WAN and Application Optimization Solution Guide

Document Version 1.0
April 2008

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

Abstract

This guide describes the Cisco WAN and application optimization solution. The guide provides detailed technical information about the design and implementation of the solution.

The WAN and application optimization solution combines Cisco products and technologies to deliver solutions to specific WAN and application optimization challenges. This guide helps its readers understand these challenges, and design and implement networking infrastructures to meet the challenges.

Key Technologies

Application optimization, network monitoring, traffic classification, WAN optimization

Target Audience

Technical personnel who design and implement enterprise networks.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO..

Contents

Figures

Tables

1 About this Guide

This guide describes the Cisco WAN and application optimization solution. The guide provides detailed technical information about the design and implementation of the solution.

The WAN and application optimization solution combines Cisco products and technologies to deliver solutions to specific WAN and application optimization challenges. This guide helps its readers understand these challenges, and design and implement networking infrastructures to meet the challenges.

1.1 How This Guide Is Organized

This guide contains the following chapters:

- Customer Challenges

  This chapter describes the challenges customers face as the number of branch offices and their networking demands increase.

- WAN and Application Optimization Overview

  This chapter provides an overview of the WAN and application optimization solution, with a focus on business requirements.

- Cisco Monitoring Instrumentation

  This chapter describes the monitoring instrumentation provided in the WAN and application optimization solution.

- Traffic Classification

  This chapter describes how traffic is classified in the WAN and application optimization solution.

- An Overview of WAN and Application Optimization Technologies

  This chapter describes the specific technologies used in the WAN and application optimization solution.

- Network Management

  This chapter describes the network management technologies used in the WAN and application optimization solution.

- WAN and Application Optimization Design and Implementation

  This chapter provides detailed descriptions, with configuration examples, of the various deployment models used in the WAN and application optimization solution.

1.2 Intended Audience

This guide is for technical personnel involved in the specification, design, and implementation of specific WAN and application optimization solutions.

2 Customer Challenges

This chapter summarizes the challenges that enterprises face when delivering applications across corporate wide-area networks (WANs).

The WAN is the connective fabric that holds a distributed organization together. Because the WAN has bandwidth restrictions and latency issues, however, application performance suffers without WAN and application optimization. With optimization, IT organizations can substantially improve application delivery to ensure secure, cost effective, and acceptable application performance to meet business needs.

2.1 Consolidating Data Centers and Server Infrastructure

Enterprise servers and applications continue to be consolidated and centralized. For example, previously it was common for remote sites to have their own file and various application servers. The cost of maintaining servers remotely is high and new regulations and compliances such as Sarbanes Oxley (SOX) and Health Insurance Portability and Accountability Act (HIPAA) push costs even higher and drive server consolidation in the data center. IT organizations face new challenges of providing LAN-like response times across the corporate WAN even as data and processing become more centralized.

2.2 Globalization

The workforce is increasingly located outside of headquarters. These remote users demand the same quality of experience when using applications and services that their headquarters colleagues enjoy connected to a server over a LAN. Remote access should not result in lower productivity due to slower response time. IT organizations face constant challenges to achieve the same response time and “always-on” services for remote users. A survival strategy is also needed so that remote locations can function alone in the event of resource failures.

2.3 Improving Business Continuity and Disaster Recovery Processes

An enterprise’s ability to failover seamlessly from one data center to another and the ability to back up data in all remote locations is essential. This requires moving massive amount of data across a WAN in real time. At the same time, enterprises want to reduce the costs of data backup and disaster recovery. Even worse, if a scheduled backup operation spills over into regular working hours, remote users may find that their application response times become unacceptable.

2.4 Delay-Sensitive Applications

Real-time applications, such as Voice over Internet Protocol (VoIP) and interactive video, have strict requirements on transport delay, jitter, packet loss and bandwidth availability. Therefore, it is essential to prioritize different traffic types to minimize congestion risk in the end-to-end service path in order to deliver high quality voice or video, as well as provide preferential treatment to business-critical applications.

2.5 Badly Behaved Applications on the WAN

Too many businesses deploy new applications without completely understanding how the applications will work in a complex, distributed network. Many business applications are developed without considering requirements relevant to performance in a real network (for example, WAN latency and limited bandwidth).

Even worse, many application architectures, which are designed for use over a LAN, do not provide efficient performance across corporate WANs. Unfortunately, LAN protocols are "chatty." For example, an especially bad variant of "chatty" occurs when applications break messages into small data blocks. The application works in a serial manner: an acknowledgement is required for each data block before the next one can be sent. This can require many round trips to send just one message, causing significant application delay. Much of the delay comes from time on the wire. In this example, latency degrades application performance and limits application throughput. Adding bandwidth does not solve such performance issues. For example, Microsoft Exchange and Common Internet File System (CIFS), Network File System (NFS), and many web-based applications have latency issues. In fact, these applications show increasing response times the further they are deployed from the data center.

Although many applications can be altered to accommodate latency and bandwidth restrictions, modifying applications is not always viable. For example, shrink-wrapped applications usually cannot be modified. In such cases, a solution outside the applications is needed. Deploying WAN optimization and application acceleration tools in the network addresses latency and performance problems, but do not require any changes to the applications.

2.6 ″Webified" Applications

Computing is changing. We are now in the early stages of implementing "webified" applications. These new application environments demand a new type of network that can support the unique requirements of Web-based application technologies.

For example, Web-enabled applications require many more connections between the client and server. New acceleration technologies must deal with the increased number of connections to achieve better application performance. Moving HTTP and XML enables developers to include more objects, such as graphics, that increase the amount of transferred data. Migrating applications to Service Oriented Architecture (SOA) radically changes network demands. Web applications are usually worse with respect to bandwidth requirements as they have to render the screen. For example, a branch user using the SAP client will only get requested data. However, a user using SAP over the Web must receive formatting and graphical data.

2.7 Delivering Rich Content and Rolling out New Services

Large organizations struggle to ensure that employees have the latest content, whether it is training collateral, compliance documentation, email, or video. IT organizations are constantly challenged to deliver more services, such as large file transfers (e.g., medical imaging and computer-aided design (CAD) files), VoIP, and streaming video. Such applications contribute to high bandwidth growth. However, IT organizations are also expected to simultaneously reduce operational expenses (OpEx). In practice, cost bandwidth costs still represent a significant portion of recurring OpEx for many organizations. Therefore, IT organizations want to exploit WAN optimization technologies to extend constrained bandwidth resources and avoid costly bandwidth upgrades.

2.8 The Network Must Truly Support the Business

IT organizations are constantly challenged to deploy new applications to drive user productivity and gain competitive advantage. There is a direct correlation between the application environment and the network solutions required. Network architectures often need to be transformed to meet new business requirements. The Cisco “network as a platform” approach allows businesses to use the network to gain significant benefits for diverse sets of applications and infrastructure architectures. By leveraging the Cisco “network as a platform” approach, we can empower our customers to rapidly roll out new applications and services across their organizations, allowing them to maintain business competitiveness.

3 WAN and Application Optimization Overview

This chapter presents the Cisco WAN and application optimization framework, provides an overview of the solution, and introduces Cisco WAN and application optimization products and technologies. It also briefly discusses the solution deployed in different places in the network.

3.1 The Cisco Vision

In modern enterprises, the network is an essential component of application performance. Cisco Systems empowers network managers to deploy critical business applications on integrated networks to increase productivity and gain competitive advantages. Cisco delivers advanced, integrated WAN and application optimization solutions to support a broad set of applications with different requirements, from IP communications to transaction-oriented applications. Cisco continues to add optimization techniques and delivers the "network as the platform."

Security directly affects network and application performance. A complete, holistic solution delivers more than comprehensive WAN and application optimization capabilities, but also cooperates with security components to protect business against disruption. Cisco offers a network-based, end-to-end systems approach that evolves with business needs and enables the opportunities generated from future technical innovations.

Figure 3-1. WAN and Application Optimization in the Network

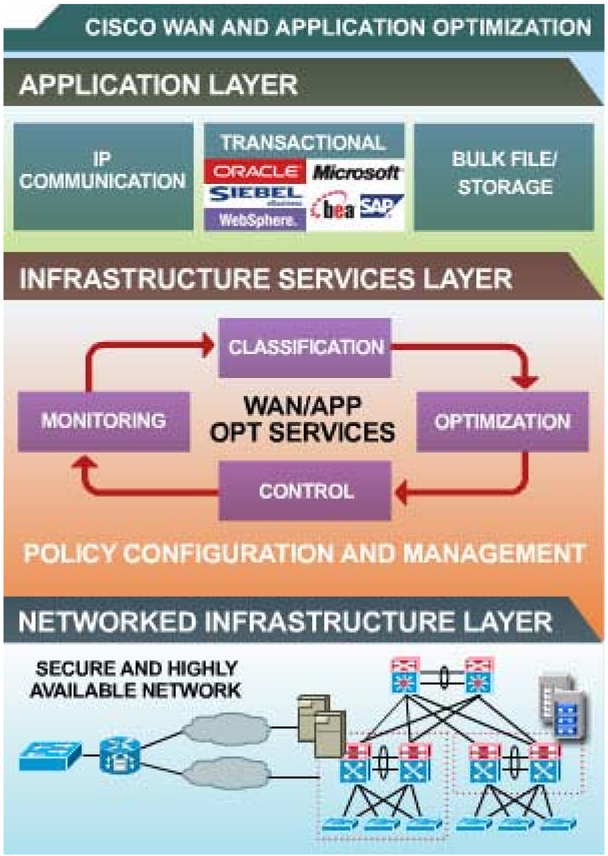
CISCO WAN AND APPLICATION OPTIMIZATION
APPLICATION LAYER
IP COMMUNICATION
TRANSACTIONAL
ORACLE
Microsoft
SIEBEL
WebSphere.
bea
SAP
BULK FILE/ STORAGE
INFRASTRUCTURE SERVICES LAYER
CLASSIFICATION
MONITORING
WAN/APP OPT SERVICES
OPTIMIZATION
CONTROL
POLICY CONFIGURATION AND MANAGEMENT
NETWORKED INFRASTRUCTURE LAYER
SECURE AND HIGHLY AVAILABLE NETWORK

Cisco WAN and application optimization is an architectural solution consisting of a set of tools and techniques working together to improve the reliability, performance, and delivery of applications securely across your network. A strategic systems approach uses the network to identify applications running in the network, gains end-to-end visibility, optimizes the network and applications, and controls and protects business critical traffic.

The Cisco WAN and application optimization solution comprises five critical components for effective application delivery. The following sections are brief descriptions of the five architectural components and the associated techniques and technologies. Subsequent chapters (4 through 8) provide more details of each of the components.

3.1.1 Classification

An intelligent network must evolve to become an active participant in application delivery. The network must be application-aware to assess and control application performance to ensure that valuable shared network resources are used efficiently. Prior to controlling traffic, the network needs to learn the requirements of and automatically discover applications running on the network. Techniques must go beyond simple IP address or TCP port recognition by supporting dynamic and migration port assignments using deep packet inspection technologies.

3.1.2 Optimization

Several techniques, when applied to network traffic, dramatically improve application performance and availability/reliability, decrease latency, improve bandwidth utilization, and bolster security:

- TCP Flow Optimization (TFO) – Improves the TCP stack and brings uniformity to TCP sessions. Mitigates the inherent lack of performance in TCP slow start and general flow control, which can slow data transfers. TFO techniques fill the pipe and reduce latency, resulting in faster transfers and optimal bandwidth use.
- Advanced Compression – Data redundancy elimination (DRE) replaces matching byte streams with a signature to significantly reduce the amount of data sent over the WAN. Signatures are maintained in libraries on opposite sides of the peering devices and enable up to 100:1 compression ratios. Standard (LZ) compression further compresses nonredundant data for maximum compression.
- Path Optimization – Each networked application is matched to the best path, ensuring application availability.
- Server Optimization – Reduces server workloads using techniques such as server load balancing (SLB), connection management, and offloading Secure Socket Layer (SSL).
- Secure WAN – Firewalls, SSL encryption, and techniques that minimize denial-of-service and other threats protect applications and critical business information assets.
- Secure VPN – Technologies promote low-latency paths by enabling direct spoke–to-spoke communications.
- DNS Optimization – Accelerating DNS lookups helps to ensure speedy application delivery.
- Enterprise Content Delivery Network (ECDN) – Improves the performance and reliability of content and application delivery across the WAN. ECDN typically comprises caching, policy-based distribution, redirection, and content management. Together, these components enable enterprises to efficiently distribute content to its remote branch offices.

3.1.3 Control

Quality of service (QoS) techniques ensure that business-critical traffic is not negatively affected by less important traffic, and that controls conform with established business policies and priorities.

3.1.4 Monitoring

Successful application delivery requires IT organizations to continuously identify applications on the network, ensuring acceptable business-critical application performance while controlling or eliminating non-critical applications.

Controlling performance requires visibility into network and application behavior. Not only does monitoring verify that policies are correctly implemented, but data acquired through monitoring can drive the generation and enforcement of new dynamic policies.

3.1.5 Network Management

Management tools gather network application- and network-performance information, which is integrated into a series of comprehensive reports to provide visibility into the network and applications. Configuration management tools also centrally define policies and perform system-based change and configuration management.

3.2 Solution Components

Cisco WAN and application optimization provides a comprehensive solution comprising several products and technologies. This section lists the Cisco products and technologies that implement the five architectural components described in the preceding sections. These architectural components are implemented in dedicated appliances and blades, and in network router features.

3.2.1 Classification

- IOS Network Based Application Recognition (NBAR)

3.2.2 Optimization

- Cisco Wide Area Application Services (WAAS) or Wide Area Application Engine (WAE)
- IOS Performance Routing (PfR)
- Cisco Application Control Engine (ACE)
- IOS Dynamic Multipoint Virtual Private Network (DMVPN)

3.2.3 Control

- IOS QoS

3.2.4 Monitoring

- IOS NetFlow
- IOS IP Service Level Agreement (SLA)

- Cisco WAAS Flow Agent

3.2.5 Network Management

- Cisco Network Analysis Module-2 (NAM-2) for Cisco Catalyst 6000 Series
- NetQoS SuperAgent
- NetQoS ReporterAnalyzer

3.3 Deploying WAN and Application Optimization

WAN and application optimization solutions are primarily deployed in the data center and branch. As the Cisco WAN and application optimization solution evolves, it will touch more places in the network.

A "network as a platform" approach uses the network to identify applications on the network, gains end-to-end visibility, optimizes applications, and controls and protects business-critical traffic.

Figure 3-2. End-to-End WAN and Application Optimization

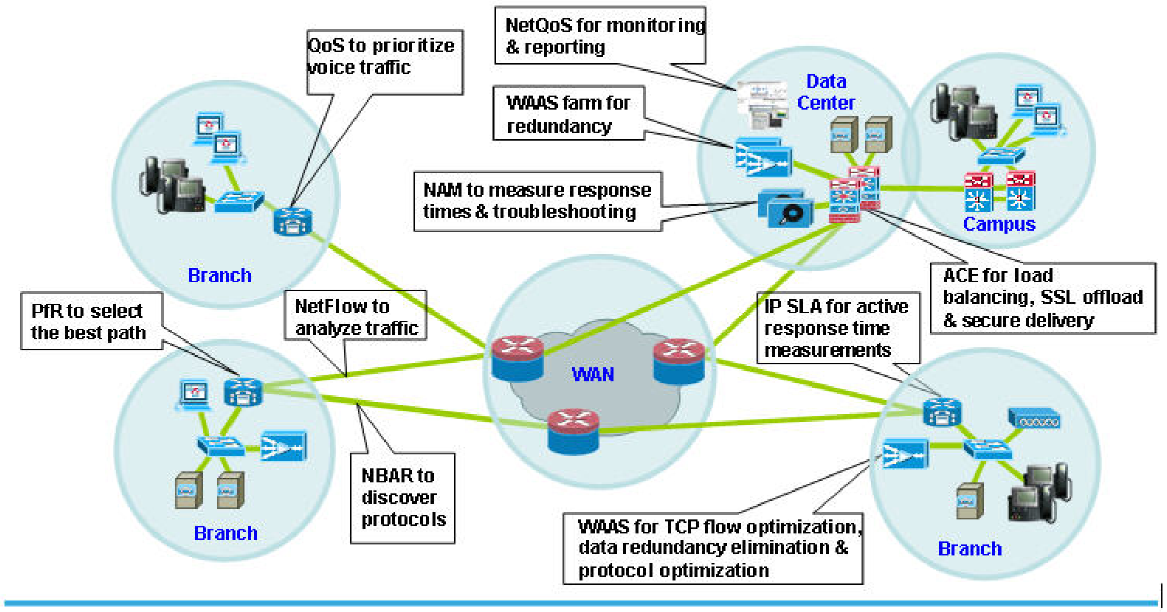
NetQoS for monitoring & reporting
QoS to prioritize voice traffic
Data Center
WAAS farm for redundancy
NAM to measure response times & troubleshooting
Campus
Branch
PfR to select the best path
NetFlow to analyze traffic
ACE for load balancing, SSL offload & secure delivery
IP SLA for active response time measurements
WAN
NBAR to discover protocols
Branch
WAAS for TCP flow optimization, data redundancy elimination & protocol optimization
Branch

As discussed in the preceding sections, WAN and application optimization is not a single technique. It is a collection of techniques and tools working cooperatively to improve application performance. For example, in Figure 3-2, various techniques and tools are enabled in different places in the network.

Inside the branch, NetFlow and NBAR are enabled in the branch access router to provide extensive visibility into the network and applications. With visibility into the applications and their utilization, IT operations can apply QoS policies in the branch router to establish transmission priorities of the application mix. A WAAS appliance can be deployed to apply a suite of WAN optimization and application acceleration technologies to dramatically improve application performance. When the branch has dual links, performance can be further enhanced by selecting the optimal path by using PfR.

Inside the data center, ACE is deployed to improve application performance, from SSL acceleration to load balancers. For example, ACE can make intelligently decide which server can send requests to yield further performance improvement. SSL acceleration is also enabled to handle the processing required to decrypt or encrypt traffic in order to offload the server.

In addition, performance management tools are deployed to support and protect business goals and objectives on an ongoing basis. NAM is deployed in the data center to measure application response times and troubleshooting. NetQoS Performance Center is used for centralized monitoring and reporting.

4 Cisco Monitoring Instrumentation

Understanding and addressing application performance issues brings visibility into how the business actually uses the network resources, and with abilities to measure how well applications are performing.

This chapter summarizes the key monitoring instrumentation technologies that provide essential information and sources of data for meeting the needs of the key performance management disciplines that optimize the networks and applications. Chapter 8 will describe the performance monitoring tools that consume this monitoring instrumentation data.

Figure 4-1 below outlines a general process that can be used to incrementally increase understanding of one's network and progressively deploy measurable improvements and adjustments as required.

Figure 4-1. WAN and Application Optimization Life Cycle

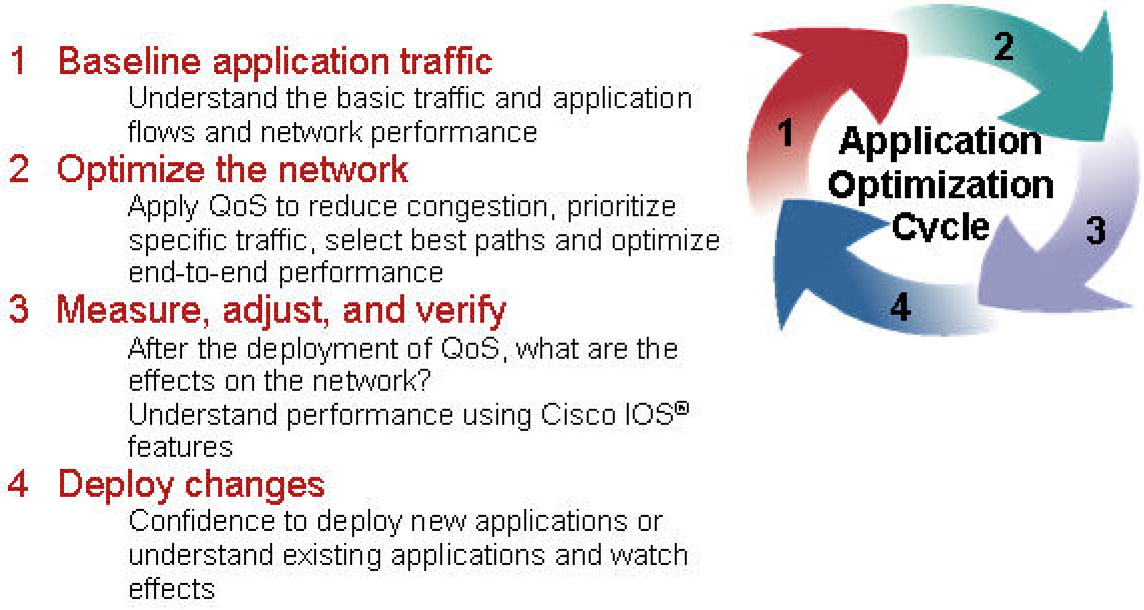

1 Baseline application traffic
Understand the basic traffic and application flows and network performance

2 Optimize the network
Apply QoS to reduce congestion, prioritize specific traffic, select best paths and optimize end-to-end performance

3 Measure, adjust, and verify
After the deployment of QoS, what are the effects on the network?
Understand performance using Cisco IOS® features

4 Deploy changes
Confidence to deploy new applications or understand existing applications and watch effects

4.1 Profiling and Baselining

The first step to WAN and application optimization is to profile network activity by establishing a reference from which service quality and application delivery effectiveness can be measured.

