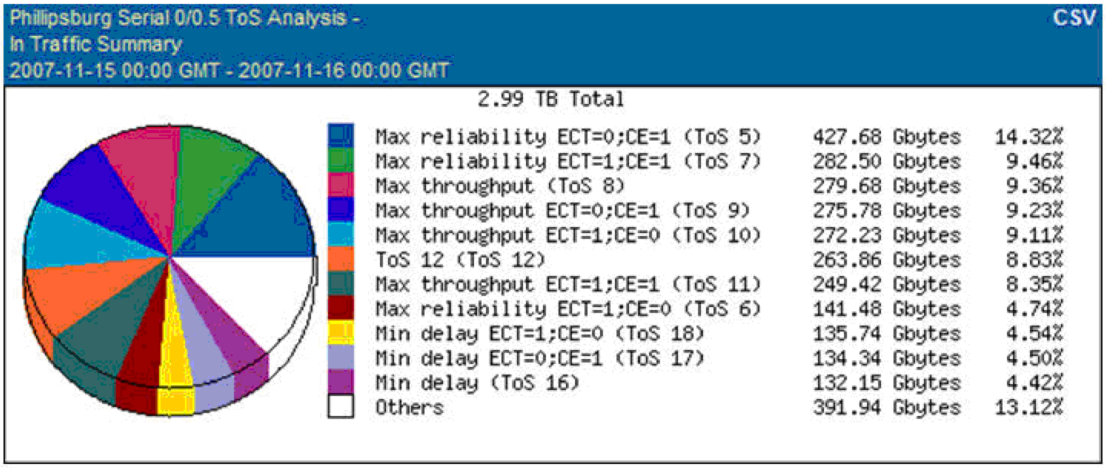
Phillipsburg Serial 0/0.5 ToS Analysis -
In Traffic Summary
2007-11-15 00:00 GMT - 2007-11-16 00:00 GMT
CSV
187592

2.99 TB Total

| | | |
|---|---|---|
| Max reliability ECT=0;CE=1 (ToS 5) | 427.68 Gbytes | 14.32% |
| Max reliability ECT=1;CE=1 (ToS 7) | 282.50 Gbytes | 9.46% |
| Max throughput (ToS 8) | 279.68 Gbytes | 9.36% |
| Max throughput ECT=0;CE=1 (ToS 9) | 275.78 Gbytes | 9.23% |
| Max throughput ECT=1;CE=0 (ToS 10) | 272.23 Gbytes | 9.11% |
| ToS 12 (ToS 12) | 263.86 Gbytes | 8.83% |
| Max throughput ECT=1;CE=1 (ToS 11) | 249.42 Gbytes | 8.35% |
| Max reliability ECT=1;CE=0 (ToS 6) | 141.48 Gbytes | 4.74% |
| Min delay ECT=1;CE=0 (ToS 18) | 135.74 Gbytes | 4.54% |
| Min delay ECT=0;CE=1 (ToS 17) | 134.34 Gbytes | 4.50% |
| Min delay (ToS 16) | 132.15 Gbytes | 4.42% |
| Others | 391.94 Gbytes | 13.12% |

- Flow forensics reports, which use a wizard (see the example shown in Figure 8-12) to create reports that analyze raw NetFlow data to provide insight into any protocol, host, and conversation on the network.

*Figure 8-12 ReporterAnalyzer Flow Forensics Wizard*

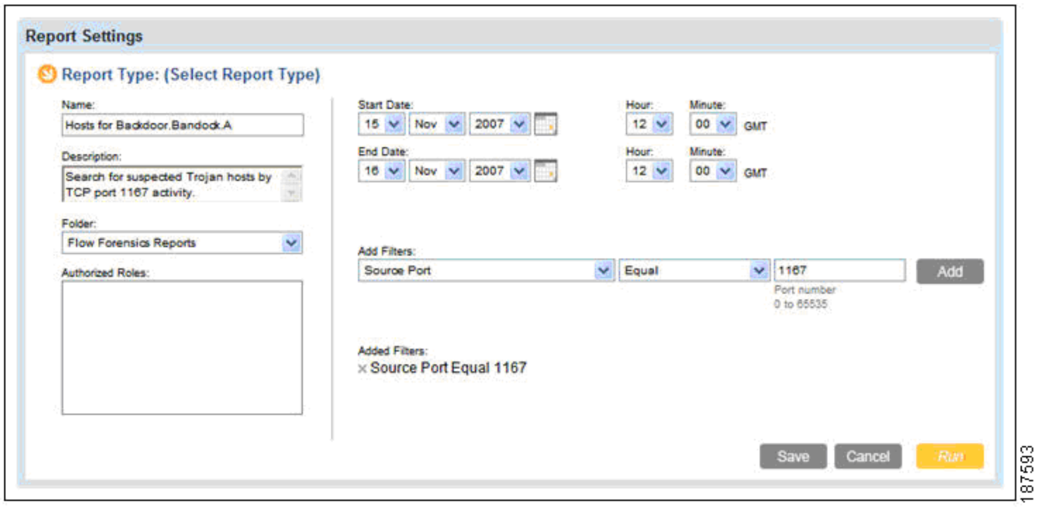

Report Settings

Report Type: (Select Report Type)

Name: Hosts for Backdoor.Bandock.A

Description: Search for suspected Trojan hosts by TCP port 1167 activity.

Folder: Flow Forensics Reports

Authorized Roles:

Start Date: 15 Nov 2007 Hour: 12 Minute: 00 GMT

End Date: 16 Nov 2007 Hour: 12 Minute: 00 GMT

Add Filters: Source Port Equal 1167

Port number 0 to 65535

Added Filters:
× Source Port Equal 1167

- Analysis reports, which use a wizard to compare collected data to an established threshold, and can be run on a schedule, or at any specified time.

# 8.4 NetQoS NetVoyant: Monitoring Device Performance and IP SLA

NetQoS NetVoyant is an appliance-based Simple Network Management Protocol (SNMP) monitoring and reporting module that helps users identify trends and potential problems early enough to potentially avoid lost productivity. NetVoyant takes advantage of the comprehensive SNMP performance data provided by Cisco devices to help IT administrators provide more consistent application delivery, optimize the network infrastructure for better application performance, and equip their staff to solve complex problems.

(SNMP) instrumentation from routers, switches, servers, frame relay circuits, logical segments, and wide area links can provide essential data for troubleshooting, capacity planning, and management reporting. Through SNMP instrumentation, Cisco products provide comprehensive device performance statistics, service level metrics, and application data. For example, Cisco WAE devices can report service statistics, configuration data, and resource consumption through SNMP.