The profile of a network describes the traffic patterns and resource bottlenecks of a network. This identifies for the network operator the links and protocols that are the best candidates for optimization. Through profiling, a network engineer can focus on only those network components whose optimization will help improve and develop baselines as a performance benchmark.

Baselining is the establishment of acceptable network behavior. This includes understanding available bandwidth, identifying a normal pattern of network behavior such as network delays and what applications are running on the network, understanding each application's behavior (and requirements) on the network, and measuring application response times. For example, while not consistent with a daily average, baselining should capture and account for behaviors such as non-working weekend days that are less stressful on the network. Network administrators need to know the acceptable range for network performance before they can make reliable conclusions about possible performance degradation. With proper baselining, administrators can differentiate between consistent network behavior and anomalous (candidates for improvement) network behavior.

A few of the goals in baselining are as follows:

4.1.1 Ensure Network Stability

Complete internetwork communications can be easily obstructed if a network device such as a server or a single segment in a LAN becomes unreachable. The same is true if a server behind a router within the campus LAN environment or even behind the WAN cannot be contacted. Many different scenarios can cause problems in a large network and being able to maintain stability is a paramount concern of network managers.

4.1.2 Ensure Network Reliability

Many upper-layer applications present in today’s enterprise networks require connection-based processing during communications from one device to another. Maintaining a consistent connection is essential when critical communications take place between network devices, such as a workstation and a server. Being able to maintain low latency between a database and client machine, for instance, would be very important for applications that rely on constant access to the database.

Cisco IOS instrumentation provides a good starting point for creating a network performance baseline through the following components:

- NetFlow
- IPSLA
- NBAR
- CBQoS MIB

NetFlow provides a good source of traffic flow information for capturing normal and abnormal behaviors on the network. Additionally, standardized SNMP MIBs from individual devices provide basic information about the network such as traffic volume by byte, errors, utilization on interfaces, etc. NBAR, a traffic identification and classification engine built into IOS, can discover the types of applications that are present on the network. Together, NetFlow, MIBs, and NBAR provide a comprehensive baseline about the physical network and the paths application flows take as they utilize the network.

Creating response time baseline is important to the success of an IT organization in establishing service quality levels. Active and passive response time measurements are two methodologies for measuring application response times. Cisco IP SLA is the active method. Cisco WAAS Flow agent, Cisco NAM and NetQoS SuperAgent implement the passive method.

There is no one single source of information for baselining your network and applications. IT organizations will need to use different monitoring instrumentation data in order to gain a solid understanding of the normal behavior of the applications, the network, and IT resources.

4.1.3 Optimize the Network

Once you have end-to-end visibility of the network and the applications, you can then determine which optimization tools and technologies to utilize to best meet the requirements. The second step is to apply the optimization or control techniques to enhance application performance.

4.1.4 Measure, Adjust, and Verify

The third step is to assess the effectiveness of each successive WAN optimization initiative. This includes continuously monitoring and collecting information about the network and application behavior, and comparing the behavior before and after successive WAN optimization initiatives.

For example, when new QoS policies have just been deployed, you want to measure the effects of the network. CBQoS MIB from individual devices provide information about the network before and after applying the QoS policies. Similarly, after deploying WAAS, you want to determine the effectiveness of WAAS before and after compression and acceleration. WAAS Flow agent provides such information.

Measuring application response times for key applications both before and after WAN optimization and control techniques allows IT organizations to determine if the changes achieve desirable results. At the same time, it allows IT organizations to determine if the changes cause unacceptable impact on the company’s other key applications.

Together, CBQoS, MIB, WAAS Flow agent, IP SLA, and NAM can serve as useful tools for measurement, adjustment, and verification of WAN optimization initiatives.

4.1.5 Deploy Changes

The fourth step is to deploy changes. IT organizations regularly deploy new applications and updates to existing applications to meet changing business needs. As new applications are deployed or changes are made, new baselines need to be established. The application optimization cycle must start all over again.

4.2 Monitoring Instrumentation Overview

Continuous performance monitoring is key to optimized application performance. Whether traffic is generated synthetically and metrics from an end host generating and receiving traffic is monitored actively, or natural network traffic is monitored passively but with lower network overhead, network and application performance data can be retrieved from a wide variety of data sources, each offering a different level of granularity and relative value. The subsequent subsections provide detail description of key monitoring instrumentation.

As networks grow in size and complexity and enterprise requirements grow, a need for greater visibility arises. IT directors and managers need tools that can help identify the various segments of their network that need improvement to allow a more efficient distribution of limited budget resources. Cisco products come packaged with tools that provide the platform to build detailed network monitoring abilities.

4.3 IOS Instrumentation

This section describes monitoring information built into IOS, such as:

- Cisco IP service level agreements (IP SLA)

- Cisco NetFlow
- NBAR
- CBQoS MIB

4.3.1 IP SLA

IP SLA is a feature set in Cisco IOS software that enables users to analyze service levels for IP applications and services. IP SLA uses reliable, scheduled continuous traffic generation to measure network performance. IP SLAs can perform network assessments, verify quality of service (QoS), ease the deployment of new services, and assist administrators with network troubleshooting.

Important IP SLA highlights include:

- Monitoring network performance:
  - Ability to measure jitter, packet loss, packet ordering, packet corruption and delay
- Network availability monitoring:
  - Test connectivity of network resources
- Network troubleshooting:
  - Troubleshoot network elements through consistent and reliable measurement

IP SLA has two key components: a source device that generates, receives, and analyzes traffic, and the target device for which SLA measurements are gathered. Additional accuracy and detail for the measurements can be achieved using the optional IP SLA Responder function on the target device. The IP SLA responder enables the target device to mark the arrival and departure times of SLA probes, so that any local processing latency on the responder is mitigated. For example, with regular ICMP echo and echo reply, the echo target can choose to process ICMP traffic in a slow, deprioritized path. Without the SLA responder-associated special arrival and departure timestamps, the additional latency added by this slow path would be indistinguishable from actual network latency

4.3.1.1 IP SLA Network Management Support

IP SLA, described in detail in Chapter 8, can be managed by third party tools such as NetVoyant from NetQoS. IP SLA has a very strong SNMP-based configuration and data collection interface, and NetVoyant offers an easy GUI for managing Simple Network Management Protocol (SNMP) devices using a central console, rather than managing each device individually. The MIB browser in the NetVoyant console supports direct access to the MIB tables of a device.

4.3.1.2 IP SLA Operations

There are several key IP SLA operations:

- Internet Control Message Protocol (ICMP) echo
- User Datagram Protocol (UDP) echo
- Domain Name System (DNS) request
- Hypertext Transfer Protocol (HTTP) requests

4.3.1.3 IP SLA Configuration

This section provides configuration examples.

4.3.1.3.1 General Configuration Commands:

```
Router(config)#:ip sla <operation number>
```

Begin configuration for an ip sla operation and enter IP SLA monitor mode.

```
Router(config)#:ip sla monitor schedule <operation number> <start-time><age out>
<recurrence>
```

Configure the scheduling parameters for an individual IP SLA. This command must be run before an IP SLA will begin.

4.3.1.3.2 General Show Commands

```
Router#sh ip sla configuration <operation number>
```

This example shows the configuration parameters set for the current IP SLA by the specified operation number.

Example

```
Router#sh ip sla configuration 1
IP SLAs, Infrastructure Engine-II.
Entry number: 3
Owner: ICMP Echo - 100.1.1.161 - 60.1.1.100
Tag: WANOPT ICMP ECHO
Type of operation to perform: icmp-echo
Target address/Source address: 60.1.1.100/0.0.0.0
Operation timeout (milliseconds): 5000
Type Of Service parameters: 0x0
Vrf Name:
Request size (ARR data portion): 28
Verify data: No
Schedule:
   Operation frequency (seconds): 60  (not considered if randomly scheduled)
   Next Scheduled Start Time: Start Time already passed
   Group Scheduled : FALSE
   Randomly Scheduled : FALSE
   Life (seconds): Forever
   Entry Ageout (seconds): 3600
   Recurring (Starting Everyday): FALSE
   Status of entry (SNMP RowStatus): Active
Threshold (milliseconds): 5000
Distribution Statistics:
   Number of statistic hours kept: 2
   Number of statistic distribution buckets kept: 1
   Statistic distribution interval (milliseconds): 4294967295
History Statistics:
   Number of history Lives kept: 0
   Number of history Buckets kept: 15
   History Filter Type: None
Enhanced History:
Router#sh ip sla statistics 1
```

This command shows basic statistics gathered by the IP SLA specified.

Example

```
Router#sh ip sla statistics 1
Round Trip Time (RTT) for       Index 1
        Latest RTT: 60 milliseconds
Latest operation start time: 15:18:20.255 EST Tue Dec 4 2007
Latest operation return code: OK
Number of successes: 58
Number of failures: 0
Operation time to live: Forever
```

4.3.1.3.3 Configuring ICMP Echo

```
Router#: conf t
Router(config)#:ip sla 1
Router(config-ip-sla)#: icmp-echo 52.1.1.100
Router(config-ip-sla)#ip sla schedule 1 start-time now life forever
Router(config-ip-sla)#:end
Router#:wr
```

4.3.1.3.4 Configuring UDP Echo

```
Router#: conf t
Router(config)#:ip sla 1
Router(config-ip-sla)#:udp-echo 52.1.1.100 443
Router(config-ip-sla)#ip sla schedule 1 start-time now life forever
Router(config-ip-sla)#:end
Router#:wr
```

4.3.1.3.5 Configuring DNS Echo:

```
Router#: conf t
Router(config)#:ip sla 1
Router(config-ip-sla)#:dns www.cisco.com name-server 52.1.1.100
Router(config-ip-sla)#ip sla schedule 1 start-time now life forever
Router(config-ip-sla)#:end
Router#:wr
```

4.3.1.3.6 Configuring HTTP Echo:

```
Setup:
Router#: conf t
Router(config)#:ip name-server 52.1.1.100
Router(config)#:ip domain-list wanopt4.cisco.com
Router(config)#:ip domain-name wanopt4.cisco.com
Router(config)#:exit
Router#:wr
Router# conf t
Router(config)# ip sla 1
Router(config-ip-sla)#http get url http://www.cisco.com
Router(config-ip-sla)#ip sla schedule 1 start-time now life forever
```

4.3.1.4 ICMP Echo

ICMP is usually the first tool used in network troubleshooting to verify connectivity between two points on the network. With ICMP echo, several ICMP echo packets are sent to a destination, which then responds with the ICMP echo-replies. A bidirectional check such as ICMP echo can quickly verify connectivity to the target device.

ICMP echo operation can also monitor end-to-end response time between a Cisco router and a network resource or IP host. To compute response time, the time between sending an ICMP echo request and receiving an ICMP echo reply is measured. Only complete transaction round-trip time (RTT) is measured.

4.3.1.5 UDP Echo

UDP echo can determine round-trip delay times for UDP packets and test connectivity to both Cisco and non-Cisco devices, which can be very useful in troubleshooting certain business-critical applications. UDP echo offers more detailed reporting than ICMP operations, such as one-way delay measurement, but only when used with the IP SLA responder.

4.3.1.6 DNS Request

DNS is commonly used to translate hostnames into IP addresses, and to translate IP addresses into hostnames. DNS request measures the amount of time it takes to send a DNS request to a DNS server and receive a response to the request. This request can contain either an IP address or hostname depending on which is specified when the SLA is setup. DNS operations are a critical element for determining a network's overall performance as most IP services depend heavily on DNS name resolution.

4.3.1.7 HTTP Operation

HTTP operation centers around monitoring the response time between the source device and the HTTP server. Three values are measured to calculate response time:

1. DNS lookup: Round Trip Time (RTT) of a DNS lookup
2. TCP Connect: RTT of a TCP connect to the HTTP server
3. HTTP Transaction Time: RTT taken from request to response from the HTTP server

The HTTP SLA has two requests that can be configured: HTTP Get and HTTP RAW. For HTTP Get requests, the IP SLA formats the request based on the specified URL. For RAW requests, the entire content of the HTTP request must be specified. This allows RAW requests control over fields such as authentication.

4.3.2 NetFlow

Cisco IOS NetFlow is an integral technology in IOS network statistics gathering. NetFlow collects packets, maps them into flows, and counts the collected flow statistics as the packets enter and exit an interface. These flow statistics can then be exported to a NetFlow collector for storage and analysis.

The key components of NetFlow are the cache stores that hold IP flow information and an export mechanism that can send NetFlow data to a remote collector such as the Cisco NetFlow Collector. NetFlow operates by creating a NetFlow cache entry for each active flow and maintains a separate flow record within the cache for active flows. Each of these flow records contain multiple data fields which themselves are exported to the NetFlow Collector.

Figure 4-2. NetFlow Collector

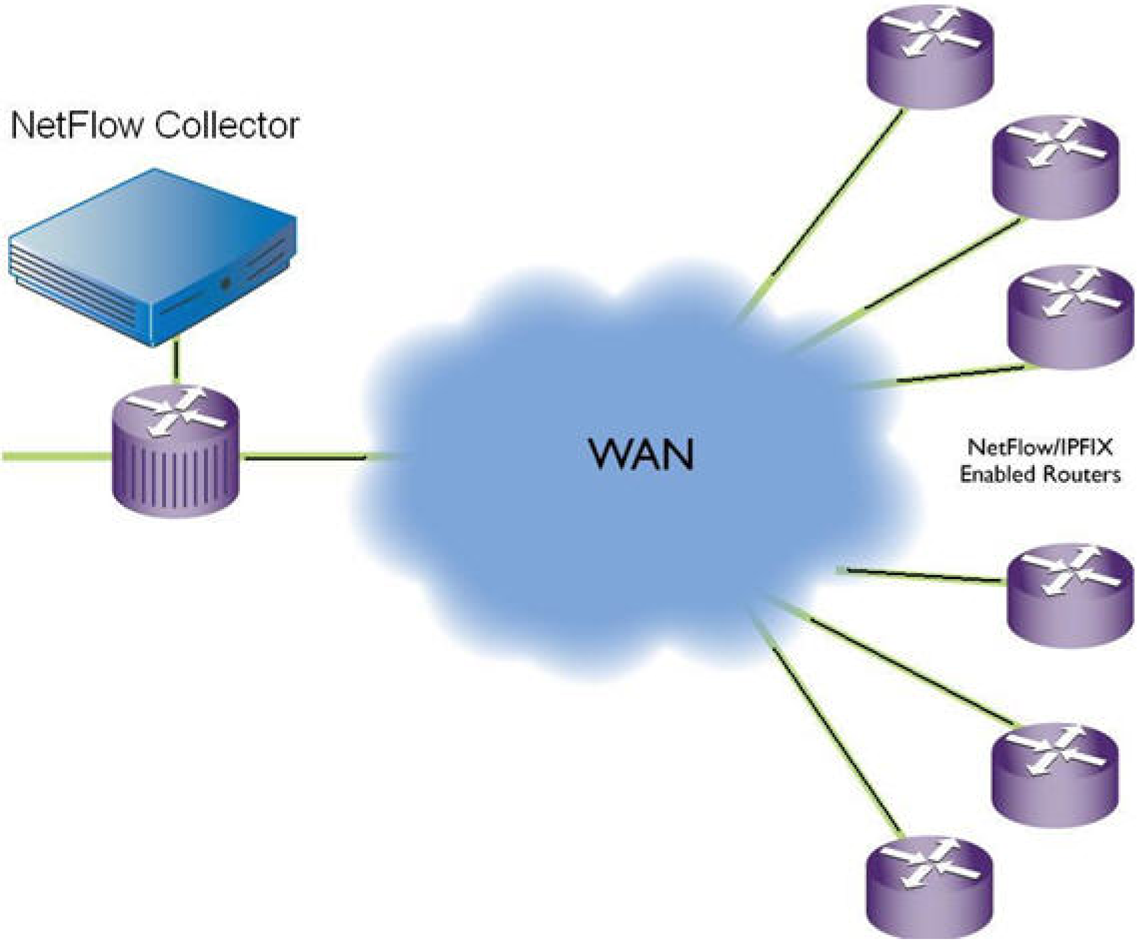

NetFlow identifies packet flows for IP Packets by looking at a number of fields in the data packet. A flow is defined as a set of packets having common properties. NetFlow defines a flow as the combination of the following seven key-fields, which determine how a flow is identified:

1. Source IP Address
2. Destination IP Address
3. Source port number
4. Destination port number
5. Layer 3 protocol type (e.g., ICMP, TCP, UDP)
6. ToS byte
7. Logical input interface (ifIndex)

Each flow record is created by grouping packets with the same characteristics into a flow. This method of determining a flow is ideal because a large amount of network information is condensed into a database of NetFlow information called the NetFlow cache. If any of these fields are different from another flow, it is considered a different flow. NetFlow operates by creating a NetFlow cache entry that contains information for each active flow.

Figure 4-3. NetFlow Cache Entry

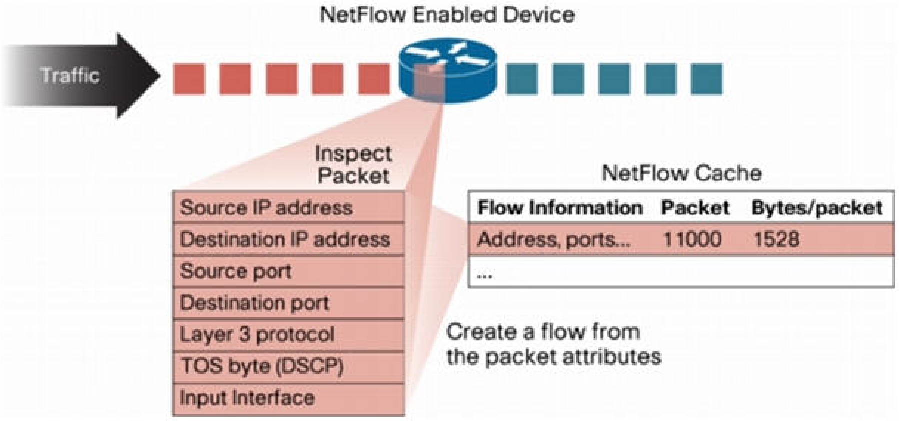

Flow Information	Packet	Bytes/packet
Address, ports...	11000	1528
...		

4.3.2.1 NetFlow Cache

The attributes of active flows can be analyzed by displaying the NetFlow Cache. This makes NetFlow a powerful troubleshooting tool, even without flow exporting.

4.3.2.2 Show Command

```
Router#: sh ip cache flow
```

Figure 4-4. NetFlow Cache Entries

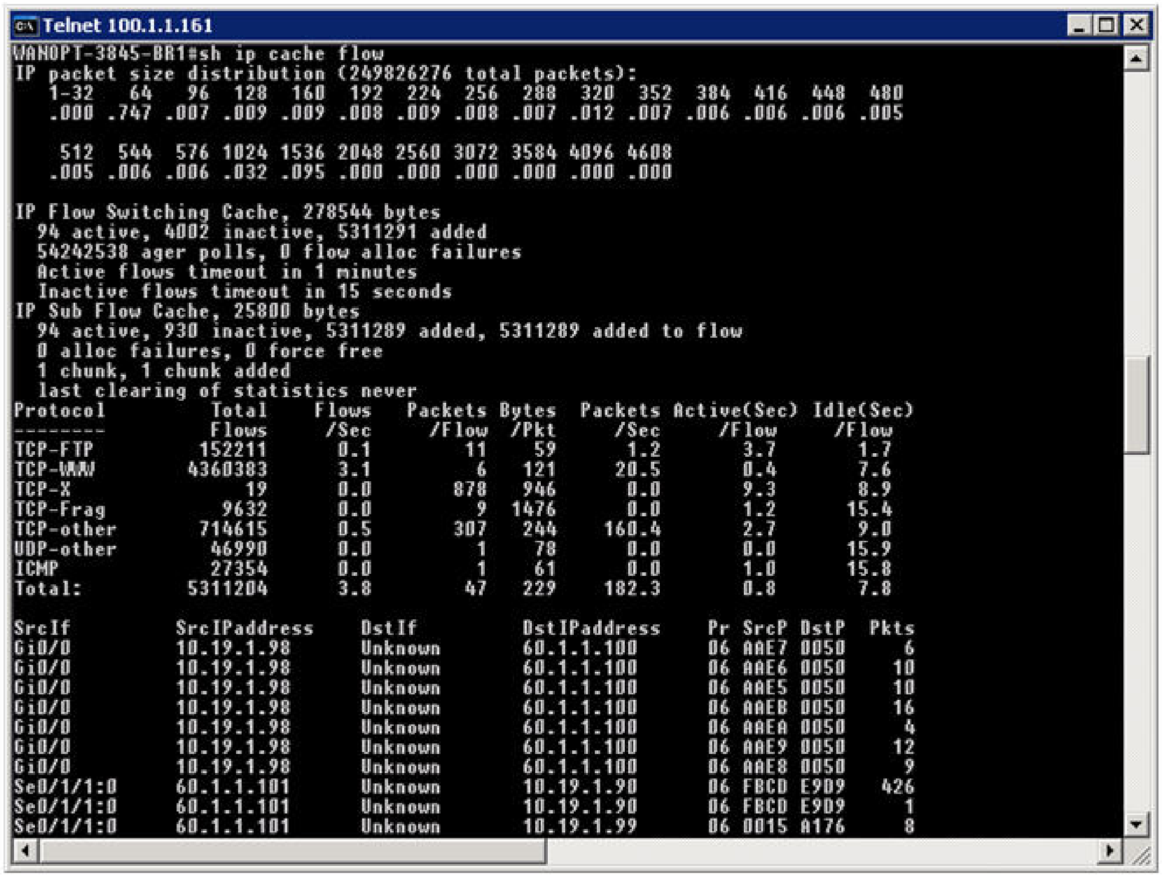

```
WANOPT-3845-BR1#sh ip cache flow
IP packet size distribution (249826276 total packets):
   1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
   .000 .747 .007 .009 .009 .008 .009 .008 .007 .012 .007 .006 .006 .006 .005

    512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
   .005 .006 .006 .032 .095 .000 .000 .000 .000 .000 .000

IP Flow Switching Cache, 278544 bytes
  94 active, 4002 inactive, 5311291 added
  54242538 ager polls, 0 flow alloc failures
  Active flows timeout in 1 minutes
  Inactive flows timeout in 15 seconds
IP Sub Flow Cache, 25800 bytes
  94 active, 930 inactive, 5311289 added, 5311289 added to flow
  0 alloc failures, 0 force free
  1 chunk, 1 chunk added
  last clearing of statistics never
Protocol         Total    Flows   Packets Bytes  Packets Active(Sec) Idle(Sec)
--------         Flows     /Sec     /Flow  /Pkt     /Sec     /Flow     /Flow
TCP-FTP         152211      0.1        11    59      1.2       3.7       1.7
TCP-WWW        4360383      3.1         6   121     20.5       0.4       7.6
TCP-X               19      0.0       878   946      0.0       9.3       8.9
TCP-Frag          9632      0.0         9  1476      0.0       1.2      15.4
TCP-other       714615      0.5       307   244    160.4       2.7       9.0
UDP-other        46990      0.0         1    78      0.0       0.0      15.9
ICMP             27354      0.0         1    61      0.0       1.0      15.8
Total:         5311204      3.8        47   229    182.3       0.8       7.8

SrcIf         SrcIPaddress    DstIf          DstIPaddress    Pr SrcP DstP  Pkts
Gi0/0         10.19.1.98      Unknown        60.1.1.100      06 AAE7 0050     6
Gi0/0         10.19.1.98      Unknown        60.1.1.100      06 AAE6 0050    10
Gi0/0         10.19.1.98      Unknown        60.1.1.100      06 AAE5 0050    10
Gi0/0         10.19.1.98      Unknown        60.1.1.100      06 AAEB 0050    16
Gi0/0         10.19.1.98      Unknown        60.1.1.100      06 AAEA 0050     4
Gi0/0         10.19.1.98      Unknown        60.1.1.100      06 AAE9 0050    12
Gi0/0         10.19.1.98      Unknown        60.1.1.100      06 AAE8 0050     9
Se0/1/1:0     60.1.1.101      Unknown        10.19.1.90      06 FBCD E9D9   426
Se0/1/1:0     60.1.1.101      Unknown        10.19.1.90      06 FBCD E9D9     1
Se0/1/1:0     60.1.1.101      Unknown        10.19.1.99      06 0015 A176     8
```

The various segments break down to:

- Packet size distribution
- General statistics about the state of the NetFlow cache
- The time a particular flow remains active in the cache before it is discarded.
- Flow breakdown by some well known protocols
- Actual NetFlow Cache entries

4.3.2.3 Aging Flows

On the NetFlow accounting device, the rules for expiring flow records and exporting them from cache entries to a flow collector are the following:

- Inactive/Active Timer: Flows that have been idle for a specified time are expired and removed. The default setting for this timer is fifteen seconds of traffic inactivity and can be configured between 10 and 600 seconds. On the other hand, long-lived flows are also expired and removed from the cache based on a different timer, called the active timer. The cutoff time for active flow expiration is thirty minutes and can be configured between 1 and 60 minutes.

- Full Cache: If a cache approaches full, emergency expiration will occur. The cache size can be configured by the network operator.
- End of a TCP connection: TCP connections at the end of a byte stream (FIN) or have been reset (RST) automatically expire on software platforms.

4.3.2.4 NetFlow Versions

4.3.2.4.1 NetFlow Version 5

NetFlow Version 5, generally considered the most popular NetFlow version, includes BGP Autonomous System information and flow sequence numbers.

Figure 4-5. Typical NetFlow Export Datagram Format for Versions 1, 5, 7, and 8

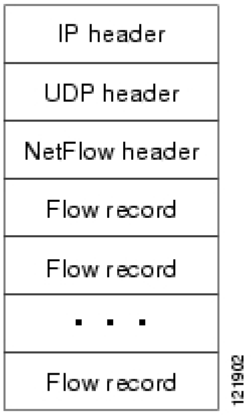

4.3.2.4.2 Configuration Commands

```
Router(config)#: ip flow-export source loopback 0
```

Set flow source destination as local device

```
Router(config)#: ip flow-export version 5 peer-as
```

Set NetFlow export version

```
Router(config)#: ip flow-export destination 52.1.1.22 9995
```

Specify the NetFlow collector for exported records.

4.3.2.4.3 Show Commands:

```
Router#: sh ip flow export
```

This command displays the NetFlow Version 5 configurations as well as other interesting statistics, such as the number of exported flow records, the number of exported packets, the number of packets that were not exported, and the reason for failures.

Figure 4-6. IP Flow Export Statistics

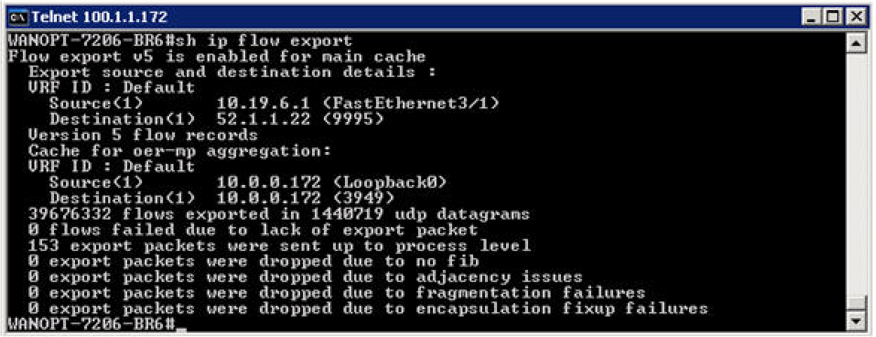

```
WANOPT-7206-BR6#sh ip flow export
Flow export v5 is enabled for main cache
  Export source and destination details :
  VRF ID : Default
    Source(1)       10.19.6.1 (FastEthernet3/1)
    Destination(1)  52.1.1.22 (9995)
  Version 5 flow records
  Cache for oer-mp aggregation:
  VRF ID : Default
    Source(1)       10.0.0.172 (Loopback0)
    Destination(1)  10.0.0.172 (3949)
  39676332 flows exported in 1440719 udp datagrams
  0 flows failed due to lack of export packet
  153 export packets were sent up to process level
  0 export packets were dropped due to no fib
  0 export packets were dropped due to adjacency issues
  0 export packets were dropped due to fragmentation failures
  0 export packets were dropped due to encapsulation fixup failures
WANOPT-7206-BR6#
```

4.3.2.4.4 NetFlow Version 9

The distinguishing feature of the NetFlow Version 9 format is that it is *template based*. Templates provide an extensible design to the record format, a feature that should allow future enhancements to NetFlow services without requiring concurrent changes to the basic flow-record format and collector code. Version 9 also incorporates new features such as multicast, MPLS, BGP next hop, and IPv6. Using templates with NetFlow Version 9 provides several key benefits:

- Third-party business partners who produce applications that provide collector or display services for NetFlow will not be required to recompile their applications each time a new NetFlow export field is added. Instead, they may be able to use an external data file that documents the known template formats.
- New features can be added to NetFlow more quickly, without breaking current implementations.
- NetFlow is considered "future-proofed" against new or developing protocols, because the Version 9 format can be adapted to provide support for them and other non-Flow based data measurements.

In the following NetFlow version 9 flow captured using WireSharks, Figure 4-7 depicts the NetFlow v9 flow template that identifies the fields that will be present in the actual flow record while Figure 4-8 depicts the actual flow record.

Figure 4-7. NetFlow version 9 Flow Template

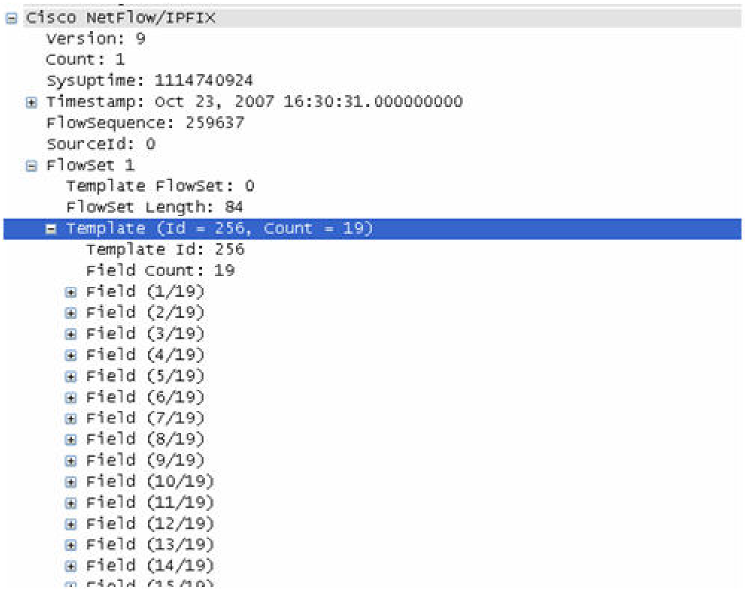

```
⊟ Cisco NetFlow/IPFIX
    Version: 9
    Count: 1
    SysUptime: 1114740924
  ⊞ Timestamp: Oct 23, 2007 16:30:31.000000000
    FlowSequence: 259637
    SourceId: 0
  ⊟ FlowSet 1
      Template FlowSet: 0
      FlowSet Length: 84
    ⊟ Template (Id = 256, Count = 19)
        Template Id: 256
        Field Count: 19
      ⊞ Field (1/19)
      ⊞ Field (2/19)
      ⊞ Field (3/19)
      ⊞ Field (4/19)
      ⊞ Field (5/19)
      ⊞ Field (6/19)
      ⊞ Field (7/19)
      ⊞ Field (8/19)
      ⊞ Field (9/19)
      ⊞ Field (10/19)
      ⊞ Field (11/19)
      ⊞ Field (12/19)
      ⊞ Field (13/19)
      ⊞ Field (14/19)
      ⊞ Field (15/19)
```

Highlighted here is the flow template that shows the template ID for this flow record and the number and type of fields included in the record.

Figure 4-8. NetFlow version 9 Flow Record

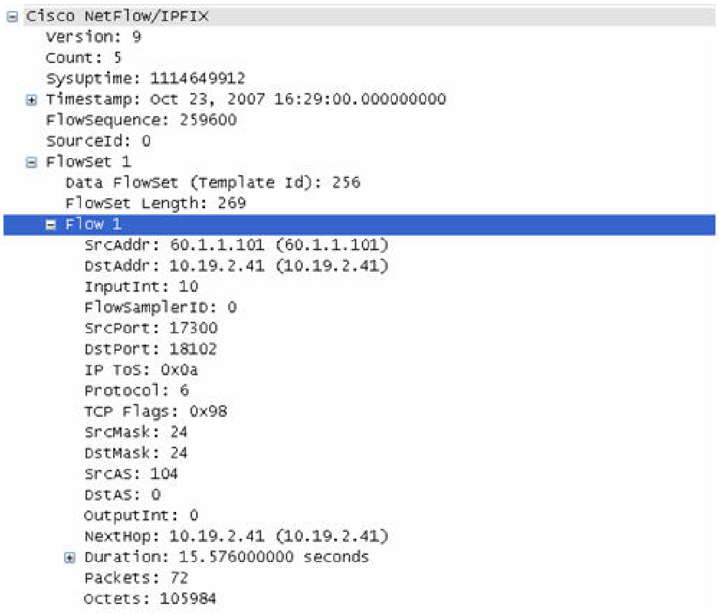

```
⊟ Cisco NetFlow/IPFIX
    Version: 9
    Count: 5
    SysUptime: 1114649912
  ⊞ Timestamp: Oct 23, 2007 16:29:00.000000000
    FlowSequence: 259600
    SourceId: 0
  ⊟ FlowSet 1
      Data FlowSet (Template Id): 256
      FlowSet Length: 269
    ⊟ Flow 1
        SrcAddr: 60.1.1.101 (60.1.1.101)
        DstAddr: 10.19.2.41 (10.19.2.41)
        InputInt: 10
        FlowSamplerID: 0
        SrcPort: 17300
        DstPort: 18102
        IP ToS: 0x0a
        Protocol: 6
        TCP Flags: 0x98
        SrcMask: 24
        DstMask: 24
        SrcAS: 104
        DstAS: 0
        OutputInt: 0
        NextHop: 10.19.2.41 (10.19.2.41)
      ⊞ Duration: 15.576000000 seconds
        Packets: 72
        Octets: 105984
```

A more detailed view of how the NetFlow Version 9 flow template and flowsets match up is depicted in Figure 4-9.

Figure 4-9. NetFlow version 9 Flow Breakdown

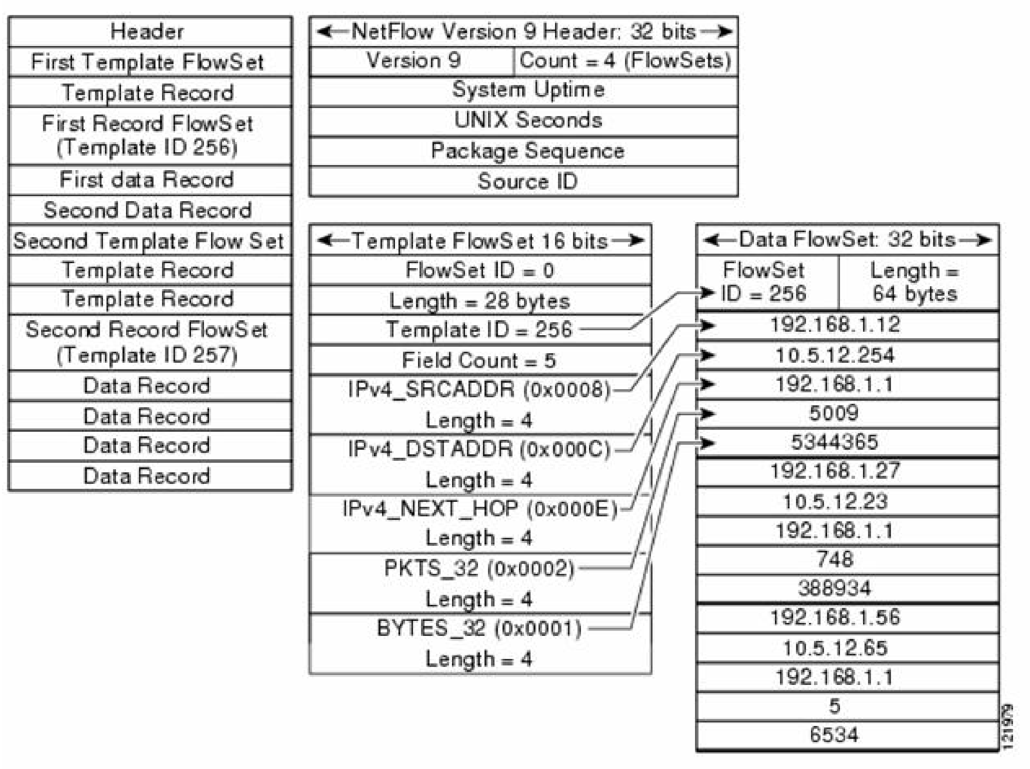

4.3.2.4.5 Configuration Command

```
Router#:conf t
Router(config)#: ip flow-export version 9
Router(config)#: ip flow-export destination  52.1.1.22 9995
Router(config)#: ip flow-export source loopback 0
```

4.3.2.4.6 Show Commands

```
Router#: show ip flow export
```

Same as netflow version 5

```
Router#: show ip flow export template
```

This command shows the number of templates and the number of active templates.

4.3.3 NBAR

Network Based Application Recognition (NBAR) provides network traffic classification. NBAR can recognize a very wide variety of applications by doing IP packet inspection up to OSI Layer 7. It can, for instance, differentiate between Web-based HTTP and Skype traffic, which can both use TCP port 80.

When an application is recognized, NBAR classifies the traffic for performance and accounting purposes. This function gives an operator the ability to invoke any range of services for that specific application, whether offering more or less bandwidth, latency queuing, or completely blocking certain packets.

NBAR also provides a special Protocol Discovery (PD) feature that determines which applications and protocols are traversing the network at any given time. PD captures key statistics that are associated with each protocol based on IP flows. Like Cisco NetFlow, NBAR defines IP flows as a unidirectional flow of IP packets that share the following five values:

- Source IP address
- Destination IP address
- Source port
- Destination port
- L3 protocol type

NetFlow and NBAR both leverage L3 and L4 header information. However, unlike NetFlow, NBAR also examines data from L3–L7. NBAR uses L3 and L4 and packet inspection for classification, and supports stateful inspection of dynamic-port traffic. NBAR also requires a set number of packets before making a protocol distinction.

4.3.3.1 NBAR Protocol Discovery

NBAR PD provides an easy way to discover application traffic flowing through an interface by providing a Protocol Discovery (PD), MIB, which gives it expanded capabilities through SNMP. This includes:

- Enable/Disable Protocol Discovery on a per interface basis
- Monitoring both ingress and egress traffic
- Display statistics on a per-protocol basis

Table 4-1 NBAR Protocol Discovery MIB Details

Table	Description
cnpdSupportedProtocols	List of all supported protocols NBAR supports
cnpdAllStats	All NBAR statistics per interface such as • Packet counters (inbound/outbound) • Byte counters (inbound/outbound) • Bit rate (inbound/outbound)
cnpdTopNStats	Top-N table statistics
cnpdThresholdhistory	History of falling or rising events
cnpdStatus	Enable or disable NBAR per interface, including time stamp
cnpdTopNConfig	Configure the Top-N table by interface
cnpdThresholdConfig	Protocol threshold configuration
cnpdNotificationsConfig	Enable traps
cnpdMIBNotifications	Rising or falling events

4.3.3.2 Configuration

4.3.3.2.1 Configuring Protocol Discovery on a Specific Interface

```
router#: conf t
router(config)#: interface gigabitEthernet 0/1
```

This configures NBAR to discover traffic and keep traffic statistics for all protocols known to NBAR on a particular interface

```
router(config-if)#:ip nbar protocol-discovery
router(config)#: exit
```

Optional

```
router(config)#: ip nbar port-map
```

This configures NBAR to search for a protocol or protocol name using port number(s) other than an already known port. Up to 16 different port numbers can represent a protocol.

```
router(config)#:ip nbar custom protocol-name [destination | source] [tcp|udp]
```

Configures NBAR to classify and monitor additional static port applications. The parameters are the following:

- protocol-name: Specifies the name of the user defined protocol.
- number: The byte location of the value to be searched in the payload (0 to 255)(optional).
- destination: Inspects destination flows only (optional)
- tcp: Defines up to 16 explicit TCP port numbers or a range of a maximum of 1000 TCP ports
- udp: Defines up to 16 explicit UDP port numbers or a range of a maximum of 1000 UDP ports.

```
router(config)#:ip nbar pdlm
```

Extends the list of protocols by loading a new PDLM (providing the full path to the PDLM) New PDLM versions are provided on the Cisco website at http://www.cisco.com/go/nbar

4.3.3.3 Show Commands

```
router#: show ip nbar protocol-discovery
```

This command displays the statistics gathered by the NBAR Protocol Discovery feature. By default, all statistics for all interfaces are displayed.

Note: Egress Traffic statistics are gathered before policing features.