NetVoyant monitors performance metrics from a variety of sources:

- Network and datacenter devices can be polled to report CPU and memory statistics, interface utilization, availability, configuration data, and other key metrics.
- IP SLA agents on Cisco routers can be configured and polled to provide latency, jitter, connection tests, and other statistics between a host router and specified targets.
- Network Based Application Recognition (NBAR), Remote Network Monitoring (RMON) metrics, and application response time (ART) from Cisco Network Analysis Modules (NAMs) can be polled.
- Class-based quality of service (CBQoS) management information bases (MIBs) on Cisco routers can be polled to provide utilization statistics for each class of service.

Figure 8-13 illustrates a typical NetVoyant distributed configuration. NetVoyant can be deployed on a single server or in a distributed configuration as shown.

*Figure 8-13 NetVoyant Device Performance Monitoring Architecture*

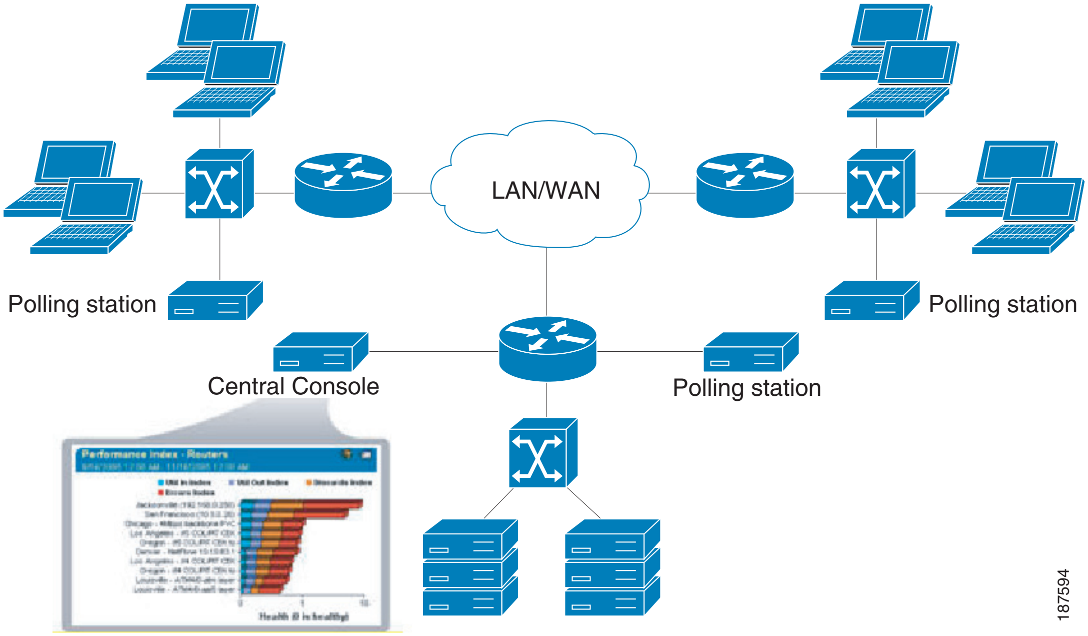

In a distributed configuration, a NetVoyant Master Console performs administration and reporting while remote Polling Stations discover and poll devices on the network. Because the volume of data collected from polling stations is much smaller than the volume of received SNMP data, on large networks it can be advantageous to place polling stations near the devices that they monitor to minimize management traffic on the WAN. During initial configuration, a wizard helps the administrator configure SNMP settings, initialize discovery processes, and determine the types of devices to monitor.

NetVoyant organizes reporting features according to user tasks:

- Management reports provide comprehensive survey views of devices and network performance using scorecards and summaries by device (as shown in Figure 8-14) and network segment.

*Figure 8-14 NetVoyant Management Views*

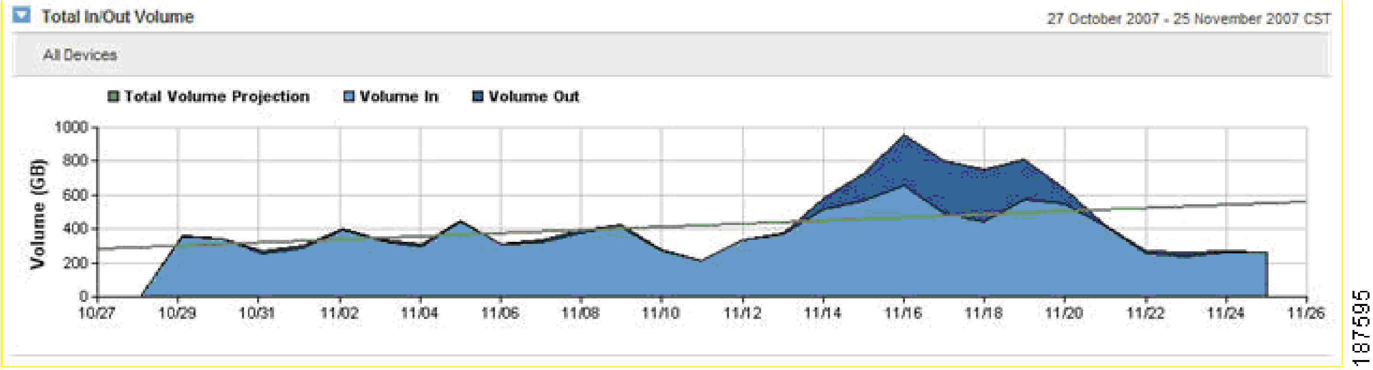

- Capacity planning reports (see the example shown in Figure 8-15) support planning for device and network link upgrades and help identify the resources with the fastest growth and characteristics approaching defined thresholds.

*Figure 8-15 Figure 8-15. NetVoyant Capacity Planning*

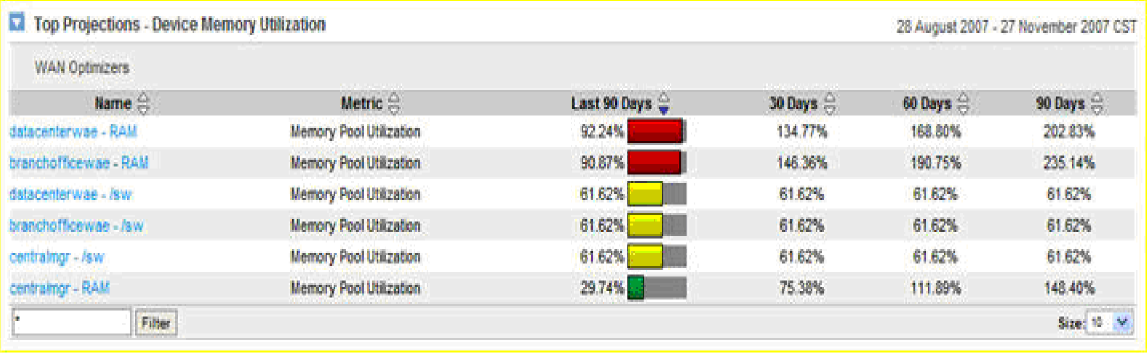

Top Projections - Device Memory Utilization 28 August 2007 - 27 November 2007 CST