Figure 4-10. Sample Output from PD Show Command

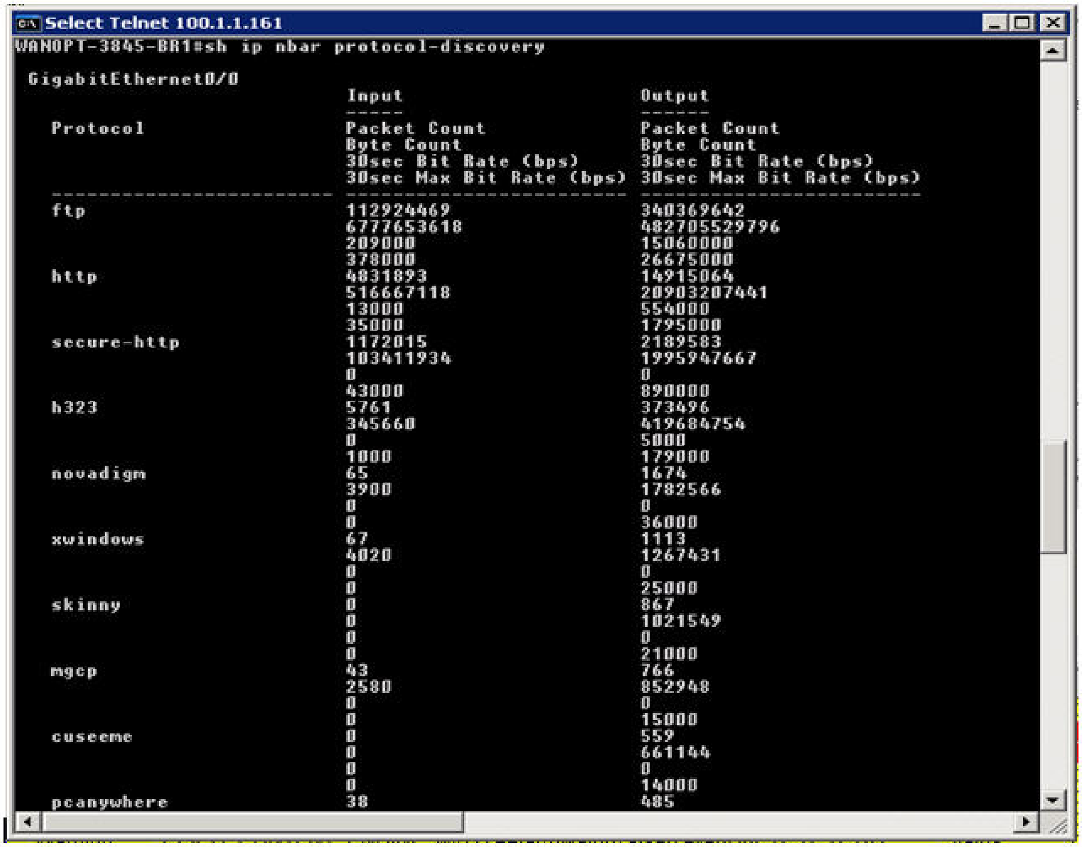

```
WANOPT-3845-BR1#sh ip nbar protocol-discovery

 GigabitEthernet0/0
                            Input                    Output
                            -----                    ------
   Protocol                 Packet Count             Packet Count
                            Byte Count               Byte Count
                            30sec Bit Rate (bps)     30sec Bit Rate (bps)
                            30sec Max Bit Rate (bps) 30sec Max Bit Rate (bps)
   ------------------------ ------------------------ ------------------------
   ftp                      112924469                340369642
                            6777653618               482705529796
                            209000                   15060000
                            378000                   26675000
   http                     4831893                  14915064
                            516667118                20903207441
                            13000                    554000
                            35000                    1795000
   secure-http              1172015                  2189583
                            103411934                1995947667
                            0                        0
                            43000                    890000
   h323                     5761                     373496
                            345660                   419684754
                            0                        5000
                            1000                     179000
   novadigm                 65                       1674
                            3900                     1782566
                            0                        0
                            0                        36000
   xwindows                 67                       1113
                            4020                     1267431
                            0                        0
                            0                        25000
   skinny                   0                        867
                            0                        1021549
                            0                        0
                            0                        21000
   mgcp                     43                       766
                            2580                     852948
                            0                        0
                            0                        15000
   cuseeme                  0                        559
                            0                        661144
                            0                        0
                            0                        14000
   pcanywhere               38                       485
```

```
Router#: show ip nbar protocol-discovery
```

interface stats protocol top-n

- interface-spec: specifies an interface to display
- stats: specifies the byte count, bit rate, or packet count is to be displayed
- protocol: specifies that statistics for a specific protocol.
- top-n: specifies that a Top-N of most active protocols is displayed

Figure 4-11. Sample Output from PD Interface Show Command

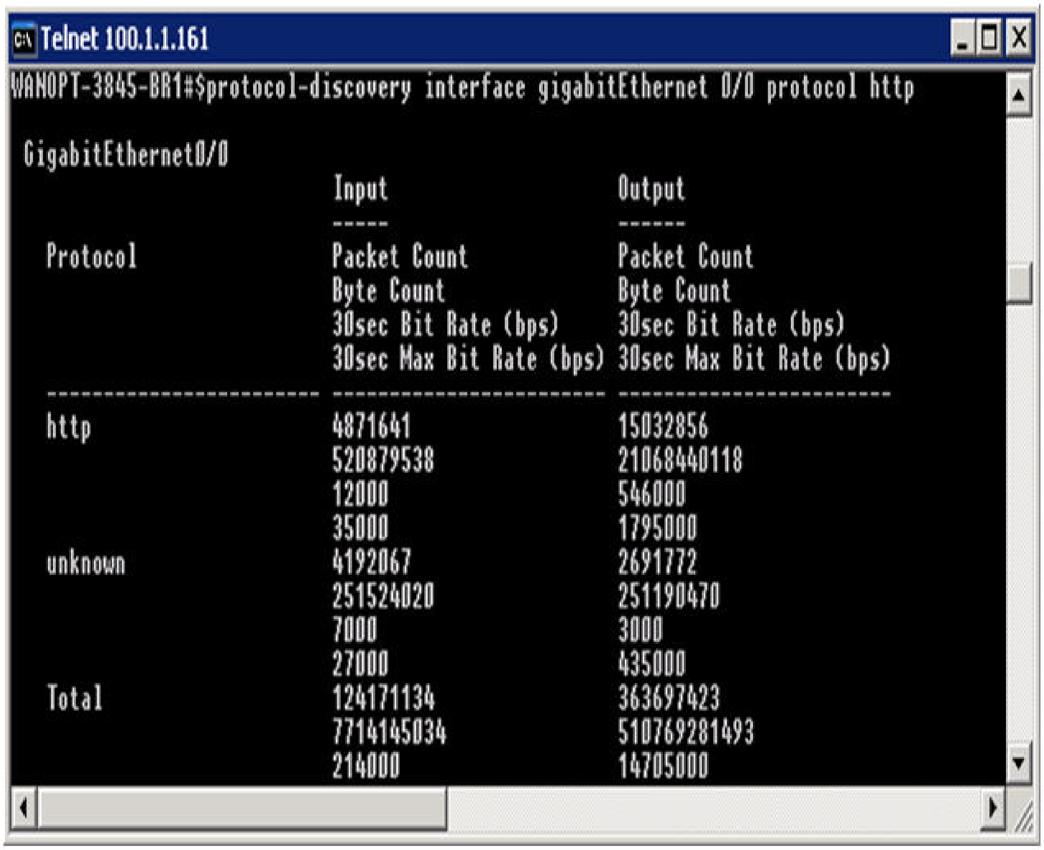

```
WANOPT-3845-BR1#$protocol-discovery interface gigabitEthernet 0/0 protocol http

 GigabitEthernet0/0
                          Input                   Output
                          -----                   ------
  Protocol                Packet Count            Packet Count
                          Byte Count              Byte Count
                          30sec Bit Rate (bps)    30sec Bit Rate (bps)
                          30sec Max Bit Rate (bps) 30sec Max Bit Rate (bps)
  ------------------------ ------------------------ ------------------------
  http                    4871641                 15032856
                          520879538               21068440118
                          12000                   546000
                          35000                   1795000
  unknown                 4192067                 2691772
                          251524020               251190470
                          7000                    3000
                          27000                   435000
  Total                   124171134               363697423
                          7714145034              510769281493
                          214000                  14705000
```

4.3.4 CBQoS MIB

The Cisco Class-Based QoS (CBQoS) MIB supplies QoS information for Cisco network elements that support the Modular QoS command-line interface (MQC). CBQoS provides configuration capabilities and monitoring statistics that include summary counts and rates by traffic class before and after the enforcement of QoS policies. It also provides detailed feature-specific statistics that are available for select PolicyMap features. Policy actions are defined per interface and traffic direction, whether ingress or egress. The CBQoS MIB supports both 32 bit and 64 bit counters.

The following is a list of relevant MIB tables for QoS and contain only statistical information.

- **cbQosClassMapStats**- Statistical information about class maps, such as pre/post-policy packet/byte counts, bit rates, drop packet/bytes and no-buffer drops.
- **cbQosMatchStmtStats-** Statistical information about match statement-specific information, such as prepolicy packet/byte
- **cbQosPoliceStats-** Statistical information about police actions, such as conformed or exceeded packet/byte counters and bit rates.
- **cbQueueingStats**- Statistical information about queuing actions, such as the various queue depth and discar packet/byte counters.
- **cbQosTSStats**- Statistical information about traffic-shaping actions, such as various delay and drop packet/byte counters, state of feature, and queue size.

- **cbQosREDClassStats-** Statistical information about per-precedence weighted random early detection actions, such as random packet/byte counters and tail drop packet/byte counters.
- **cbQosPoliceActionCfg-** Required objects to display class-based QoS objects' configuration information.

4.4 Additional Instrumentation

Additional instrumentation includes:

- Cisco WAAS Flow Agent
- Connection State and Operation Statistics Reports

4.4.1 Cisco WAAS Flow Agent

Cisco and NetQoS jointly developed monitoring instrumentation to allow accurate end-to-end response time measurements, from the client to the server, over links optimized by Cisco WAAS devices. Central to this instrumentation is a Cisco software feature called the FlowAgent, a flow monitoring module integrated in the WAE. The FlowAgent captures important packet information and sends it across the network to a third-party monitoring agent (e.g., NetQoS SuperAgent). It exports all data necessary to report application response times experienced at remote sites served by WAAS and the detailed performance metrics for each optimized link.

The FlowAgent is part of the standard software image for Cisco WAAS 4.0.13 and newer. When configured on a WAE, the FlowAgent collects relevant flow information for optimized TCP transactions and transmits it to the SuperAgent Aggregator, a device that is dedicated solely to FlowAgent data collection. The SuperAgent Aggregator processes the flow information and sends it to a SuperAgent Management Console for storage and reporting. FlowAgent data collection can be configured on the Cisco WAAS Central Manager selectively for a single WAE device, or for multiple (or all) WAE devices using device groups.

Figure 4-12. Cisco WAAS FlowAgent

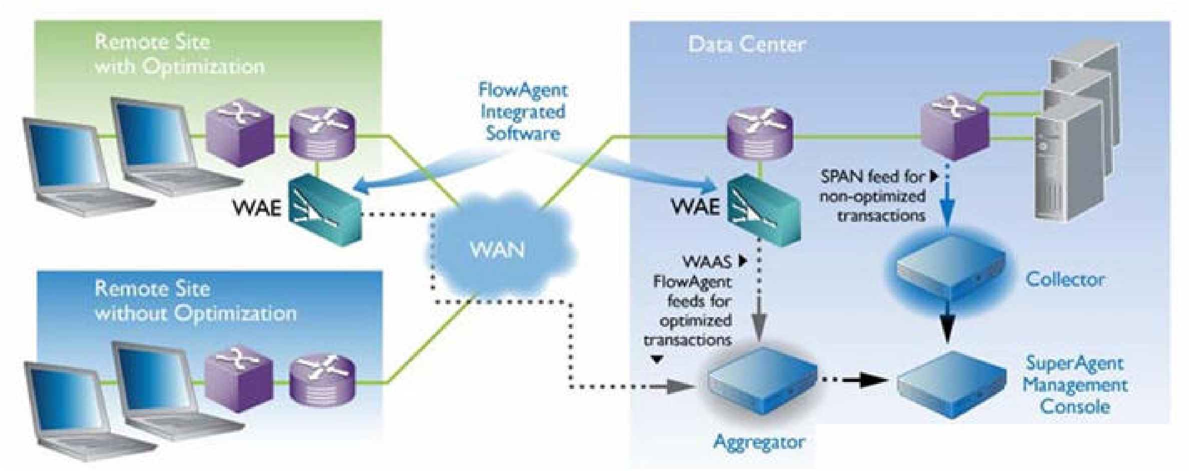
Remote Site with Optimization
WAE
FlowAgent Integrated Software
WAN
Remote Site without Optimization
Data Center
WAE
SPAN feed for non-optimized transactions
WAAS FlowAgent feeds for optimized transactions
Collector
SuperAgent Management Console
Aggregator

4.4.1.1.1 Sample Export of Flow Records for Optimized Traffic

The FlowAgent captures the following information on optimized traffic and sends it over to the configured SuperAgent Aggregator:

- Source IP address
- Destination IP address
- Source TCP port
- Destination TCP port
- TCP Sequence number
- TCP Acknowledgement number
- TCP payload byte count
- Packet arrival time in milliseconds
- IP identifier
- WAE MAC address Src/Dst flag
- TCP flags

4.4.1.2 FlowAgent Configuration

The SuperAgent FlowAgent monitoring agent is composed of two modules: the console (or host) and the collector. The console IP address is configured on the WAE through the WAE CLI or through the Central Manager GUI. The WAE initiates a temporary connection to the monitoring agent console. This temporary connection is referred to as the control connection. The control connection uses TCP port 7878, and its purpose is to obtain the IP address and port number of the collector to which the WAE is assigned. The WAE also pulls the configuration information regarding which servers are to be monitored over the control connection. Once the WAE obtains the IP address and port number information of the collector, the WAE opens a persistent connection to the collector. Collected summary data for the servers that are being monitored is sent over this persistent connection. The WAE send summary data only to Aggregators for servers assigned to it.

Configuration for flow monitoring with NetQoS involves the following tasks:

1. From the WAE CLI or Central Manager GUI, enter the SuperAgent Master Console IP address in the **tcpstat-v1** Host field on your WAE appliances. If you are configuring multiple appliances through a device group, wait for the configuration to propagate to all the appliances in the device list.

2. From the NetQoS SuperAgent console, assign a WAE to a SuperAgent Aggregator (known as the collector in WAAS terminology) and configure the NetQoS Networks, Servers, and Applications entities.

4.4.1.3 Enabling the FlowAgent on the WAE using the Central Manager GUI

To configure flow monitoring on your WAEs using the Central Manager GUI, follow these steps:

1. From the NetQoS SuperAgent console, assign a WAE to a SuperAgent Aggregator (known as the collector in WAAS terminology) and configure the NetQoS Networks, Servers, and Applications entities.

2. Create a new device group to be used for configuring flow monitoring on multiple devices. To create a device group, choose **Devices > Device Groups > Create New Device Group**.

   When you create the device group, check the auto assign all newly activated devices to this group check box to enable this option.

3. Add your existing WAE devices to this new device group.

4. From the Device Group listing window, click the **Edit** icon next to the name of the flow monitoring configuration device group that you want to configure.

5. In the Contents pane, choose **General Settings > Notification and Tracking > Flow Monitor**. The Flow Monitor Settings for Device Group window appears.

6. Check the **Enable** check box.

7. In the tcpstat-v1 Host field, enter the IP address of the monitoring agent console.

   This configuration allows the WAE to establish a temporary connection (a control connection) to the console for obtaining the IP address of the SuperAgent aggregator. You must configure the aggregator IP address information from the console device. (See the configuration documentation for your third-party flow monitoring application software.)

8. Click **Submit** to apply the settings to the devices in this device group.

4.4.1.4 Enabling the FlowAgent on the WAE using the CLI

To configure flow monitoring on your WAEs using the CLI, two commands are required (Figure 4-13. Enabling FlowAgent on the WAE):

1. Register the WAE with the IP address of the monitoring agent console by using the **flow monitor tcpstat-v1 host** global configuration command. The WAE needs to know the location of the SuperAgent Management Console

   ```
   #flow monitor tcpstat-v1 host <IP_Address>
   ```

   This configuration allows the WAE to establish a temporary connection (a control connection) to the console (or host) for obtaining the IP address of the collector. You must configure the collector IP address information from the console device. (See the configuration documentation for your third-party flow monitoring application software.)

   Note that the IP Address in this command is the SuperAgent Management Console. The SuperAgent Aggregator IP address is not needed for this configuration since the SuperAgent Aggregator to WAE configuration is accomplished in the SuperAgent Management Console UI.

2. Enable flow monitoring on the WAE appliance by using the **flow monitor tcpstat-v1 enable** global configuration command.

   ```
   #flow monitor tcpstat-v1 enable
   ```

   Note that the WAE requires the IP Address of the SuperAgent Management Console before using the enable command.

Figure 4-13. Enabling FlowAgent on the WAE

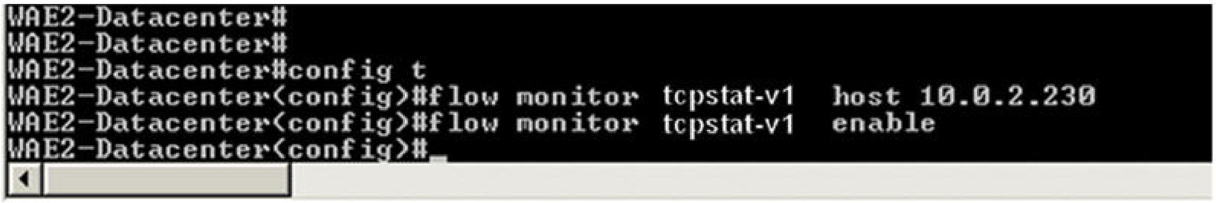

```
WAE2-Datacenter#
WAE2-Datacenter#
WAE2-Datacenter#config t
WAE2-Datacenter(config)#flow monitor  tcpstat-v1  host 10.0.2.230
WAE2-Datacenter(config)#flow monitor  tcpstat-v1  enable
WAE2-Datacenter(config)#_
```

4.4.2 Connection State and Operation Statistics Reports

The console (or host) module and the collector module may be on a single device or may be located on separate devices. These connections are independent of one another.

The state of these connections, as well as various operation statistics, is reported by the **show statistics flow monitor tcpstat-v1** EXEC mode command. This command is run from the WAE to determine the status of its connection to the SuperAgent system. As shown in Figure 4-14, this command provides the following information:

- Host Connection shows the IP of the SuperAgent Management Console. It usually shows "Not Connected" unless running it during a WAE to Console poll.

  Note: The connection between the SuperAgent Management Console and WAE is non-persistent. It will reconnect and synchronize every 5 minutes).

- Collector Connection should always show "Connected" after the IP:Port of the assigned SuperAgent Aggregator for that WAE.
- The assigned Aggregator is setup in the SuperAgent GUI. You will probably have one Aggregator even in a large network. It acts more like a Harvester than a Collector.

Figure 4-14. FlowAgent Connection Status

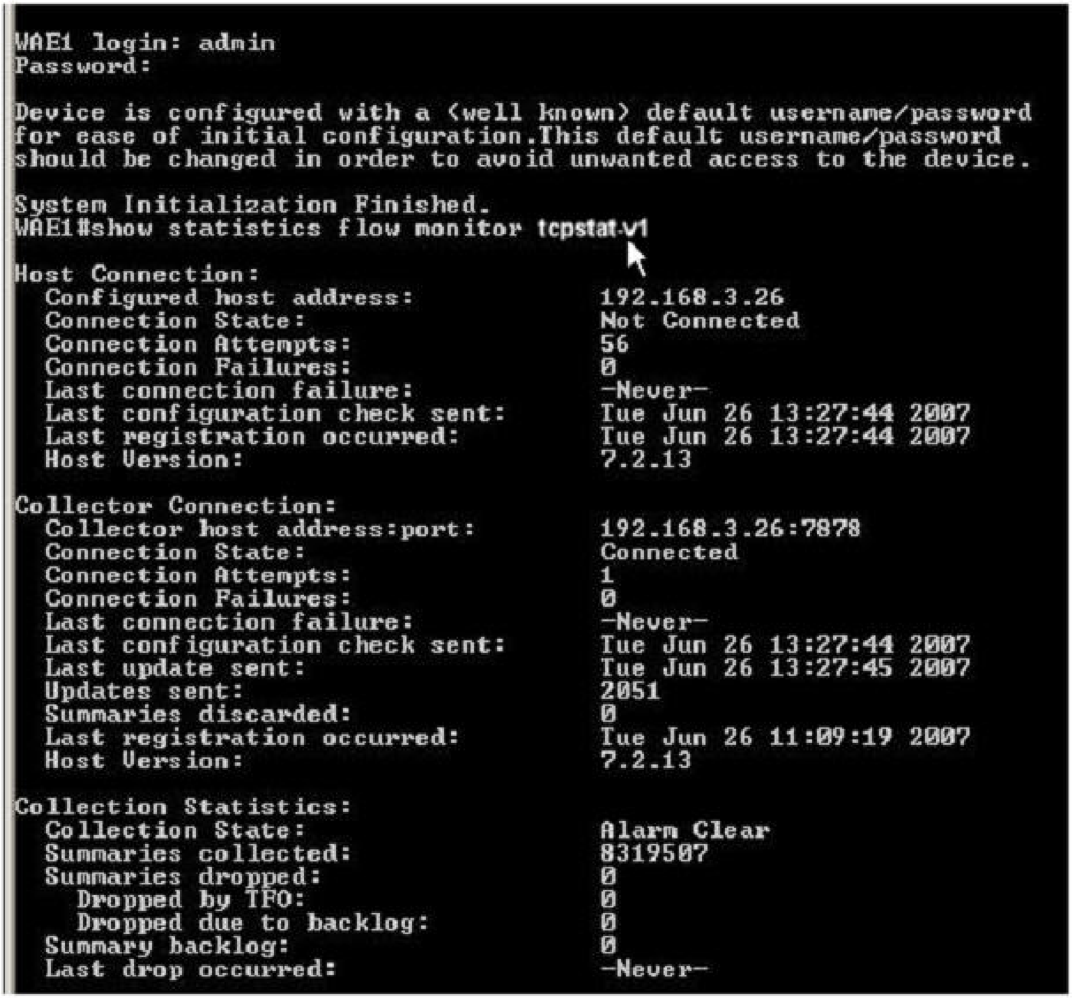

```
WAE1 login: admin
Password:

Device is configured with a (well known) default username/password
for ease of initial configuration.This default username/password
should be changed in order to avoid unwanted access to the device.

System Initialization Finished.
WAE1#show statistics flow monitor tcpstat-v1

Host Connection:
   Configured host address:            192.168.3.26
   Connection State:                   Not Connected
   Connection Attempts:                56
   Connection Failures:                0
   Last connection failure:            -Never-
   Last configuration check sent:      Tue Jun 26 13:27:44 2007
   Last registration occurred:         Tue Jun 26 13:27:44 2007
   Host Version:                       7.2.13

Collector Connection:
   Collector host address:port:        192.168.3.26:7878
   Connection State:                   Connected
   Connection Attempts:                1
   Connection Failures:                0
   Last connection failure:            -Never-
   Last configuration check sent:      Tue Jun 26 13:27:44 2007
   Last update sent:                   Tue Jun 26 13:27:45 2007
   Updates sent:                       2051
   Summaries discarded:                0
   Last registration occurred:         Tue Jun 26 11:09:19 2007
   Host Version:                       7.2.13

Collection Statistics:
   Collection State:                   Alarm Clear
   Summaries collected:                8319507
   Summaries dropped:                  0
     Dropped by TFO:                   0
     Dropped due to backlog:           0
   Summary backlog:                    0
   Last drop occurred:                 -Never-
```

Figure 4-15 illustrates a WAE that has not connected to the SuperAgent Management Console and has never received SuperAgent Aggregator assignment. An Alarm state will be raised in the State field.