WAN Optimizers

| Name | Metric | Last 90 Days | 30 Days | 60 Days | 90 Days |
|---|---|---|---|---|---|
| datacenterwae - RAM | Memory Pool Utilization | 92.24% 🟥 | 134.77% | 168.80% | 202.83% |
| branchofficewae - RAM | Memory Pool Utilization | 90.87% 🟥 | 146.36% | 190.75% | 235.14% |
| datacenterwae - /sw | Memory Pool Utilization | 61.62% 🟨 | 61.62% | 61.62% | 61.62% |
| branchofficewae - /sw | Memory Pool Utilization | 61.62% 🟨 | 61.62% | 61.62% | 61.62% |
| centralmgr - /sw | Memory Pool Utilization | 61.62% 🟨 | 61.62% | 61.62% | 61.62% |
| centralmgr - RAM | Memory Pool Utilization | 29.74% 🟩 | 75.38% | 111.89% | 148.40% |

- Service level reports (see the example shown in Figure 8-16) can help users to verify compliance with SLAs and to help track performance metrics with unexpected values, including data from IP SLA tests, VoIP statistics, and CBQoS reports.

*Figure 8-16 NetVoyant SLA Reports*

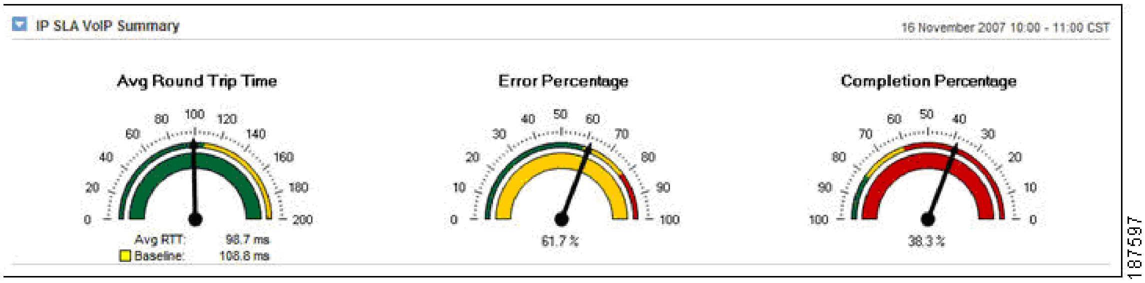

IP SLA VoIP Summary 16 November 2007 10:00 - 11:00 CST

Avg RTT: 98.7 ms
Baseline: 108.8 ms

- Operations reports (see the example shown in Figure 8-17) provide an operations-level view of devices on the network, with views that include the devices that are most unavailable, the interfaces that are most utilized, and events and alarms for SNMP data collection issues. In the view shown in Figure 8-17, the SNMP poller msfcrouter.qa.local was unavailable 10% of the time.

*Figure 8-17 NetVoyant Operations Reports*

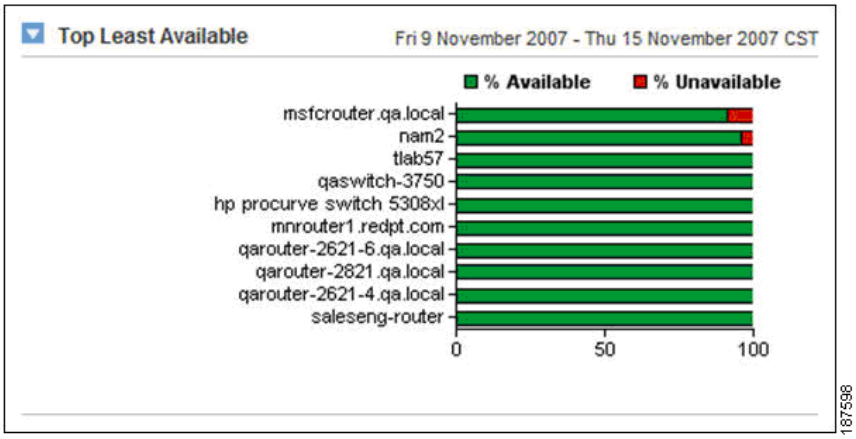

## 8.5 NAM: Granular Monitoring and Troubleshooting

Network Analysis Module (NAM) is a powerful integrated network monitoring tool designed to report how users experience network services and to help network operators ensure and improve network performance.

The Cisco WAN and application optimization solution incorporates NAM to help deliver a feature-rich, granular, and interactive troubleshooting capability that can substantially reduce the time and effort needed to isolate configuration and performance problems.

Figure 8-18 shows an example of how NAM is placed in the data center.

*Figure 8-18 Example of NAM Placement in the Data Center*

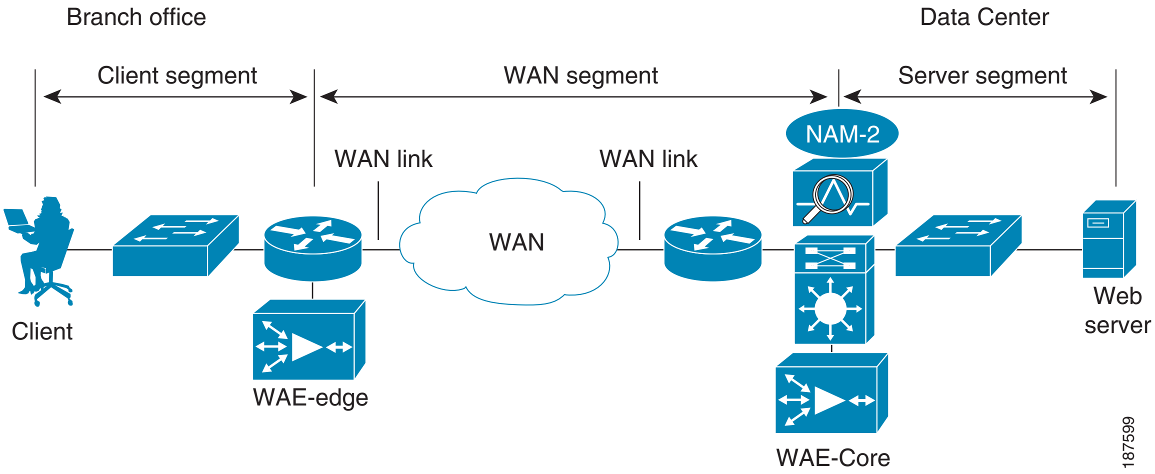

Cisco NAM is available as

- A single blade (NAM-1 and NAM-2) for the Cisco Catalyst 6500 Series Switch or the Cisco 7600 Series Router.
- A single network module (NM-NAM and NME-NAM) for the Cisco 2800 and 3800 Series Integrated Services Routers (ISR) and the Cisco 3700 Series Multiservice Access Routers.

NAM-2 is the deployment form incorporated into the current WAN and application optimization Solution architecture.

Cisco NAM combines embedded data collection and analysis features with a remotely accessible, web-based management console that delivers:

- **Real-time visibility** for troubleshooting, with the flexibility to apply filters in real-time - this sidesteps the need to pre-configure data-sources, servers, and applications in troubleshooting contexts;
- **Deep information granularity**, supporting drill-downs to individual conversations, with data-points collected over short time-frames;
- Short and long-term **traffic and performance reports** on individual conversations and hosts;
- **Performance analytics** covering TCP application response time metrics, DiffServ QoS, voice and video;
- User-defined trigger-based **packet capture**, filters and decodes;
- **Integrated NetFlow processing** for remote troubleshooting.

The NAM Traffic Analyzer GUI provides quick access to the configuration menus and presents easy-to-read reports.

# 8.6 Monitoring and Profiling Network and Application Usage

Cisco NAM can inspect packets to gather information about applications, hosts, and conversations. Application monitoring identifies each application that consumes bandwidth and how much, and detects which hosts are using which applications.

Host and conversation-pair monitoring reports bandwidth consumption per host and shows which hosts talk to each other, along with the amount of traffic each host generates. These metrics reveal usage patterns for users, and for router and switch, interface, server, and application resources.

For example, Figure 8-19 shows a report identifying the top 10 hosts on the network.

*Figure 8-19 Monitoring the Top 10 Hosts on the Network*

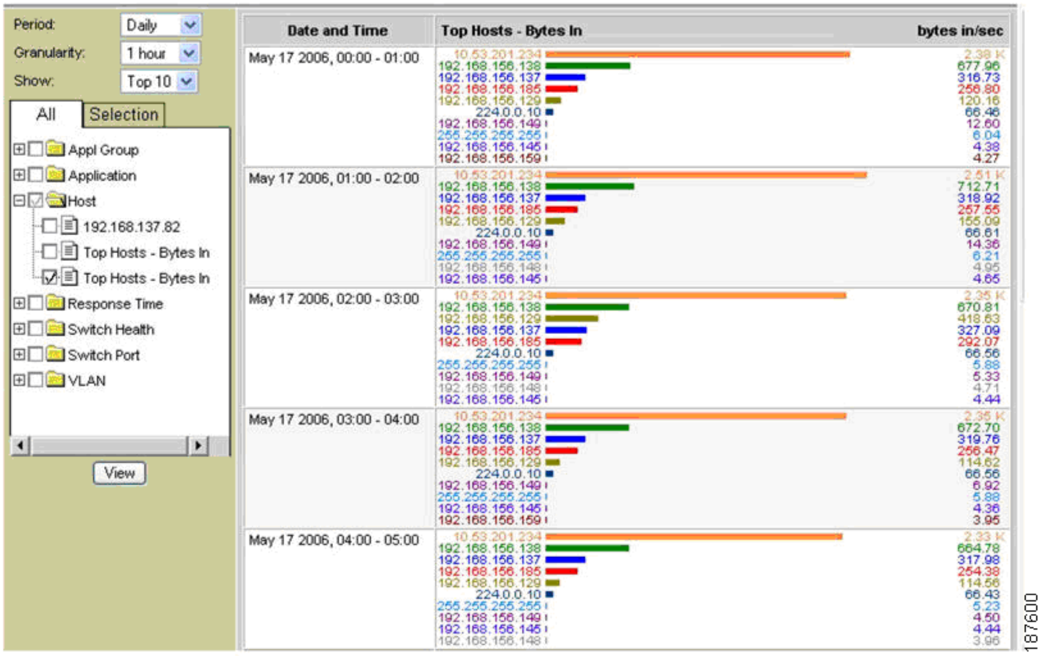

Period: Daily
Granularity: 1 hour
Show: Top 10

| Date and Time | Top Hosts - Bytes In | bytes in/sec |
|---|---|---|
| May 17 2006, 00:00 - 01:00 | 10.53.201.234 | 2.38 K |
| | 192.168.156.138 | 677.96 |
| | 192.168.156.137 | 316.73 |
| | 192.168.156.185 | 258.80 |
| | 192.168.156.129 | 120.16 |
| | 224.0.0.10 | 66.46 |
| | 192.168.156.149 | 12.60 |
| | 255.255.255.255 | 6.04 |
| | 192.168.156.145 | 4.38 |
| | 192.168.156.159 | 4.27 |
| May 17 2006, 01:00 - 02:00 | 10.53.201.234 | 2.51 K |
| | 192.168.156.138 | 712.71 |
| | 192.168.156.137 | 318.92 |
| | 192.168.156.185 | 257.55 |
| | 192.168.156.129 | 155.09 |
| | 224.0.0.10 | 66.61 |
| | 192.168.156.149 | 14.36 |
| | 255.255.255.255 | 6.21 |
| | 192.168.156.148 | 4.95 |
| | 192.168.156.145 | 4.65 |
| May 17 2006, 02:00 - 03:00 | 10.53.201.234 | 2.35 K |
| | 192.168.156.138 | 670.81 |
| | 192.168.156.129 | 418.63 |
| | 192.168.156.137 | 327.09 |
| | 192.168.156.185 | 292.07 |
| | 224.0.0.10 | 66.56 |
| | 255.255.255.255 | 5.88 |
| | 192.168.156.149 | 5.33 |
| | 192.168.156.148 | 4.71 |
| | 192.168.156.145 | 4.44 |
| May 17 2006, 03:00 - 04:00 | 10.53.201.234 | 2.35 K |
| | 192.168.156.138 | 672.70 |
| | 192.168.156.137 | 319.76 |
| | 192.168.156.185 | 256.47 |
| | 192.168.156.129 | 114.62 |
| | 224.0.0.10 | 66.56 |
| | 192.168.156.149 | 6.92 |
| | 255.255.255.255 | 5.88 |
| | 192.168.156.145 | 4.36 |
| | 192.168.156.159 | 3.95 |
| May 17 2006, 04:00 - 05:00 | 10.53.201.234 | 2.33 K |
| | 192.168.156.138 | 664.78 |
| | 192.168.156.137 | 317.98 |
| | 192.168.156.185 | 254.38 |
| | 192.168.156.129 | 114.56 |
| | 224.0.0.10 | 66.43 |
| | 255.255.255.255 | 5.23 |
| | 192.168.156.149 | 4.50 |
| | 192.168.156.145 | 4.44 |
| | 192.168.156.148 | 3.96 |

# 8.7 Granular Live and History Reporting

The live conversation reporting capability of NAM exposes potential problems that would otherwise be masked if critical report updates occurred less frequently or within aggregated traffic data. This information helps facilitate quick identification and verification of traffic anomalies so current and impending problems can be rapidly resolved.