Figure 4-15. FlowAgent connection status failure

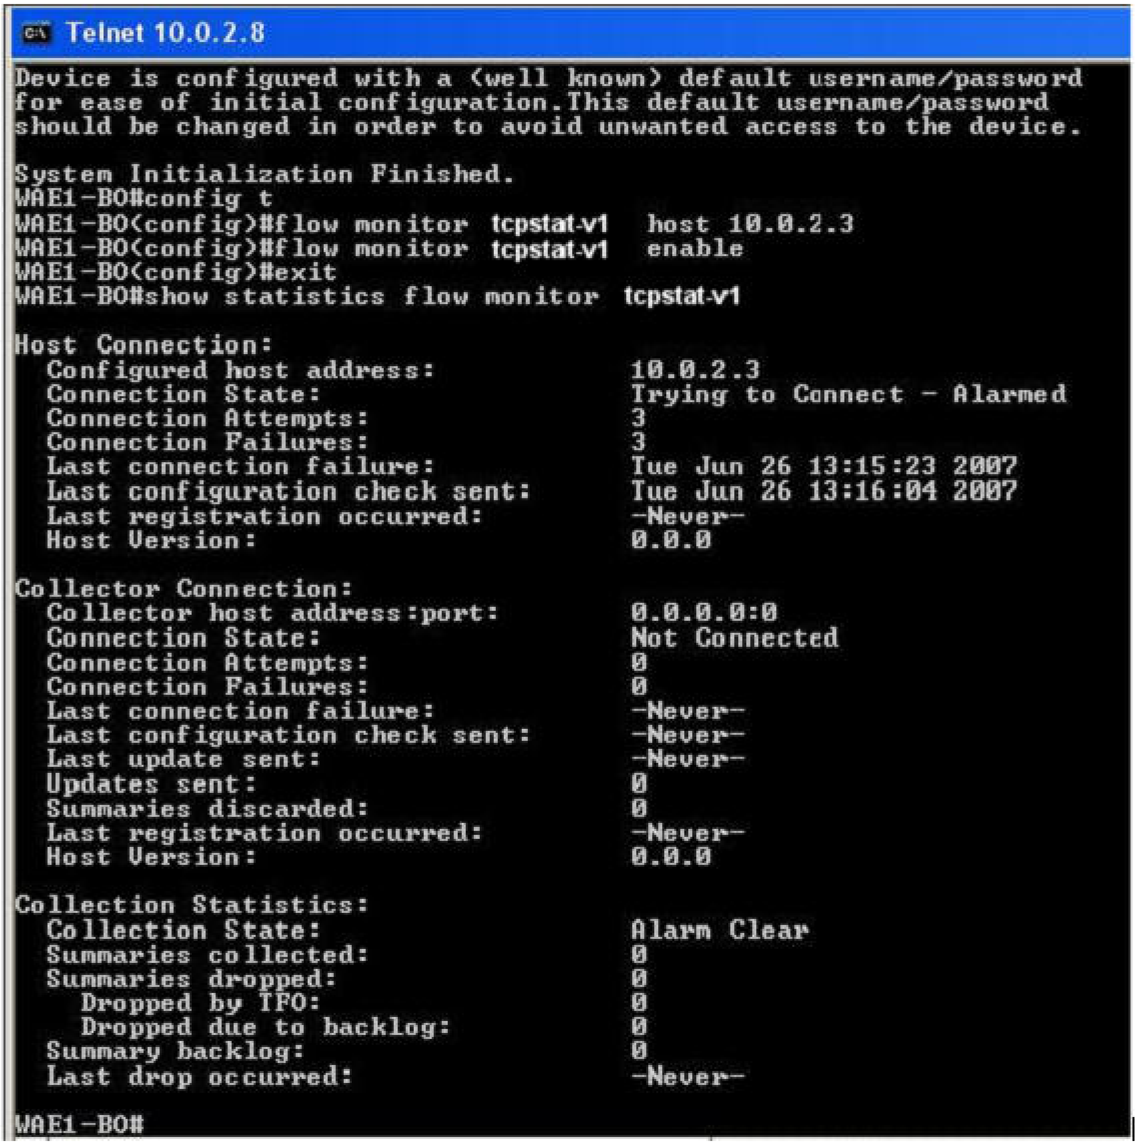

```
Device is configured with a (well known) default username/password
for ease of initial configuration.This default username/password
should be changed in order to avoid unwanted access to the device.

System Initialization Finished.
WAE1-BO#config t
WAE1-BO(config)#flow monitor tcpstat-v1  host 10.0.2.3
WAE1-BO(config)#flow monitor tcpstat-v1  enable
WAE1-BO(config)#exit
WAE1-BO#show statistics flow monitor tcpstat-v1

Host Connection:
  Configured host address:            10.0.2.3
  Connection State:                   Trying to Connect - Alarmed
  Connection Attempts:                3
  Connection Failures:                3
  Last connection failure:            Tue Jun 26 13:15:23 2007
  Last configuration check sent:      Tue Jun 26 13:16:04 2007
  Last registration occurred:         -Never-
  Host Version:                       0.0.0

Collector Connection:
  Collector host address:port:        0.0.0.0:0
  Connection State:                   Not Connected
  Connection Attempts:                0
  Connection Failures:                0
  Last connection failure:            -Never-
  Last configuration check sent:      -Never-
  Last update sent:                   -Never-
  Updates sent:                       0
  Summaries discarded:                0
  Last registration occurred:         -Never-
  Host Version:                       0.0.0

Collection Statistics:
  Collection State:                   Alarm Clear
  Summaries collected:                0
  Summaries dropped:                  0
    Dropped by TFO:                   0
    Dropped due to backlog:           0
  Summary backlog:                    0
  Last drop occurred:                 -Never-

WAE1-BO#
```

Connection errors and data transfer errors raise alarms on the WAE and in the Central Manager GUI. For debug information, use the **debug flow monitor tcpstat-v1** EXEC mode command.

The `show statistics flow filters` command identifies which filters were built based on the server assignment from the SuperAgent Management Console. In Figure 4-16, only the 10.0.3.6 server traffic is being optimized as shown in the “flow hits” column. Note that the server list comes from the SuperAgent Management Console automatically. This is the best way to validate the SuperAgent Management Console configuration against the WAE if data does not appear in the graphs.

Figure 4-16. Identifying Built Filters from the SuperAgent Management Console

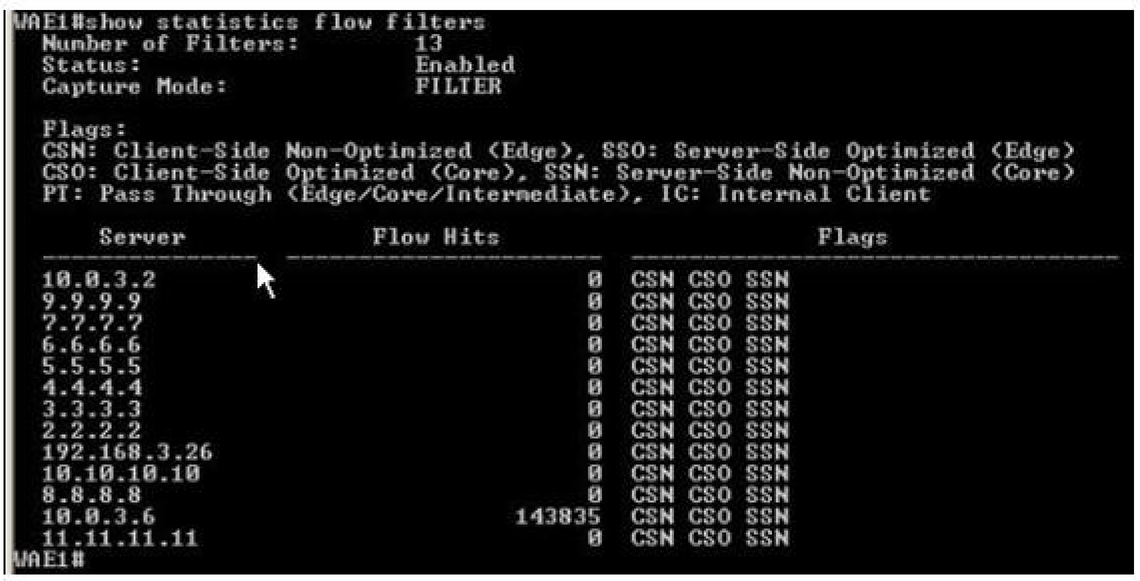

```
WAE1#show statistics flow filters
  Number of Filters:          13
  Status:                     Enabled
  Capture Mode:               FILTER

  Flags:
  CSN: Client-Side Non-Optimized (Edge), SSO: Server-Side Optimized (Edge)
  CSO: Client-Side Optimized (Core), SSN: Server-Side Non-Optimized (Core)
  PT: Pass Through (Edge/Core/Intermediate), IC: Internal Client

      Server              Flow Hits                      Flags
  ---------------  ---------------------  ----------------------------------
  10.0.3.2                             0  CSN CSO SSN
  9.9.9.9                              0  CSN CSO SSN
  7.7.7.7                              0  CSN CSO SSN
  6.6.6.6                              0  CSN CSO SSN
  5.5.5.5                              0  CSN CSO SSN
  4.4.4.4                              0  CSN CSO SSN
  3.3.3.3                              0  CSN CSO SSN
  2.2.2.2                              0  CSN CSO SSN
  192.168.3.26                         0  CSN CSO SSN
  10.10.10.10                          0  CSN CSO SSN
  8.8.8.8                              0  CSN CSO SSN
  10.0.3.6                        143835  CSN CSO SSN
  11.11.11.11                          0  CSN CSO SSN
WAE1#
```

Figure 4-17 depicts either a communication issue with the console or that there are no servers defined in SuperAgent. To resolve, ensure that servers are configured within the SuperAgent Management Console and use the **show statistics flow monitor tcpstat-v1** command to troubleshoot connectivity problems.

Figure 4-17. Problem Reported in the SuperAgent Management Console

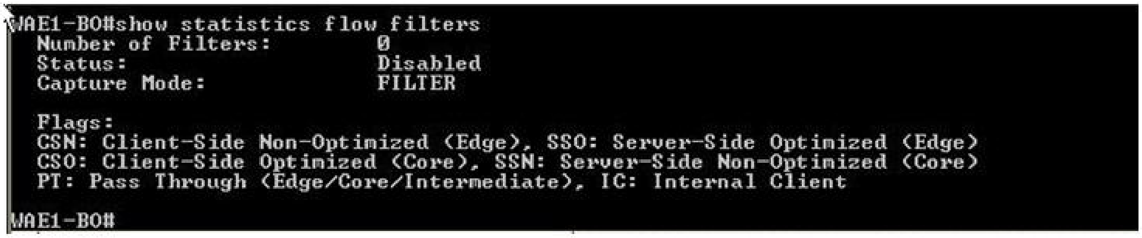

```
WAE1-BO#show statistics flow filters
  Number of Filters:          0
  Status:                     Disabled
  Capture Mode:               FILTER

  Flags:
  CSN: Client-Side Non-Optimized (Edge), SSO: Server-Side Optimized (Edge)
  CSO: Client-Side Optimized (Core), SSN: Server-Side Non-Optimized (Core)
  PT: Pass Through (Edge/Core/Intermediate), IC: Internal Client

WAE1-BO#
```

4.5 Summary

Cisco IOS measurement tools give IT directors and managers the needed tools to benchmark the various components of their network. Cisco IOS comes packaged with tools such as NetFlow and IPSLA polling that provide the platform to build detailed network monitoring abilities as well as the WAAS Flow Agent that allows network visibility into optimized TCP traffic. This visibility can reduce the amount of time required to troubleshoot network issues and bring about a resolution as cost effectively as possible.

5 Traffic Classification

In a typical network, the traffic through the network is heterogeneous and consists of flows from multiple applications and utilities. Many of these applications are unique and have their own requirements with respect to network parameters such as delay, jitter, etc. Unless these requirements are met, the quality and usability of these applications will be severely compromised. While meeting these requirements in a Local Area Network (LAN) with its huge bandwidth might be easy, it usually is a challenge to meet them on the WANs, which have bandwidth constraints.

Thus, traffic management on the WANs must exist in order to properly prioritize different applications across the limited WAN bandwidth and ensure that these requirements are met. In addition, a proper understanding of the applications and protocols in the network traffic is essential for any network manager to implement appropriate security policies. In a real network, user perception also matters. Although a user application might allow large delays or jitter, the user might be very sensitive to long wait times. Managing network traffic thus requires a judicious balance of all these priorities.

Classification of traffic is only the first step that helps identify different applications and protocols that exist in a network. Various actions, such as monitoring, discovery, control, and optimization can then be performed on the identified traffic with the end goal of improving the network performance. Typically, once the packets are classified (identified) as belonging to a particular application or protocol, they are marked or flagged. These markings or flags help the router determine appropriate service policies to be applied for those flows.

In other words:

- Classification is a technique that identifies the application or protocol, and
- Marking is the process that colors the packets (or just lets them through untouched) based on certain classification policies, which are used by the routers internally, or further downstream (depending on the kind of coloring) to provide appropriate treatment to those packets.

There are two other approaches to classifying traffic:

1. Classifying the packet based on the payload, i.e., Payload-Based Classification. In this method, packets are classified based on the fields of the payload, such as Layer 4 ports (source or destination or both).

2. Classification based on a statistical method that uses statistical analysis of the traffic behavior like inter-packet arrival, session time, etc.

The payload-based method is most prevalent. However, more often than not, it fails with encrypted and tunneled traffic. The Payload-Based Classification technique can be divided into generic or basic payload analysis or advanced payload analysis. The generic approach to traffic classification is based on information in the IP header. Typically, the following information is looked at.

Layer 3 address (IP address)
Layer 2 address (MAC)
Protocols

This technique is very rudimentary and does not provide classification for most of the applications. A classification method based on the placement of traffic (ingress interface) also exists, but is not widely used. (Therefore, it will not be examined here.)

All generic classification techniques based on Destination IP address, Source IP address, or IP protocol, etc. are limited in their ability as the inspection is limited to the IP header only. Similarly, classifying based on Layer 4 ports only is also limited. The problem with this approach is that not all current applications use standard ports. Some applications even obfuscate themselves by using well the defined ports of other applications (e.g., IM applications may run over TCP port 80, which is generally used for HTTP). Hence, the Layer 4 port mechanism of application identification is not always reliable.

Advanced classification techniques rely on deep packet inspection (DPI). There are varieties of DPI techniques, such as pattern analysis or behavior analysis, as detailed in 5.2, “Deep Packet Inspection.” These are much more reliable than the generic classification technique.

Figure 5-1 illustrates various classification methods and techniques.

Figure 5-1. Classification Methods and Techniques

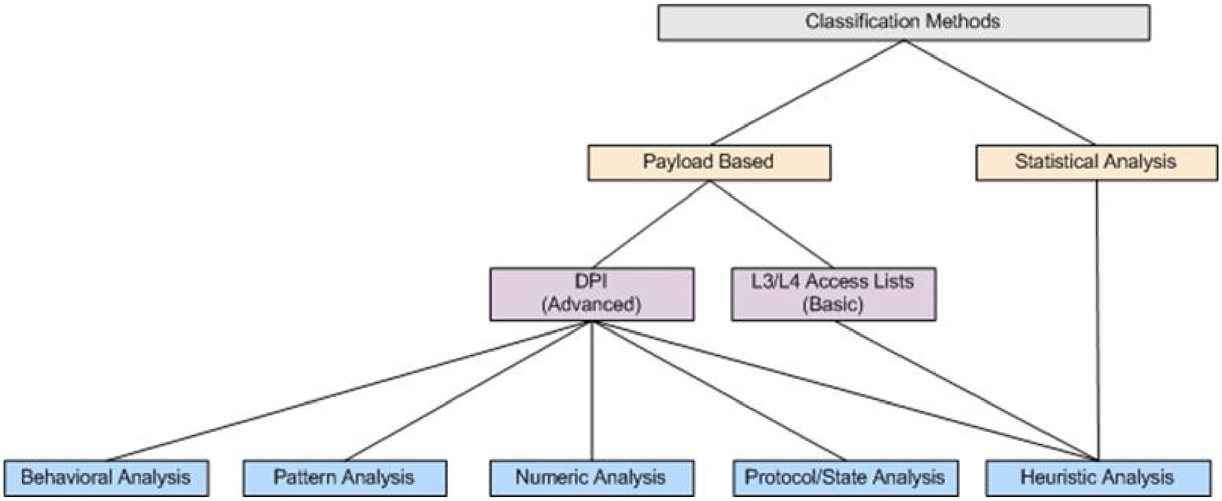

5.1 Payload-Based Traffic Classification

Payload-based classification methods can also be divided based on the processing method used for classifying traffic. Regardless of the method, of which there are four, all of them use one or more payload inspection techniques like Deep Packet Inspection to verify and classify traffic.

1. Packet-Based No State (PBNS) is the simplest and involves checking the payload for certain parameters like port numbers. It is less taxing on the CPU. For example, a simple access-list based port matching like the one below can identify all Telnet traffic.

   ```
   access-list 101 permit tcp any any eq telnet
   ```

   This method typically utilizes the basic payload-based classification technique. However, as already discussed, it is not always accurate or fully usable as the classification is on a per-packet basis without regard to an application session and is also limited by how deep inside the packet the verification of the flow goes.

2. Packet-Based per Flow State (PBFS) method is based on flows. A flow is defined as a sequence of packets from a sending application to a receiving application. In this method, a table to track each session based on the 5 tuples (source address, destination address, source port, destination port, and the transport protocol) is maintained for each flow. Since a flow has multiple packets, once a packet is marked as belonging to an application all subsequent packets in the flow need to be

marked as such. For example, in a typical VoIP call, H.323 is used for setting up the call and then RTP/RTCP is used for carrying the actual voice traffic. Once a H.323 flow is identified and marked, subsequent RTP/RTCP flows to the same source IP/destination IP pair are tagged with the same parameters.

3. Message-Based per Flow State (MBFS) method is similar to PBFS, except that this operates on messages instead of packets. A message is protocol dependent and is an information element that can span multiple packets or a single packet can contain multiple messages. Since it operates on messages there needs to be some sort of TCP normalizer to take care of IP fragments and TCP Segments. Since entire messages have to be considered, there is a considerable increase in memory requirements.

4. In the Message-Based per Protocol State (MBPS) method, not only is the application tracked but also what the application is transmitting. In other words, complete knowledge of the protocol state machine is needed to implement this method. This is the most taxing method on the CPU and has more memory requirements too.

The last three methods, PBFS, MBFS, and MBPS, utilize advanced classification techniques that are based on DPI.

5.2 Deep Packet Inspection

Although most general applications can be determined or at least guessed based on the L3 and L4 information, additional granular sub-classes within applications (like URLs) or specific kinds of messages within the application (like voice within IM streams) are required. For proper classification and sub-classification, it is necessary to do a deep packet inspection (DPI) and verify what the application is.

Most DPI mechanisms use Signature Analysis to understand and verify different applications. Signatures are unique patterns that are associated with every application. In other words, each application is studied for its unique characteristics and a reference database is created. The classification engine then compares the traffic against this reference to identify the exact applications. That means the reference needs to be updated periodically to keep current with new applications as well as new developments in existing protocols.

There are different Signature Analysis methods. The most popular methods include:

- Pattern analysis
- Numerical analysis
- Behavioral analysis
- Heuristic analysis
- Protocol/state analysis

5.2.1 Pattern Analysis

Some applications embed certain patterns (bytes/characters/string) in the payload of the packets, which can be used by the classification engine to identify such protocols. Depending on the application, these patterns may not necessarily be always located at a specific deterministic offset. The patterns might be present in any position in the packet. Still, the classification engine can identify these packets. However, not all protocols embed special pattern, string, or characters in the packets and hence this approach will not work for them.

5.2.2 Numerical Analysis

Numerical Analysis involves looking into the numerical characteristics of packets such as payload size, number of response packets, and offsets. Older Skype versions (pre-2.0) are good cases for such analysis. The request from client is an 18-byte message and the response it receives is usually 11 bytes. As the analysis may be spread over multiples packets, the classification decision might take more time.

5.2.3 Behavior & Heuristic Analysis

Occasionally, analyzing the traffic behavior would produce greater insight into the applications that may be running. This behavior can be used to classify such applications. Similarly, by doing a statistical (heuristic) analysis of the inspected packets, the underlying protocol can be classified. Behavior and Heuristic analysis typically go hand in hand and many of the anti viral programs use these techniques to identify viruses and worms.

5.2.4 Protocol/State Analysis

In some applications, the protocol follows a certain sequence of steps or actions. For example, a typical HTP GET request from a client is followed by a valid response from the server. Such protocol conformance can be used to classify such applications

As more and more applications start encrypting traffic, it becomes a challenge for any kind of classification mechanism to classify them accurately. With encryption, all upper layer information suddenly becomes invisible to DPI mechanisms. Behavior and Heuristic analysis methods can help in identifying some of these applications. Newer classification mechanism exist that use behavior and heuristic analysis methods (along with intelligent algorithms such as clustering algorithms) and can help identify encrypted traffic.

None of these methods, on their own, can provide satisfactory classification of all applications. Hence, in a typical deployment all of these are used in conjunction.