NAM delivers granular, real-time snapshots of bandwidth usage and performance, and can be easily configured to deliver continuous history views focused on specific hosts and conversations to help isolate intermittent user or conversation problems. NAM can also collect Data for a specific period and analyzed after the event to discover when an anomaly has occurred so it can be quickly resolved.

Figure 8-20 illustrates sample history reports.

*Figure 8-20 History Reports for WAN and Application Optimization Validation*

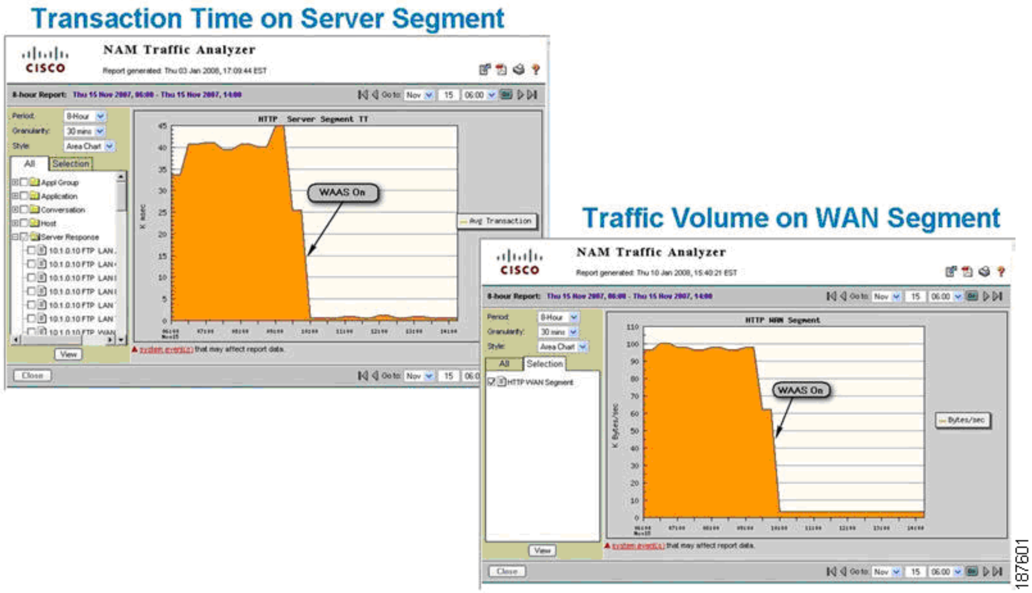

Transaction Time on Server Segment

NAM Traffic Analyzer

Report generated: Thu 03 Jan 2008, 17:09:44 EST

8-hour Report: Thu 15 Nov 2007, 06:00 - Thu 15 Nov 2007, 14:00

Period: 8-Hour
Granularity: 30 mins
Style: Area Chart

▲ system event(s) that may affect report data.

Traffic Volume on WAN Segment

NAM Traffic Analyzer

Report generated: Thu 10 Jan 2008, 15:40:21 EST

8-hour Report: Thu 15 Nov 2007, 06:00 - Thu 15 Nov 2007, 14:00

Period: 8-Hour
Granularity: 30 mins
Style: Area Chart

☑ HTTP WAN Segment

▲ system event(s) that may affect report data.

## 8.7.1 Transaction-Aware Response-Time Measurement, Monitoring, and Baselining

Cisco NAM implements newly enhanced transaction-based response-time measurement features that passively gather data on TCP-based client/server requests and acknowledgements. The response-time monitoring capability of Cisco NAM provides intelligent information about client, server, and application latency.

By giving fast, easy access to specific host and conversation response-time metrics, Cisco NAM can assist IT staff in troubleshooting application performance problems, analyzing application behavior and performing pre- and post- deployment monitoring of application optimization and acceleration services, as well as defining and tracking targeted service levels.

Figure 8-21 shows examples of granular transaction time views that could be generated to help track intermittent performance problems.

*Figure 8-21 Application Response-Time Monitoring*

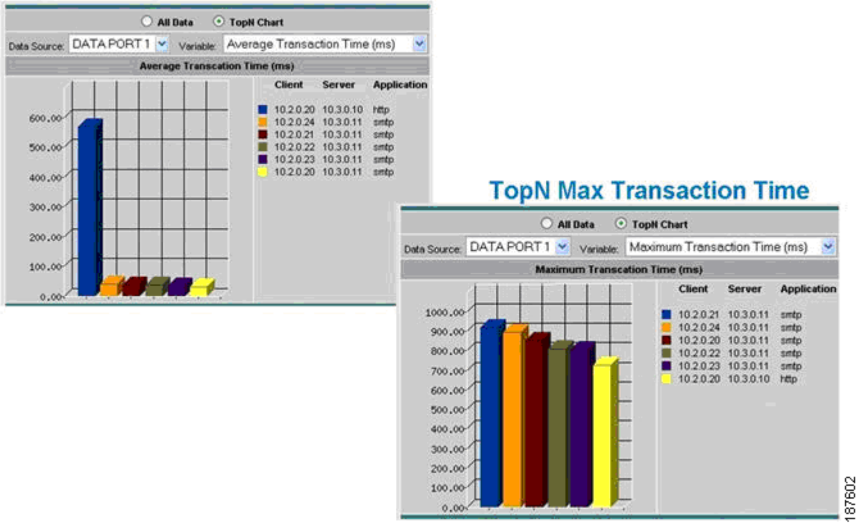

○ All Data ● TopN Chart

Data Source: DATA PORT 1 Variable: Average Transaction Time (ms)

Average Transcation Time (ms)

TopN Max Transaction Time

○ All Data ● TopN Chart

Data Source: DATA PORT 1 Variable: Maximum Transaction Time (ms)

Maximum Transcation Time (ms)

Figure 8-22 illustrates a live drill-down to an individual conversation.

*Figure 8-22 Detailed Application Response Times for a Specific Server/Client*

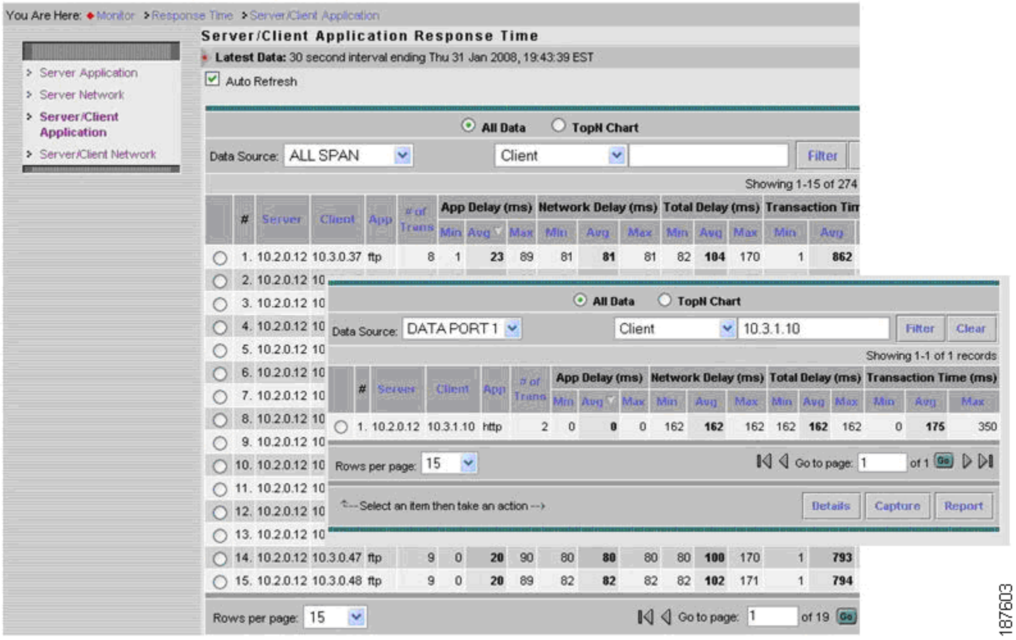

You Are Here: Monitor > Response Time > Server/Client Application

Server/Client Application Response Time

Latest Data: 30 second interval ending Thu 31 Jan 2008, 19:43:39 EST

☑ Auto Refresh

● All Data ○ TopN Chart

Data Source: ALL SPAN Client

Showing 1-15 of 274

| | # | Server | Client | App | # of Trans | App Delay (ms) | | | Network Delay (ms) | | | Total Delay (ms) | | | Transaction Tim | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | | Min | Avg | Max | Min | Avg | Max | Min | Avg | Max | Min | Avg |
| ○ | 1. | 10.2.0.12 | 10.3.0.37 | ftp | 8 | 1 | **23** | 89 | 81 | **81** | 81 | 82 | **104** | 170 | 1 | **862** |
| ○ | 2. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 3. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 4. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 5. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 6. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 7. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 8. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 9. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 10. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 11. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 12. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 13. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 14. | 10.2.0.12 | 10.3.0.47 | ftp | 9 | 0 | **20** | 90 | 80 | **80** | 80 | 80 | **100** | 170 | 1 | **793** |
| ○ | 15. | 10.2.0.12 | 10.3.0.48 | ftp | 9 | 0 | **20** | 89 | 82 | **82** | 82 | 82 | **102** | 171 | 1 | **794** |

Rows per page: 15 Go to page: 1 of 19

● All Data ○ TopN Chart

Data Source: DATA PORT 1 Client 10.3.1.10

Showing 1-1 of 1 records

| | # | Server | Client | App | # of Trans | App Delay (ms) | | | Network Delay (ms) | | | Total Delay (ms) | | | Transaction Time (ms) | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | | Min | Avg | Max | Min | Avg | Max | Min | Avg | Max | Min | Avg | Max |
| ○ | 1. | 10.2.0.12 | 10.3.1.10 | http | 2 | 0 | **0** | 0 | 162 | **162** | 162 | 162 | **162** | 162 | 0 | **175** | 350 |

Rows per page: 15 Go to page: 1 of 1

←-- Select an item then take an action -->

## 8.7.1.1 Alarms, Packet Captures and Decodes for Troubleshooting

To support isolating intermittent problems, users can set thresholds and alarms on various network parameters such as increased utilization and severe application response delays. When a potential problem is identified, the packet stream can be automatically captured and decoded to help resolve the problem.

Captures can be performed using a Web browser, and decodes can be viewed through the Traffic Analyzer GUI while data is being captured. This helps the user quickly pinpoint and resolve problems when troubleshooting. Because NAM typically uses a local Switched Port Analyzer (SPAN) configuration to perform a capture, capture traffic does not stress the operational user network.