5.3 Cisco Classification Technologies

Cisco classification technologies include QoS access lists and DPI engines

5.3.1 QoS Access Lists

5.3.1.1 QoS Software-based L3/L4 access-lists

Cisco IOS provides the ability to configure Layer 3 or Layer 4 based access lists that can be used with the QoS to classify different types of traffic. Specific QoS classes can be configured to use different access lists to match traffic and based on the match the packets can be marked. The matching can be based on Layer 3 addresses (source/destination IP), Layer 4 protocol or ports, or a combination thereof.

5.3.1.2 Classification with QoS ACLs in hardware

In addition to software based ACLs, Cisco platforms like 6500 & GSR provide the ability to do ACL lookups in hardware. For example in a 6500 platform, these ACLs can be programmed into the TCAM (Ternary Content Addressable Memory) and lookups performed against those entries. However, it should be noted that TCAMs have finite memory and if not carefully planned the resources can be exhausted. TCAM lookups are much faster than traditional software lookups as they are performed in the hardware and hence classification based on TCAM lookups is much faster.

5.3.2 DPI Engines

DPI engines can be co-resident in the software or can be dedicated hardware. Both have advantages and disadvantages. While a dedicated hardware provides speed and versatility, the cost of deploying such a box restricts their usage to high traffic volume environments like a Data Center or a large Enterprise Branch office. Cisco’s Service Control Engine (SCE) is a good example of a dedicated hardware DPI appliance. Software-based DPI engines are cost effective, but they do consume CPU cycles and hence can be deployed only at low or medium traffic volume environments such as those found in a small or medium Enterprise Branch Office.

5.3.2.1 Service Control Engine (SCE)

Cisco Service Control Engine (SCE) is a DPI device that can do DPI and detect traffic patterns at line rates. SCE incorporates many DPI technologies such as protocol/state analysis, pattern analysis, and behavioral and heuristic analysis. SCE can also do subscriber-level classification.

The Cisco SCE can be deployed in-band or out-of-band. It is typically deployed in the Data Center. If it is deployed in band, all the traffic in the network passes through SCE. If it is deployed out of band then a copy of all the traffic is passed onto SCE by the DC switch. It should be noted that in out of band mode, the SCE can only perform monitoring.

5.3.2.2 Network Based Application Recognition (NBAR)

NBAR is an application-aware classification feature in IOS. NBAR can look deep inside a packet and do stateful analysis of the information in the packet. It can recognize a number of applications, including ones that use ephemeral ports. Even with a given protocol, NBAR can look so deep inside the packets that it can categorize packets that are of the same protocol, but with different protocol-specific parameters. For example, NBAR can classify based on the URL for HTTP packets and based on ICA traffic for CITRIX ICA.

Typically, QoS and NBAR are used in conjunction. NBAR is used to recognize specific applications and QoS is used to mark them and provide appropriate treatment based on the markings.

5.4 Packet Markings

Once the flow and packets have been identified, they need to be marked so that appropriate service policies can be applied on them. The markings or flags can be set in a number of ways: for IP, Type of Service (ToS) or Differentiated Services Code Point (DSCP); for Ethernet packets, VLAN priority, etc. However, L3 markings are the most widely used method.

5.4.1 L2 Packet Markings

Many popular Layer 2 technologies like Asynchronous Transfer Mode (ATM), Frame Relay (FR), and Ethernet provide options for marking the packets to help provide differentiated treatment. The disadvantage with all L2 markings is that they are limited to their L2 domain.

5.4.1.1 ATM Marking

ATM networks use simple cell marking of the Cell Loss Priority (CLP) bit on the cell header to indicate whether the cell can be dropped in times of congestion. A typical ATM Cell consists of a 5-byte header and 48 byte payload. Figure 5-2 illustrates an ATM cell header and the location of the CLP bit.

Figure 5-2. ATM Cell Header

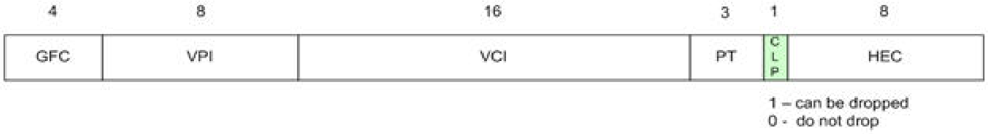

As Figure 5-2 indicates, if the CLP bit is set to “1,” the cell can be dropped in times of congestion. Traditionally, an ATM switches is used to mark the CLP bit. However, this functionality was added to CISCO IOS as a part of enhanced QoS feature set. In a typical network, the user may choose to mark non-critical traffic flowing through the ATM switches with the CLP bit. This will ensure bandwidth availability for critical traffic in times of congestion.

5.4.1.2 Frame Relay Marking

Similar to the CLP bit on the ATM header, the Frame Relay Header also has a bit called Discard Eligible (DE) to indicate whether the frame can be dropped in times of congestion. Figure 5-3 illustrates a typical Frame Relay Header.

Figure 5-3. Frame Relay Header

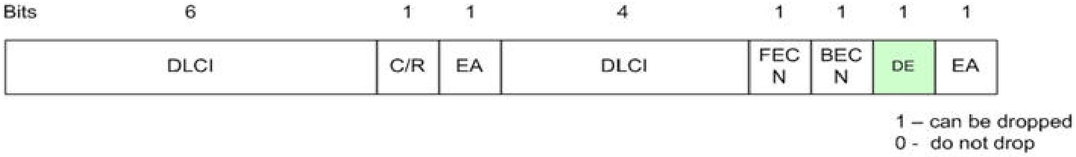

As in the ATM case, the DE bit can be set for non-critical traffic to help alleviate congestion.

5.4.1.3 Ethernet Marking

The IEEE 802.1p standard provides traffic class expediting and dynamic multicast filtering. It enables Layer 2 switches to prioritize traffic. The 802.1p defines 3 bits in the header for classification, which helps classifying traffic into eight different traffic classes. It should be noted that 802.1p is an extension of 802.1Q standard and they work together. Figure 5-4 shows an Ethernet 802.1q frame and the TAG byte where the Priority bits are located.

Figure 5-4. Ethernet 802.1Q Frame

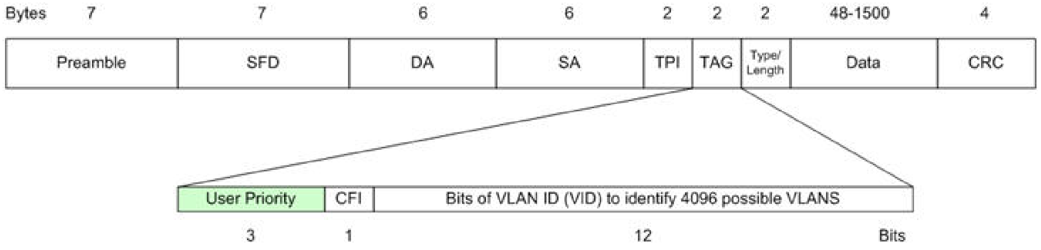

IEEE has put forth recommendations on various traffic types, corresponding traffic classes, and priorities to be used with 802.1p standard. They are listed in Table 5-1.

Table 5-1. Traffic Classes to Priority mapping

Traffic Type	**Traffic class**	**Priority**
Bulk transfers, Games etc.	Background	1
Less than 10 millisecond delay	Voice	2
Less then 100 millisecond delay	Video	3
Some important application	Controlled Load	4
Best Effort for important users	Excellent Effort	5
Ordinary LAN priority	Best Effort	6
High requirement to get through to maintain and support the network infrastructure	Network Control	7

5.4.2 L3 Packet Markings

Similar to Layer 2 headers, the IP header has fields that can be used to classify traffic into treatment groups. The most widely used L3 marking techniques are Type of Service (ToS) and DSCP. Figure 5-5 below shows a typical IP Header.

Figure 5-5. IP Header

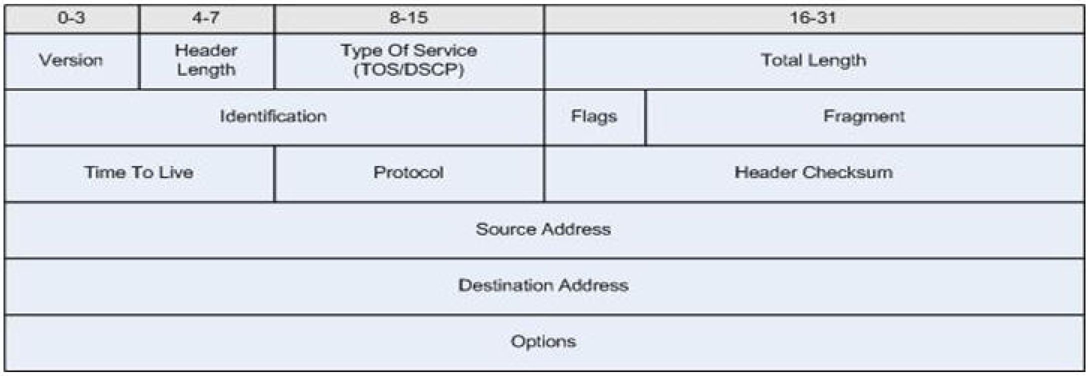

5.4.2.1 ToS

ToS was originally defined in RFC 791 and 795 and was further modified/updated by other RFCs like RFC 1122, RFC 1123, and RFC1349. Although the field has been there for quite sometime, it has not been widely used. Its use has been superseded by DSCP today. The ToS fields are shown in Figure 5-6:

Figure 5-6. ToS Fields

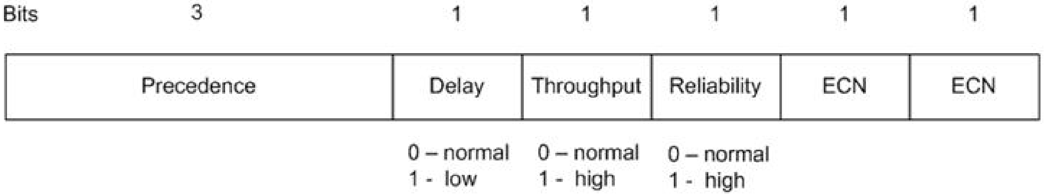

IP Precedence - Table 5-2 details the precedence bits and their possible values:

Table 5-2. ToS Precedence Bits and their values

Binary	Decimal	Classification
000	0	Routine
001	1	Priority
010	2	Immediate
011	3	Flash
100	4	Flash Override
101	5	Critical
110	6	Internetwork Control
111	7	Network Control

- Delay - when set to 1, the packet requests low delay.
- Throughput - when set to 1, the packet requests high throughput.

- Reliability - when set to '1,' the packet requests high reliability.

5.4.2.2 DSCP

Differentiated Service Code Point (DSCP) was defined in RFCs RFC 2474 and RFC 2475. DiffServ (DS) has more priority levels than that of ToS since it has more priority bits. DiffServ uses the same three most significant like ToS to define priority, but uses the next three bits to further refine them. DS fields are used to determine the per-hop behavior (PHB) of the packet. The following byte layouts illustrates the difference between ToS and DiffServ.

5.4.2.2.1 ToS Byte

P2	P1	P0	T2	T1	T0	ECN	ECN

- IP precedence—three bits (P2 to P0)
- Delay, Throughput and Reliability—three bits (T2 to T0)
- ECN — two bits

5.4.2.2.2 DiffServ Field

DS5	DS4	DS3	DS2	DS1	DS0	ECN	ECN

- DSCP—six bits (DS5-DS0)
- ECN—two bits

The ECN bits were not in the original DSCP RFCs. They were later added later by RFC 3168 to allow for congestion notification in the path.

RFC 2597 for DiffServ defines Assured Forwarding (AF) PHB that can be used by a service provider to provide different forwarding assurances based on different AF classes. There are four different AF classes each with three different drop probabilities.

RFC 2598 for DiffServ defines Expedited Forwarding (EF) PHB. "The EF PHB can be used to build a low loss, low latency, low jitter, assured bandwidth, and end-to-end service through DS (Diffserv) domains. Such a service appears to the endpoints like a point-to-point connection or a "virtual leased line." This service has also been described as “Premium service." Codepoint 101110 is recommended for the EF PHB.

All the PHBs require vendor support to implement and not all vendors fully support them. Table 5-3 shows the guidelines for DSCP service classes as defined by RFC 4594.

Table 5-3. DSCP to Service Class Mapping

Service Class Name	DSCP Name	DSCP Values	Application Example
Network Control	CS6	110000	Network Routing
Telephony	EF	101110	IP Telephony Bearer
Signalling	CS5	101000	IP Telephony Signalling
Multimedia Conferencing	AF41, AF42,AF43	100010,100100,100110	H.323/V2 video conferencing (adaptive)
Real-Time Interactive	CS4	100000	Video conferencing and Interactive gaming

Multimedia Streaming	AF31, AF32, AF33	011010, 011100,011110	Streaming video and audio on demand
Broadcast Video	CS3	011000	Broadcast TV & live events
Low-Latency Data	AF21, AF22, AF23	010010, 010100, 010110	Client/server transactions, Web-based ordering
OAM	CS2	010000	OAM & P
High-Throughput Data	AF11, AF12, AF13	001010, 001100, 001110	Store and Forward Applications
Standard	DF (CS0)	000000	Undifferentiated applications
Low-Priority Date	CS1	001000	Any flow that has no BW assurance

5.5 Summary

Classification involves proper identification of different application flows and packets in the traffic and their appropriate marking. Once the packets are classified, the router can apply appropriate service policies for those packets. Typically, QoS is used to provide appropriate treatment to different traffic based o the configured policies. As was discussed earlier, each application has its own characteristics and requirements. With the limited WAN bandwidth, QoS policies help in providing different treatments for different traffic classes. Various QoS mechanisms such as Congestion Management, Congestion Avoidance, Traffic Policing/Shaping, and Link Efficiency exist that can be used to manage the WAN bandwidth.

6 WAN and Application Optimization Technologies

WAN and application optimization comprises a framework of technologies that improves the network application experience at the branch, and makes better use of limited network resources. In some cases, the user experience simply needs to be maintained while other changes occur. For example, during server consolidation (a process where branch-based servers are relocated to a centralized data center over the WAN, the reliance as well as stresses on the WAN network increase but the user experience needs to be at least maintained. Additionally, given resource contention on the WAN, business criticality must map closely to network usage and access.

Various mechanisms are available for optimizing the WAN, ranging from technologies that allow horizontal scalability (the ability to cluster multiple devices, rather than vertical scalability- which requires more power in each device), such as server load sharing, advanced compression technologies, dynamic routing to place traffic on the best path, and intelligent flow replication reduction tools, to name a few.

This chapter explores this optimization framework, with an emphasis on basic issues being targeted, and on individual technology solutions. Note that not all of the technologies described in this chapter are part of the currently packaged "WAN-Optimization" solution.

6.1 Areas of Interest

As shown in Figure 6-1, the network between the branch and the datacenter can be divided into multiple spheres of work, each with its own challenges and opportunities for optimization.

Figure 6-1. Simplified View of a Typical WAN Topology

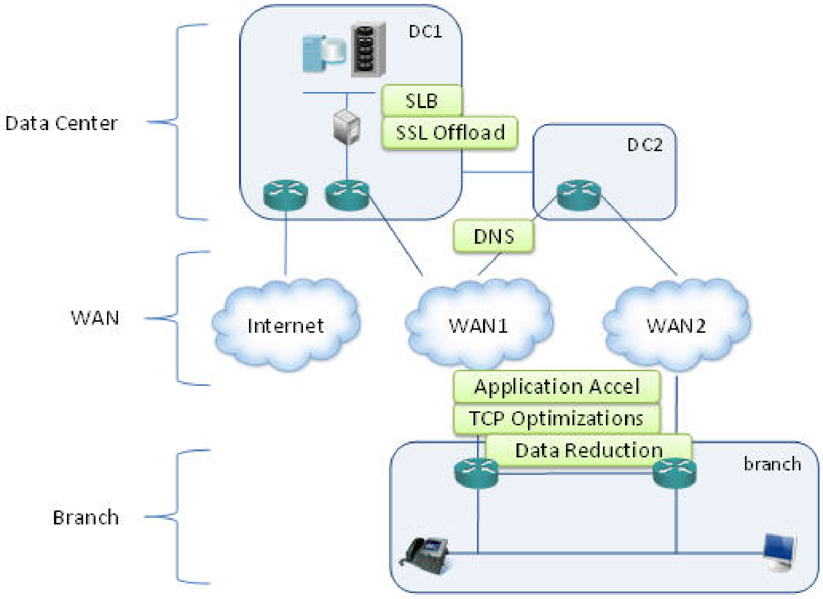

6.1.1 Layer 3 End Point Optimization and Server Selection

Because network communications is about connecting parties so that meaningful transactions can take place, one of the first and easiest optimizations is selecting the specific end-hosts involved in the transactions. When a network application user wants to use the application, he or she will interact with a server or a set of servers. The specific server is not relevant to the network application, as long as the server can handle the transaction.

Therefore, it is possible to have multiple servers that can each service the network application. These servers can even be geographically dispersed so that they are not only redundant and share the overall load, but can also be closer to the various clients. The capability to match a specific server to a client request is a form of optimization. After the server Layer 3 (L3) endpoint is established, it is generally very disruptive to the user, the network application, and the transaction to try to move the transaction to a different L3 address.

6.1.2 DNS-Based Optimization

When starting a network transaction, one of the first things that a client attempts to do is identify the L3 address for the remote peer. Generally, this L3 translation is done using Domain Name System (DNS), where a human-readable name (for example, erp.example.com) is converted to a routable IP address (such as 209.165.201.1 or 2001:0DB8::1). Along with translations to IPv4 addresses (called A records) and IPv6 addresses (AAAA records), DNS can translate services into specific L4 endpoints. For example, DNS service (SRV) records enable a client to locate a LDAP server (for example, Query sent for_ldap._tcp.example.com).

Because the authoritative DNS server returns the mapping between the DNS name and the L3 address, there is an opportunity to select a specific geographic location, or even a specific server, that is best positioned to handle that specific client. For example, assume that the client is located in Australia, and that servers in India, Germany, and the United States can all handle the transaction equally well.

An optimizing authoritative DNS solution, such as the Cisco Global Site Selector (GSS) product, can measure the latency between each server location and the client, and consider liveliness, loads, and so on. GSS can then determine which server is best able to answer the client query. Using our previous example, the Australian client might be directed to India or the United States, depending on the location of the lowest latency and best network conditions.

6.1.3 IOS DNS Views feature

In some cases, DNS can direct the client to a suboptimal server because of how the DNS query is sent. In many enterprise networks, clients are configured with the address of a DNS caching forwarder that queries the actual DNS system on behalf of the client. This approach is a problem when the DNS forwarder queries the optimizing authoritative DNS solution; the address and network location of the DNS forwarder is used instead of the address and network location of the client. Therefore, the client is sent to the server that is best for the corporate DNS forwarder, not for the client, which might actually talk to the server over a completely different path.

Figure 6-2 shows such an example: the server that the client needs to work with is on the Internet path available through the branch router. The corporate DNS forwarder query to the Internet returns the best third-party server address for the corporate network. With the IOS DNS views feature, the IOS router can intelligently send queries to the corporate DNS forwarder instead of the Internet based DNS forwarder. The separation of DNS queries is based on the domain name.

For example, suppose that the third party server is at purchasing.thirdparty.com and the corporate email cluster is at email.example.com. The DNS views feature can send queries in the example.com domain to the corporate DNS servers, while everything else is sent to Internet servers using the branch site IP addressing. Additionally, this feature can be used for non-Internet cases, such as when the branch is connected over a single link to a MPLS VPN that provides extranet services to multiple companies that each operate their own optimizing authoritative DNS servers.

Figure 6-2. DNS Views Feature

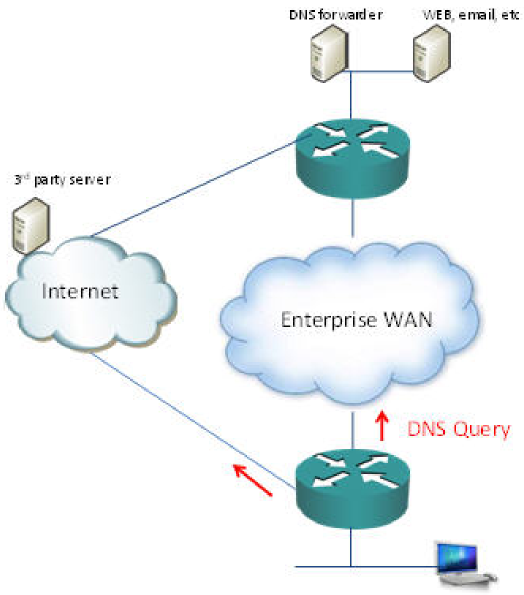

6.1.4 Anycast Addressing

Another, much simpler but more limited mechanism for redirecting clients to a better server is to use IP anycast addressing schemes. With anycast, the same L3 address is assigned to multiple physical servers, and the network routing protocol selects the “nearest” server. Because the same IP address is used for each server, this severely limits the kind of connection that can be established. The network cannot guarantee that the next packet sent to the server IP address will be received by the same server. While anycasting has been used with great success to optimize communication and distribute the load for some specific applications (for example, the root DNS nameservers use anycasting), it should be noted that the application and network protocol being used must be compatible with anycasting technology and is not commonly deployed.

6.1.5 Layer 7 Redirection

In L7 redirection, a client can initially connect to a well-known L3 address, and then is redirected to the actual best server using the L7 protocol. This approach, provided by the Cisco Application Control Engine (ACE) product, is often used with HTTP redirects. For example, a user might try to connect to http://wwwin.example.com/erp, but the server at wwwin.example.com knows that the erp network application is best served from another server called www-erp1.example.com, which hosts the erp application, is lightly loaded, and is geographically closer to the client. The wwwin.example.com server sends an HTTP redirect message to the client, asking it to connect to the www-erp1.example.com server. The initial connection and subsequent redirect request can be obtrusive and alarming to the user, and many non-HTTP applications do not support L7 redirection.

6.1.6 Local Server Load Balancing

Server load balancing (SLB) is based on the concept of front-ending a number of similar real servers with a virtual IP address. The client establishes an L3 (and possibly an L4) relationship with a server load balancer which, at connection initiation time, selects the best real server to terminate the connection on. This provides optimization, in that the client gets access to the network application in a resilient manner (if a real server is down, the SLB device takes that server out of the pool), that can also consider performance.

Applications or subapplications that are known globally by the same IP address can be divided into multiple pools of real servers that service the application. In the multipool case, the SLB device, such as Cisco IOS or Cisco ACE, parse beyond L4 (TCP/UDP port information) and try to identify the subapplication in order to map the session to a real server. For HTTP, a specific URL can denote the subapplication, as shown in Figure 6-3.

Because SLB devices hide the L3 information about the real server from the client, the client is never aware of the specific real server that that serviced it. While for a specific TCP session the same real server will definitely be used, it is less guaranteed that across multiple TCP sessions that the same real server will be used for a client. Depending on the implementation of the network application, ‘stickiness’ to a specific real server may be required. The SLB devices generally provide a variety of mechanisms for stickiness including client IP, insertion of HTTP cookies etc. It is important to understand that while the network operator may not have a choice about whether to apply stickiness, it will limit the flexibility of loadsharing amongst the real servers.

Figure 6-3. SLB Example

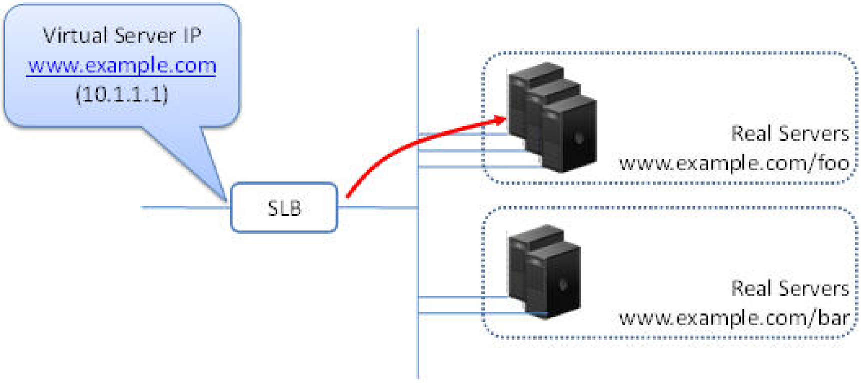

Many SLB devices, such as Cisco ACE, incorporate a Secure Socket Layer (SSL) offload feature that is not directly related to L3 endpoint selection and optimization. Using the SSL offload feature, the SLB device terminates the SSL session and exposes only a cleartext session to the real server. Other common high-CPU utilization functions, such as HTTP compression, are often available on SLB devices. This frees valuable compute cycles on the real servers. In the case of HTTP compression, there are also benefits of WAN bandwidth reduction and faster load times for the HTTP client.

It should be noted that many of the various techniques that lock down transactions between a client and a specific physical server can be used concurrently. For example, an optimizing intelligent DNS authoritative nameserver can be used with L7 redirects, and with local server load balancers.

6.1.7 Path Optimization

After the L3 endpoints are established, communication between the client and server becomes a matter of IP packets traversing a path between the two. However, there are typically multiple possible paths with many different attributes. Some paths might have high bandwidth and low latency, while other paths have little bandwidth but also provide low latency. The challenge of path optimization (implemented as a feature in Cisco IOS) is to manage traffic so that specific traffic classes are placed on the best network path.

Traditional routing is only responsible for announcing reachability to an L3 prefix. Reachability is based on the ability of the routers along the path to maintain adjacencies to each other. Therefore, even though the routing protocols advertise reachability for a prefix, actual reachability might not exist, or might even be degraded (for example, congestion along a path results in dropped packets). Traditional routing protocols also use static metrics for evaluating paths and do not account for changing network conditions. In practice, congestion, delay, and varying link utilization loads affect network application performance, even though regular routing does not take these network path attributes into account.

Because the WAN edge of a network is also a flow aggregation point, possibly a bandwidth bottleneck, and usually the last point of control in routing policy, the WAN edge becomes the natural place to make path selection decisions. Traffic at this point is already treated with optimization technologies, so the flow bandwidth observed by the WAN edge router is after optimization.

Generally, delay, bandwidth, and other types of issues are less prevalent and much easier to solve within a local branch or campus network. In a local network, it is relatively straightforward to add additional links or upgrade the bandwidth. However, on the WAN side, increasing bandwidth can add significant continuing costs, and there is little control over congestion and latency, especially in shared infrastructure networks, such as MPLS-VPN networks, and in cases where the networked sites are geographically dispersed.

In traditional packet forwarding systems, only a limited number of mechanisms are available to distribute traffic among links. Placing packets in round-robin fashion is not viable: different paths have varying latencies that cause out of order packets, eventually resulting in far lower performance for the network application.

Hash based mechanisms, for example, EtherChannel or Cisco Express Forwarding (CEF), were devised so that flows would be statistically distributed, based on mathematical functions, among different paths. However, the traffic rate for a flow is not considered, and it is possible for multiple high bandwidth flows to be sent on an already congested link while other links remain underutilized. Finally, link bonding technologies, such as Inverse Multiplexing over Asynchronous Transfer Mode (ATM-IMA) and Multilink Point to Point Protocol (MLPPP) try to slice packets so that they each packet is simultaneously sent on multiple links.

Link bonding technologies generally place a substantial load on the fragmenting device and the reassembly device, and are sensitive to intramember-link delay variation. Finally, the remote peer must be the same type of device, limiting the use of multiple WAN providers.

Figure 6-4. Path Optimization for Voice and Email Traffic

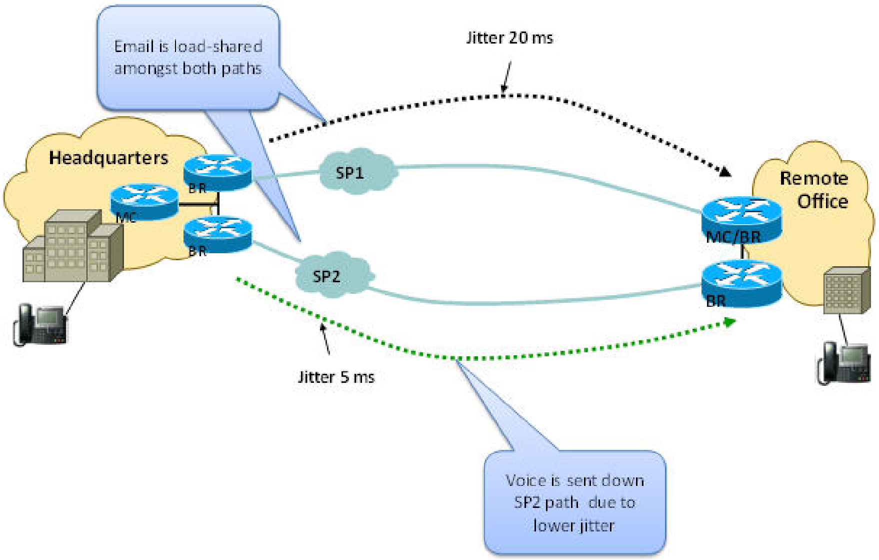
Email is load-shared amongst both paths
Jitter 20 ms
Headquarters
BR
MC
BR
SP1
SP2
Remote Office
MC/BR
BR
Jitter 5 ms
Voice is sent down SP2 path due to lower jitter

Figure 6-4 illustrates path optimization for different traffic classes. Voice traffic requires low latency; if a path provides acceptable minimum network latency compared to other paths, voice traffic is best placed on that low-latency path. However, other traffic classes, such as email, are not sensitive to latency, and could be placed on a more varied set of paths. Identifying traffic classes, measuring network attributes such as latency and packet loss, and placing traffic on the most appropriate path is the role of path optimization.

Path optimization can have varying levels of granularity, ranging from manipulating only the path to an L3 destination network, to per-flow manipulation and per-application path selection so that traffic uses the best available path. Using flow selection mechanisms and awareness of exit link utilization enables path optimization technologies, such as Cisco IOS Performance Routing (PfR), to distribute flow load across multiple paths on a per flow basis.

Using path optimization, the full aggregate bandwidth for a network site can be effectively used by performing link utilization based traffic load-sharing. Additionally, pockets of free bandwidth are created by the traffic distribution across all the links to support temporary data bursts (called microbursts) that exist in most real-world application profiles.

6.2 Layer 4 Optimizations

L4 optimizations include TCP stack optimization and payload compression.

6.2.1 TCP Stack Optimization

The previously described optimization technologies involved endpoint and path selection. Notably, these technologies do not change the of the IP traffic payload. However, there are opportunities for optimization at the TCP stack level, on the actual payload, and finally on exploiting behaviors and patterns at the network application.

There are many implementations of the TCP stack across operating systems and even among different versions of the same operating system. Not all the stacks are alike, and some perform better under certain conditions than others. Of course, there are also cases where TCP stacks are severely out of date and have not been maintained. Over the years, there have been many advances in TCP technology, and only some advances have trickled down to implementations. TCP stack optimization uniformly applies the behavior of an advanced TCP stack across network applications.

TCP is responsible for reliably transporting a message stream from one computer to another. As data is acknowledged by the remote host, succeeding data in the stream can be sent. This enables a network application to work on good in-order data without missing parts.

TCP controls how much data can be in transit and not acknowledged at any particular time, and this effectively controls the overall data transmission rate. The number of outstanding packets allowed without acknowledgement is known as the maximum window size (MWS). Over time, if there are no lost packets, the MWS expand to take more of the available bandwidth in the network path.

A network path, depending on its bandwidth and end-host to end-host latency, can hold a certain amount of traffic in transit, called bandwidth delay product (BDP). As shown in Figure 6-5, a network path with 155Mb/s of bandwidth and a round-trip time (RTT) of 10ms can hold 193KB in transit. Meanwhile another path (possibly on the same link, but to another host much further away) with the same bandwidth but with a 200ms RTT can hold 3860 KB in transit.

Because the TCP header used to specify window size contains two bytes, the MWS is only 65KB. The example in Figure 6-5 requires a much higher window size to pack the long network path full of data. RFC 1323 introduces a mechanism for scaling window sizes that supports much larger window sizes.

Figure 6-5. Comparing BDPs

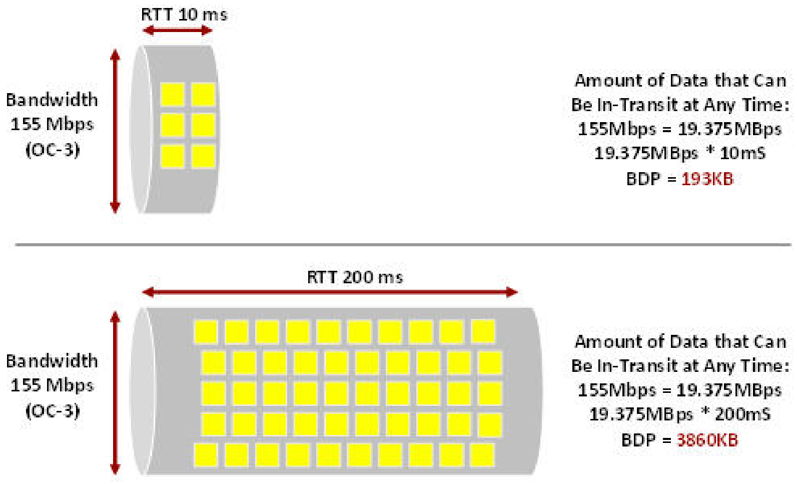

Because the capacity of the network path is unknown, starting at too high a rate (a large MWS) can cause severe congestion and packet loss not only for a specific session, but also for all sessions sharing the same path. For this reason, traditional TCP employs a technique known as slow-start that initiates the session with a small window that expands as good network behavior is observed. However, for higher latency links the RTT can slow such expansion so much that the bandwidth potential of the link is never realized. RFC 3390 supports larger initial windows so that a higher data rate is set when a connection is initiated.

When packets are lost, presumably because of congestion, traditional TCP aggressively reduces the MWS and slowly increases the data rate. This is known as congestion avoidance. For higher latency paths, it can take quite a long time to build to the proper data rate while causing minimal impact to other flows on the network path.

A multipacket loss event in traditional TCP can also result in unnecessary data retransmission or multiple RTT of delay). This is because in traditional TCP, only properly ordered packets are acknowledged. RFC 2018 introduced a mechanism to perform selective acknowledgements, and other mechanisms for advanced congestion avoidance are available. Figure 6-6 illustrates this.

Figure 6-6. Cumulative Traditional TCP Stack Delays and Underutilized Links

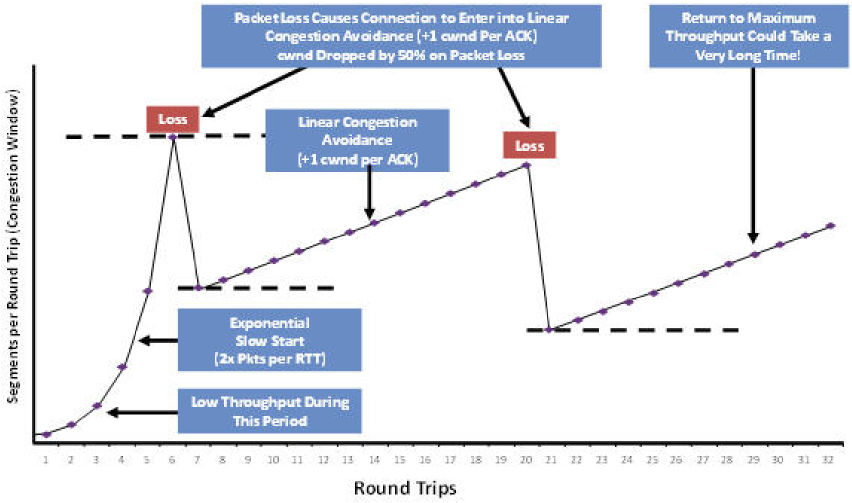

These are only some of the added capabilities of advanced TCP stacks that are often not available on a particular operating system version, disabled by default, or not implemented on any client-facing operating system. As TCP sessions traverse the network, they can be intercepted at an aggregation point and ”upgraded” with additional properties, regardless of the capabilities of actual end host. Cisco WAAS includes an intercepting TCP stack enhancer, called TCP Flow Optimization (TFO), which can apply the described TCP stack improvements to network flows.

6.2.2 Layer 4 Payload Compression

Data compression on limited bandwidth WAN links has been available for several years. Compression not only increases effective data transfer rates, but also reduces buffering, which can effectively reduce latency.

Data compression has been implemented at the lower data communications layers for a long time. For example, V.42bis on dialup modems and various schemes (RFC 1974, RFC 1967, and so on) used the Point-to-Point Protocol (PPP) Compression Control Protocol (CCP) to negotiate compression protocols.

In Frame Relay (FR), FRF.9 can compress data on a per-FR private virtual circuit (PVC). However, as these compression schemes are tied to the lower network layers, they are useful only for single L3 hops. Additionally, such compression technologies were limited by the size of the dictionary table used to store compressing patterns. The dictionary size had to be balanced against the practice of communicating over a transport system that did not reliably deliver data in order. A smaller dictionary size (usually limited to a packet) reduced the range of the data stream that could be referenced.

At the higher IP layers, there have been nondata payload compression schemes that act on TCP/IP headers, for example, RFC 1144, and compression schemes that act on specific application traffic, such as Real-time

Transport Protocol (RTP), for example, RFC 2508). In practice, the techniques in these examples are effective only on extremely low bandwidth paths, such as T1and slower, where bandwidth savings constitute higher percentage of the overall WAN speed. Both RFC 1144 and RFC 2508 are limited to single hop L3.

The IPComp protocol (RFC 3173) added a capability to compress the IP payload between two hosts or routers. Unlike the previously described compression methods, IPComp can take place over multiple L3 hops, so IPComp can be practically deployed in MPLS VPNs and in VPNs created over the Internet. However, IPComp still works over an unreliable transport and maintaining a large dictionary is a significant task. For combined IPSec and IPComp, the compression dictionary is limited to single packets and severely limits potential compression. In addition, in many deployment scenarios, traditional IPComp would require the creation of tunnels.

In recent years, TCP-based network compression has been developed to address many of these issues. TCP-based network compression combines large RAM based pattern matching dictionaries and large disks (several hundred gigabytes). Because TCP provides reliability, the compression dictionary and history are not limited to a packet and extremely high compression ratios can be achieved. This form of compression is different from object-caching mechanism (for example, HTTP caching). TCP-based network compression can work effectively when changes occur in the object.

For example, with TCP stream compression or data reduction, a change in a Word document would transfer only the changed content with compressed references to the rest of the content. In HTTP object-caching, the object must first be recognized (usually based on a name), and the entire object is transferred. Because TCP-based network compression is performed at L4, compression can occur over multiple router hops. This technology can be deployed in many kinds of networks, including MPLS VPN, Internet, and so on.

In a Cisco implementation, the WAAS product line, TCP data reduction technology is available in the Data Redundancy Elimination (DRE) function, which stores variable-length patterns on the WAAS hard drive. LZ compression is performed, after DRE (across multiple packets), for new patterns that DRE has not yet learned, or for micropatterns, as illustrated in Figure 6-7.

Figure 6-7. A WAAS Device Performing DRE and LZ Compression

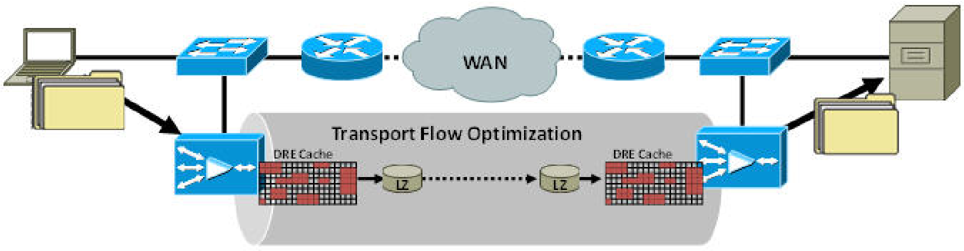

6.3 Layer 7 Optimizations

The L4 based, TCP-based compression method described in the previous section does not examine the actual traffic payload. The method simply observes patterns in the data stream and performs a compression algorithm. In certain cases, however, when the context of the data and the protocol is well-understood, much higher bandwidth, message, and latency (because of fewer RTTs) reductions can be achieved.

Many network applications were first developed for, and deployed in, only high bandwidth, low-latency networks, such as a high-speed LAN. Networking assumptions that were built into these applications usually do not hold true in WAN deployments. L7 optimizations attempt to improve the performance of such network applications on the WAN. L7 optimizations use a variety of techniques, such as aggregating messages, caching, and advanced compression algorithms.

6.3.1 HTTP Compression

The HTTP/1.1 specification (RFC 2616) allows the use of various types of compression on HTTP payloads. This enables servers to directly send compressed data to a web browser. (Almost all browsers support such encoding.) A decompressing network node at the network edge is not required... While many web servers can perform this compression, web servers are usually better used for dynamic web page creation and interactions with backend servers.

The more mundane task of compression can be offloaded to network devices that can perform the compression in hardware without adding much additional latency. As in the TCP case, a compression dictionary and history can be highly effective within the scope of an HTTP session because references can be passed across multiple packets in the stream.

However, unlike DRE technology, the compression history in HTTP/1.1 applies only to one HTTP session. Downloads by other PCs at the same site do not see any improvement because the compressed content can only be used once. In the DRE case, traffic for multiple computers traffic passes through an L4 optimizer, which keeps a copy of the data patterns for future reuse and referencing. Therefore, HTTP compression is generally less effective than DRE.

6.3.2 Application Acceleration

There are many cases where the implementation of a network application or protocol does not work well over the bandwidth limited and higher latency WAN environments. Application acceleration is designed to understand a network application and reinvent its protocol, optimized for the WAN. This protocol reinvention can be very simple, such as simply reordering message payloads, or much more complicated.

A good candidate for application optimization is Common Internet File System (CIFS), which Microsoft developed for Windows network file services as a variation of the Server Message Block (SMB) protocol. Over time, many versions of CIFS have been developed. Each new version of the Windows operating system includes a slightly expanded CIFS vocabulary. Currently, there are over 120 commands across the various Windows versions. When many CIFS operations are initiated, several handshaking operations must be performed, such as file permissions checking, file locking, and so on. Because these operations depend on each other, they are done serially. Each operation incurs delay across the WAN.

The CIFS application acceleration technology implemented in WAAS tries to group CIFS messages and proxy them on the optimization node on the WAN egress. This reduces the effects of the round-trip times. Additionally, the benefits of DRE, LZ compression, and other TCP stack optimizations are performed on the CIFS transactions.

6.3.3 Prepositioning

Prepositioning content where it is more convenient to the user (from a networking perspective is a very common form of optimization. Several different technologies can perform prepositioning, but the benefit comes from topological closeness. The ability to host and serve the content with minimal user interaction is also important. Ensuring that prepositioned content remains “fresh,” that is, that prepositioned content matches the content at the originating server is essential.

The L3 End Point Optimization (EPO) technology makes very frequent use of prepositioning. EPO uses DNS to send a user to an HTTP cache that is topologically closer to the user, and that is already prepositioned with the content.