*Figure 8-23 Using NAM to Capture and Decode Packets*

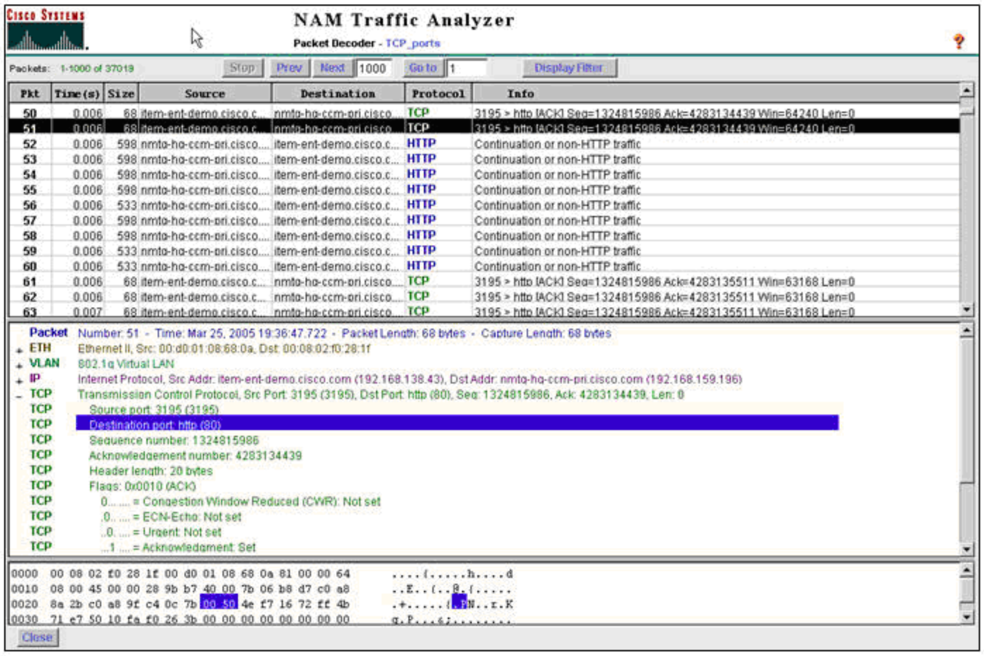

| Pkt | Time (s) | Size | Source | Destination | Protocol | Info |
|---|---|---|---|---|---|---|
| 50 | 0.006 | 68 | item-ent-demo.cisco.c... | nmto-hq-ccm-pri.cisco.... | TCP | 3195 > http [ACK] Seq=1324815986 Ack=4283134439 Win=64240 Len=0 |
| 51 | 0.006 | 68 | item-ent-demo.cisco.c... | nmto-hq-ccm-pri.cisco.... | TCP | 3195 > http [ACK] Seq=1324815986 Ack=4283134439 Win=64240 Len=0 |
| 52 | 0.006 | 598 | nmto-hq-ccm-pri.cisco.... | item-ent-demo.cisco.c... | HTTP | Continuation or non-HTTP traffic |
| 53 | 0.006 | 598 | nmto-hq-ccm-pri.cisco.... | item-ent-demo.cisco.c... | HTTP | Continuation or non-HTTP traffic |
| 54 | 0.006 | 598 | nmto-hq-ccm-pri.cisco.... | item-ent-demo.cisco.c... | HTTP | Continuation or non-HTTP traffic |
| 55 | 0.006 | 598 | nmto-hq-ccm-pri.cisco.... | item-ent-demo.cisco.c... | HTTP | Continuation or non-HTTP traffic |
| 56 | 0.006 | 533 | nmto-hq-ccm-pri.cisco.... | item-ent-demo.cisco.c... | HTTP | Continuation or non-HTTP traffic |
| 57 | 0.006 | 598 | nmto-hq-ccm-pri.cisco.... | item-ent-demo.cisco.c... | HTTP | Continuation or non-HTTP traffic |
| 58 | 0.006 | 598 | nmto-hq-ccm-pri.cisco.... | item-ent-demo.cisco.c... | HTTP | Continuation or non-HTTP traffic |
| 59 | 0.006 | 533 | nmto-hq-ccm-pri.cisco.... | item-ent-demo.cisco.c... | HTTP | Continuation or non-HTTP traffic |
| 60 | 0.006 | 533 | nmto-hq-ccm-pri.cisco.... | item-ent-demo.cisco.c... | HTTP | Continuation or non-HTTP traffic |
| 61 | 0.006 | 68 | item-ent-demo.cisco.c... | nmto-hq-ccm-pri.cisco.... | TCP | 3195 > http [ACK] Seq=1324815986 Ack=4283135511 Win=63168 Len=0 |
| 62 | 0.006 | 68 | item-ent-demo.cisco.c... | nmto-hq-ccm-pri.cisco.... | TCP | 3195 > http [ACK] Seq=1324815986 Ack=4283135511 Win=63168 Len=0 |
| 63 | 0.007 | 68 | item-ent-demo.cisco.c... | nmto-hq-ccm-pri.cisco.... | TCP | 3195 > http [ACK] Seq=1324815986 Ack=4283135511 Win=63168 Len=0 |

```
  Packet  Number: 51 - Time: Mar 25, 2005 19:36:47.722 - Packet Length: 68 bytes - Capture Length: 68 bytes
+ ETH     Ethernet II, Src: 00:d0:01:08:68:0a, Dst: 00:08:02:f0:28:1f
+ VLAN    802.1q Virtual LAN
+ IP      Internet Protocol, Src Addr: item-ent-demo.cisco.com (192.168.138.43), Dst Addr: nmto-hq-ccm-pri.cisco.com (192.168.159.196)
- TCP     Transmission Control Protocol, Src Port: 3195 (3195), Dst Port: http (80), Seq: 1324815986, Ack: 4283134439, Len: 0
  TCP       Source port: 3195 (3195)
  TCP       Destination port: http (80)
  TCP       Sequence number: 1324815986
  TCP       Acknowledgement number: 4283134439
  TCP       Header length: 20 bytes
  TCP       Flags: 0x0010 (ACK)
  TCP         0... .... = Congestion Window Reduced (CWR): Not set
  TCP         .0.. .... = ECN-Echo: Not set
  TCP         ..0. .... = Urgent: Not set
  TCP         ...1 .... = Acknowledgment: Set
```

```
0000  00 08 02 f0 28 1f 00 d0 01 08 68 0a 81 00 00 64     ....(.....h....d
0010  08 00 45 00 00 28 9b b7 40 00 7b 06 b8 d7 c0 a8     ..E..(..@.{.....
0020  8a 2b c0 a8 9f c4 0c 7b 00 50 4e f7 16 72 ff 4b     .+.....{.PN..r.K
0030  71 e7 50 10 fa f0 26 3b 00 00 00 00 00 00 00 00     q.P...&;........
```

## 8.7.1.2 Analyzing QoS

NAM supports QoS troubleshooting, through its support of the Differentiated Services Monitoring (DSMON) MIB, which monitors traffic by differentiated services code point (DSCP) allocations defined by QoS policies.

Using its DiffServ monitoring capabilities, NAM can help identify the hosts, conversations, and applications participating in each grouping of DiffServ classes. This information can be used to validate and tune QoS allocations in detailed troubleshooting workflows, to detect incorrectly marked or unauthorized traffic, and in other troubleshooting scenarios. Figure 8-24 provides a QoS monitoring example.

*Figure 8-24 QoS Monitoring Using DSMON*

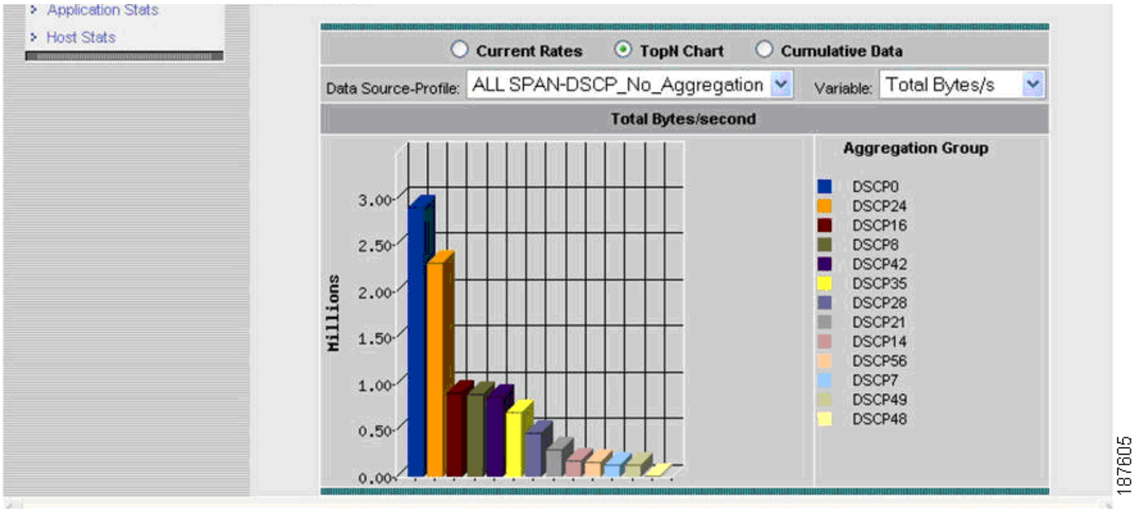

○ Current Rates ● TopN Chart ○ Cumulative Data

Data Source-Profile: ALL SPAN-DSCP_No_Aggregation

Variable: Total Bytes/s

# 8.8 Configuration Management

The WAN and application optimization solution spans multiple network devices and IOS features. However, WAN optimization is only one of the powerful features delivered by those network devices and appliances. Because decisions on configuration will need to take place within a wider context, this section only briefly outlines the functions and products most central to WAN optimization deployments.

## 8.8.1 General Configuration Management Functions

General configuration covers the generic, network-wide configuration of devices and IOS features as the network evolves.

- **Inventory management** helps to enables network operations to maintain a comprehensive, up-to-date record of devices and modules on the enterprise network, and to generate tailored inventory reports.
- **Image management** helps supports the distribution of software images to the network devices, and the maintenance of up-to-date and consistent images with role-based access control.
- **Change management** supports the implementation of changes to the network inventory, configuration, and images, and reports on current and historical modifications.
- **Network backup** supports image and configuration backups and archiving.
- **Compliance support** helps enforce enterprise configuration policies, and comply with regulatory requirements such as Sarbanes-Oxley (SOX), using predefined, configurable role-based workflows and detailed policy and regulation compliance reports.

The CiscoWorks product suite is available to support the general configuration management for enterprise users. Refer towww.cisco.com for details.

## 8.8.2 Dedicated Configuration Management

Several products and features in the WAN and application optimization solution require dedicated configuration management to aid or enable customer the configuration of rich features on non-IOS devices, or to increase automation and user guidance for configuring complex IOS features, such as QoS.

Of these, the product most relevant to the WAN and application optimization solution is the Cisco WAAS Central Manager. WAAS Central Manager is a scalable, secure, and simple function that runs on WAE appliances. WAAS Central Manager provides a centralized configuration mechanism, along with basic WAAS deployment monitoring and reporting. WAAS Central Manager can be accessed from a Web browser, allowing secure remote management.

More information about these and other dedicated management products are available on www.cisco.com.

*Figure 8-25 A View of Detailed Application Response Times for a Specific Server/Client*

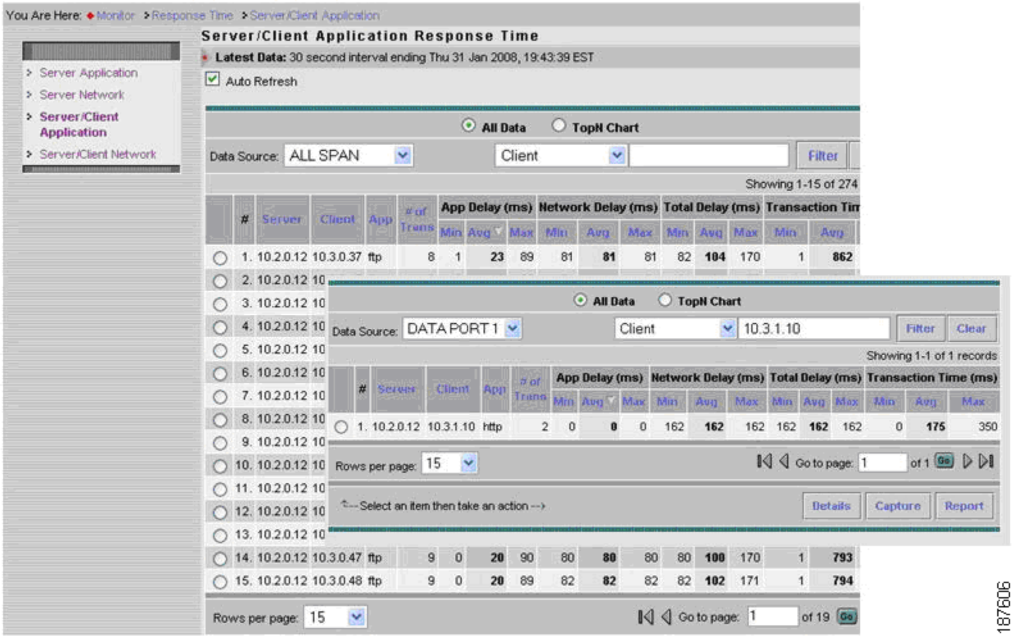

You Are Here: ◆Monitor >Response Time >Server/Client Application

Server/Client Application Response Time

Latest Data: 30 second interval ending Thu 31 Jan 2008, 19:43:39 EST

☑ Auto Refresh

● All Data ○ TopN Chart

Data Source: ALL SPAN Client

Showing 1-15 of 274

| | # | Server | Client | App | # of Trans | App Delay (ms) | | | Network Delay (ms) | | | Total Delay (ms) | | | Transaction Tim | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | | Min | Avg | Max | Min | Avg | Max | Min | Avg | Max | Min | Avg |
| ○ | 1. | 10.2.0.12 | 10.3.0.37 | ftp | 8 | 1 | **23** | 89 | 81 | **81** | 81 | 82 | **104** | 170 | 1 | **862** |
| ○ | 2. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 3. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 4. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 5. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 6. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 7. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 8. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 9. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 10. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 11. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 12. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 13. | 10.2.0.12 | 10 | | | | | | | | | | | | | |
| ○ | 14. | 10.2.0.12 | 10.3.0.47 | ftp | 9 | 0 | **20** | 90 | 80 | **80** | 80 | 80 | **100** | 170 | 1 | **793** |
| ○ | 15. | 10.2.0.12 | 10.3.0.48 | ftp | 9 | 0 | **20** | 89 | 82 | **82** | 82 | 82 | **102** | 171 | 1 | **794** |

Rows per page: 15 Go to page: 1 of 19

● All Data ○ TopN Chart

Data Source: DATA PORT 1 Client 10.3.1.10

Showing 1-1 of 1 records

| | # | Server | Client | App | # of Trans | App Delay (ms) | | | Network Delay (ms) | | | Total Delay (ms) | | | Transaction Time (ms) | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | | Min | Avg | Max | Min | Avg | Max | Min | Avg | Max | Min | Avg | Max |
| ○ | 1. | 10.2.0.12 | 10.3.1.10 | http | 2 | 0 | **0** | 0 | 162 | **162** | 162 | 162 | **162** | 162 | 0 | **175** | 350 |

Rows per page: 15 Go to page: 1 of 1

←-- Select an item then take an action -->