Windows file services can use local file-shares that are already prepositioned with documents, images and so on. WAAS provides prepositioning services as part of its CIFS framework.

Many audio/video files are extremely large. The Cisco Application and Content Networking System (ACNS) product supports prepositioning videos and audio files, such as training material, during off-peak times, and then using the content during peak times. This reduces the WAN load and provides users with nearly instantaneous access to the content.

6.3.4 Stream Splitting Technologies

Video and audio streaming have interesting properties with respect to optimization in that multiple parties are many times interested at the same time in the same set of data. Additionally, some applications such as stock tickers and pre-positioning are trying to deliver the same data set to a distributed population.

6.3.5 Multicast

None of the technologies described so far require changing WAN functionality or features. In this sense, they work despite inefficiencies in the WAN network. Multicast technology is akin to a controlled broadcast in that data is transmitted once over a network, specifically to interested parties.

Looking at the L4 optimizations described in “Layer 4 Optimizations,” and at Figure 6-8 , you can see that unicast streams are optimized at the branch level. That is, for multiple receivers at a branch, only one stream is sent to the branch. However, from the perspective of the stream source (the hub site), multiple identical copies are sent to each branch. If hundreds of branches are interested in the stream, there will be hundreds of duplicate data streams that will each need bandwidth on the WAN link at the hub.

Multicast can also be used with prepositioning. As described earlier in this chapter, prepositioning places the same data at points closer to its users. One way to update distributed preposition destinations is to use individual unicast streams and update each site one by one. If this is done concurrently, the individual streams use a lot of bandwidth on WAN link of the origin server. If done serially, the prepositioned update can take a long time to complete. With a multicast enabled WAN, however, the origin server can update the prepositioned sites concurrently, and WAN bandwidth is used more efficiently.

A multicast enabled WAN enables the highly efficient duplication of the data to multiple sites. If the WAN link from the hub and the WAN itself is multicast-enabled, only a single stream of data must be sent from the hub site. The WAN duplicates the flow as necessary along the path to the sites. This not only reduces the outgoing bandwidth use at the hub, it also optimizes bandwidth use in the WAN.

Figure 6-8. Multicast-Enabled WAN

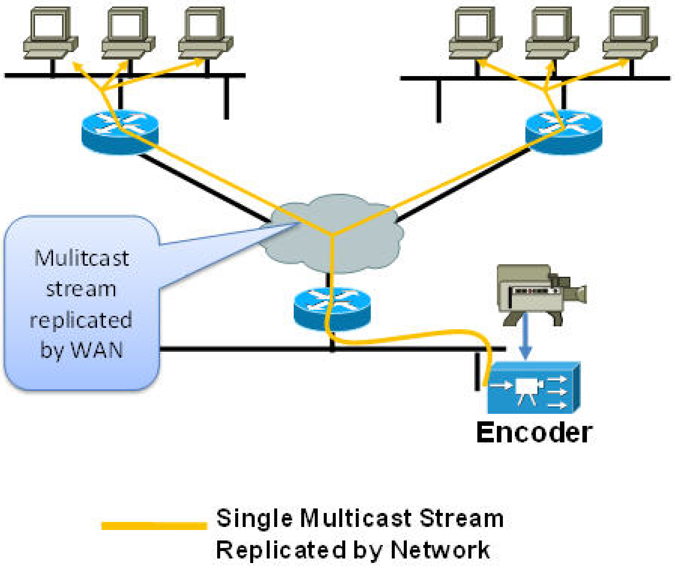
Mulitcast stream replicated by WAN
Encoder
Single Multicast Stream Replicated by Network

Multicast can also be useful when the WAN is not fully multicast enabled, but is capable of multicast transport. Generally, in a WAN that is only capable of multicast transport, the WAN edge router must replicate the packets. To create such a network, generic routing encapsulation (GRE) tunnels can be deployed between the hub and branch sites, or, in a FR WAN, for example, each branch site can have a PVC to the hub.

This also means that neither the WAN nor the hub will see bandwidth savings for multisite traffic. However, having the WAN edge router perform the replication removes this burden from the server. The server does not need to expend resources for duplicating packets, and those server resources are free for other tasks.

6.3.6 Multicast Translation and Unicast Stream Splitting

The one-to-many or many-to-many type of data distribution model maps perfectly to multicast. However, multicast is not enabled in many parts of enterprise networks, nor is multicast generally enabled in the Internet. This lack of multicast capability causes users in the same branch to receive the same audio or video stream from the streaming server, with each stream requiring its own bandwidth on the local branch, and on the WAN.

In this optimization case, the video/audio optimization device (Cisco ACNS) intercepts the streaming data request and only one request is sent over the WAN. The origin server therefore transmits only one stream to the site. After the data arrives from the WAN to the optimization device, the device can either multicast the stream locally to the interested parties, or replicate the unicast single stream of data into multiple unicast streams, once for each interested party.

This technique dramatically reduces the WAN bandwidth used by streaming video and audio. In fact, this technique might be required to enable multiple users to receive such streams; multiple receivers can easily use all the available bandwidth on the WAN access link.

This technology is also very useful in reducing the amount of streaming audio, such as Internet radio, that is received over an enterprise Internet access connection. It is possible; for example, that several hundred users are listening to the same Internet radio station and that each listener is eating up bandwidth on the Internet access link.

Figure 6-9. Optimizing Unicast Streams over the WAN

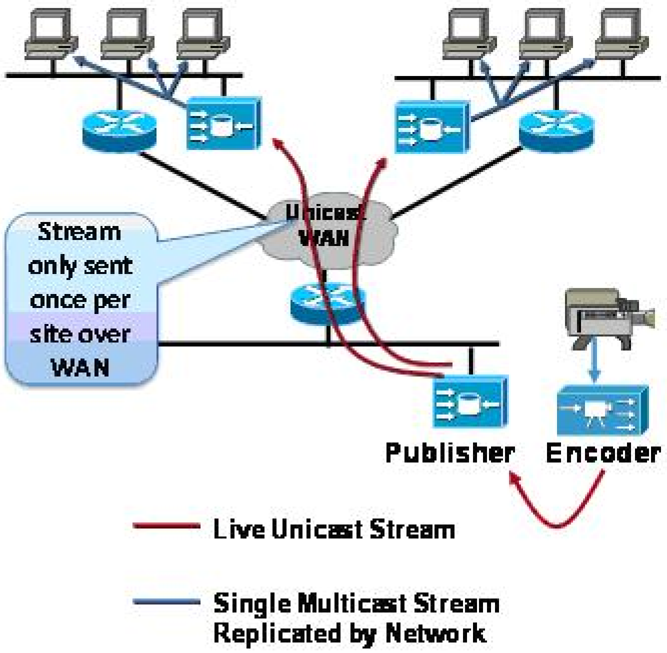
Unicast WAN
Stream only sent once per site over WAN
Publisher
Encoder
Live Unicast Stream
Single Multicast Stream Replicated by Network

7 Network Control Technologies

Different applications have different networking requirements. For example, real-time interactive voice requires a low latency network path, but does not require high bandwidth. Quality of Service (QoS) deterministically degrades network service quality to protect the user experience. In the case of voice, queuing delays can occur when there is congestion in a network path. Classifying voice, and placing it in a special queue that is serviced more often, enables the creation of a virtual path having lower latency. However, other traffic now must wait longer to be sent out a network interface. Such deterministic degradation is the fundamental aspect of control optimization. No additional bandwidth is created. Instead, a class of traffic is given special treatment based on its requirements and relative importance, as determined by the network operator.

7.1 QoS Requirements and Placement

As shown in Figure 7-1, the best place to apply QoS policy is on the branch WAN edge, which is usually the branch router WAN-facing interface. This interface has the following characteristics:

- Aggregation point for traffic
  - The WAN interface is often the first location at which many different streams are placed on the same wire. Merging this traffic- can cause contention between different traffic types. These different streams can come from many different paths: multiple interfaces on the branch edge device, optimized traffic from WAAS, and so on.
- Interface speed mismatches
  - The WAN edge interface can be the first point at which traffic must face congestion. In Figure 7-1, for example, the LAN interface might be 100MB, but the WAN facing interface could very well be a T1 (1.5Mb/s). This drastic difference in speed forces congestion and causes some traffic to be dropped.
- Features specific to access technologies
  - Some access technologies, such as Frame Relay, with forward explicit congestion notification/ backward explicit congestion notification (FECN/BECN), and cable-modems, with nondeterministic bandwidth) present unique challenges that only QoS awareness at branch edge router can address.

An operator can make relative judgments about how to treat specific traffic classes and apply QoS at the WAN edge router. This requires an awareness of all the traffic exiting a site toward the WAN link.

QoS features available in the WAN aggregation router, and WAAS network transparency, correctly identify and separate individual flows to support more granular optimization of application delivery and more effective differentiated services. QoS classifies different WAN flows by IP precedence bits, or by port levels present in their packet headers. QoS can make efficient use of WAN resources required by mission critical applications. QoS can ensure that the bandwidth and latency requirements for voice and multimedia applications are met, and that excess capacity is shared equally among lower priority applications.

Cisco routers provide an infrastructure for QoS features, including hardware based packet low latency queuing elements, shapers, and policers.

Figure 7-1. Applying QoS Policy at a WAN Congestion Point

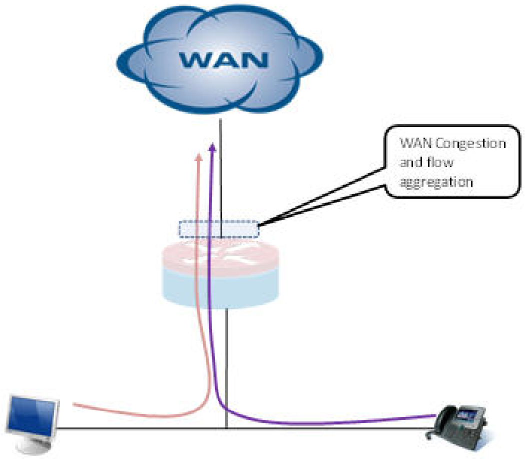

7.2 Cisco IOS QoS Model

QoS on IOS routers supports extremely granular management of valuable network and path services. QoS control functions follow the model shown in Figure 7-2.

Figure 7-2. Cisco IOS QoS Model

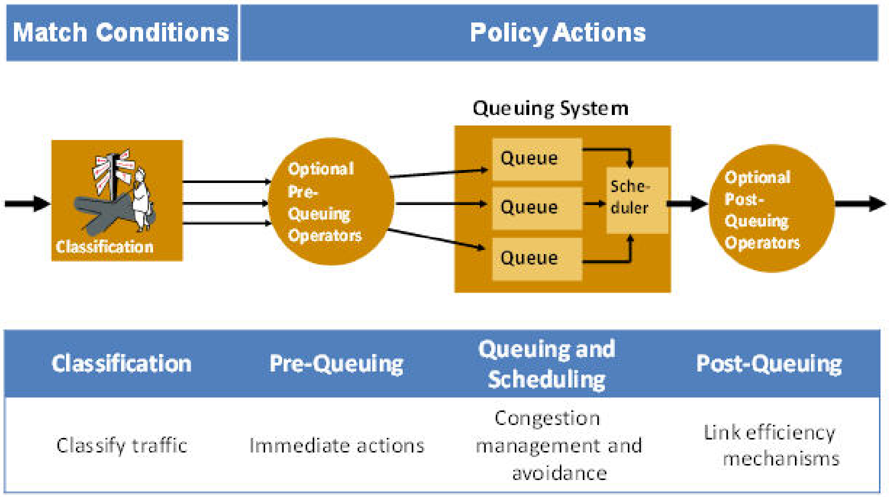

Classification	Pre-Queuing	Queuing and Scheduling	Post-Queuing
Classify traffic	Immediate actions	Congestion management and avoidance	Link efficiency mechanisms

QoS can be thought of as a collection of classification tools and actions on the classified traffic flows.

7.2.1 Classification

Classification has already been covered in chapter 4. The basic concept with classification in the context of QoS is to allow for selection of kinds of traffic for the purpose of further actions. As shown in Table 1 below, different kinds of applications ranging from interactive voice and email have radically different needs and requirements from the network. These requirements need to be well understood as well as their relationship with other network attributes. For example in the case of interactive transactional traffic loss of packets can add more delay due to retransmission. A classification system needs to be able to identify these different traffic classes by either explicit application identification or by keying off an agreed upon control point (such as DSCP) that has been previously set.

7.2.2 Prequeuing

After classification has been performed, there is an opportunity to filter specific traffic that has been identified as superfluous or unauthorized on the network. For explicit removal, this could be as simple as an access-list or a bit more complicated such as discarding specific traffic only when exceeding a rate of transmission. In rate based removal, a QoS traffic policer could be configured. Additionally, a Denial of Service (DoS) attack can be constrained by placing drop access-lists towards the source, or if the exact match criteria have not been explicitly discovered or identified, the traffic can be rate-limited.

Table 1. Different Applications have different needs

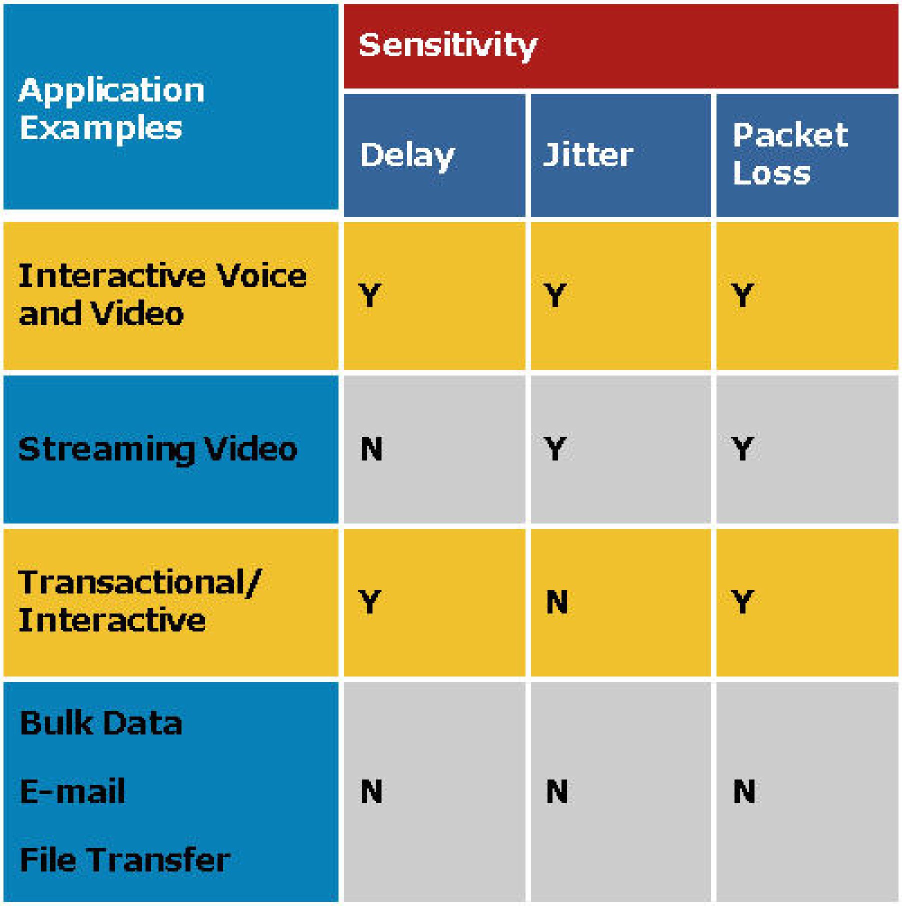

Application Examples	Sensitivity		
	Delay	Jitter	Packet Loss
Interactive Voice and Video	Y	Y	Y
Streaming Video	N	Y	Y
Transactional/ Interactive	Y	N	Y
Bulk Data E-mail File Transfer	N	N	N

7.2.3 Queuing and Scheduling

The queuing and scheduling portion of QoS is often the most powerful with the ability to shape (shaping), prioritize (low-latency queuing), provide a level of session fairness (fair-queuing) among many other features. Given elastic and inelastic traffic types, the queuing and scheduling technologies can create artificial bandwidth limitations (shaping) as well as signal the end host (e.g. dropping the packet or Explicit Congestion Notification (ECN)) such that it reduces the rate of transmission.

7.2.4 Postqueuing

As the packet is scheduled to exit the router, there are is an opportunity to do a few last actions on it. The other items in the model refrain from modifying the actual packet. They either drop or delay the packet, but do not modify the payload or the frame itself. These last bits of post-queuing functions act on the packet. For example, for RTP packets, header compression can be performed that can reduce a significant percentage of the bandwidth space that the packet was using. Large packets on slow links can cause excessive jitter as sending a large packet takes up a solid chunk of time that prevents other real-time traffic streams from being transmitted.

Post-queuing features such as multi-link ppp link fragment interleave (MLPPP-LFI) and IMA (inverse multiplex for ATM) break the large packets into smaller packets, and interleave smaller packets requiring low-latency from other traffic streams in between. In addition, post-queuing allows “bandwidth fairness” to be applied on a per-user or even at the level of per-application-per-user level of traffic management.

7.2.5 Congestion Management and Avoidance

When no congestion exists, packets are sent out the WAN interface as soon as they arrive. However, when packets are arriving at a faster rate than the router can forward them-- congestion occurs. The Congestion Management QoS features can be enabled to address congestion conditions.

Queuing algorithms are used to sort and prioritize traffic overflow before packets are transmitted onto the network. Packets are scheduled for transmission according to their assigned priority and queuing mechanism applied to the WAN interface. There are different types of queuing mechanisms such as First-In-First-Out (FIFO), Priority Queuing (PQ), and Weighted-Fair-Queuing (WFQ). Traffic flows classified as expedite forwarding will be placed in the high priority queue and transmitted first. Traffic flows requiring a certain amount of bandwidth will be sorted and placed on different queues before being transmitted. In times of congestion, some of these queues will overflow and excess traffic will be dropped.

Selective packet drop can be implemented with Weighted Random Early Detection (WRED). By assigning a drop probability to a traffic class, the operator can select which traffic class will drop packets at the time of congestion. In general, a higher drop probability is given to traffic classes considered to be of lower priority. WRED is one of the algorithms used by the Congestion Avoidance feature in IOS.

WRED also supports the Explicit Congestion Notification (ECN) bits of the IP header of packets as defined in RFC3168. The ECN bits indicate whether the endpoints of a traffic flow support the feature. If they do, WRED can mark the packet via ECN and forward the packet instead of dropping. The ECN bits will indicate to the sender that a congestion condition is present and that they should reduce their windows—effectively reducing the data transmission rate. If the endpoints do not support ECN, the packets will be processed by the WRED algorithm and depending on traffic conditions, the packets may be dropped. Traffic optimized by the WAAS TFO feature automatically supports ECN.

7.2.6 Integrated Services and RSVP

In addition to marking packets and generalized queues, that traffic can be scheduled in there is a more strict system for reserving and enforcing bandwidth available in Cisco IOS. Known as IntServ, the RSVP protocol is used to reserve resources in the network for streams of traffic by applications to allow for guaranteed QoS.

RSVP reservations can be especially effective for inelastic traffic types, such as many types of video, or, more generally, UDP based applications, which do not respond to congestion in the network path. In the previous section, ECN was mentioned as a mechanism for relaying to the end host the need reduce the traffic rate. Packet drops and ECN do not apply to inelastic traffic, and an explicit reservation of bandwidth is used to protect the quality of the traffic. If multiple applications try to reserve bandwidth for a specific queue type and the reservation cannot be met, the network denies the reservation, which prevents a subpar application experience and protects existing reservations from degradation.

7.2.7 Modular QoS CLI (MQC)

The QoS tools described in this chapter can be applied in IOS platforms using Modular QoS CLI (MQC). MQC provides a standard set of CLI commands and QoS semantics that supports a simple, straightforward way to configure QoS features. MQC uses definitions of class-maps, policy-maps, and service policies to apply QoS controls to traffic going through a router. MQC policies share a common language across IOS platforms, and are easily understood and rapidly deployed to ensure consistent behavior across a network.

A class-map defines the criteria for a traffic class. Traffic matching specific criteria is grouped into a traffic class, and is processed according to the traffic policy defined in the policy-map. The traffic policy is attached to an interface so that all traffic going through the interface in a particular direction is subject to the policy.

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

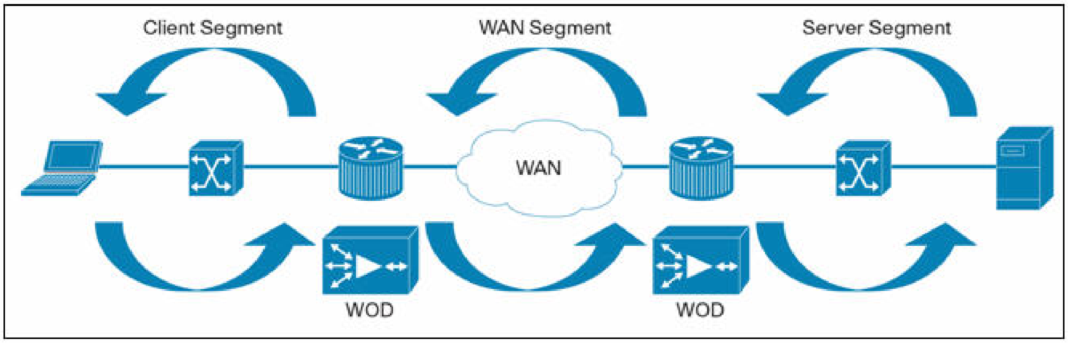

Fortunately, Cisco WAN optimization technologies are architected to help overcome the transparency issues found in competing products. For example, Cisco Wide Area Application Services (WAAS):

- 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.

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

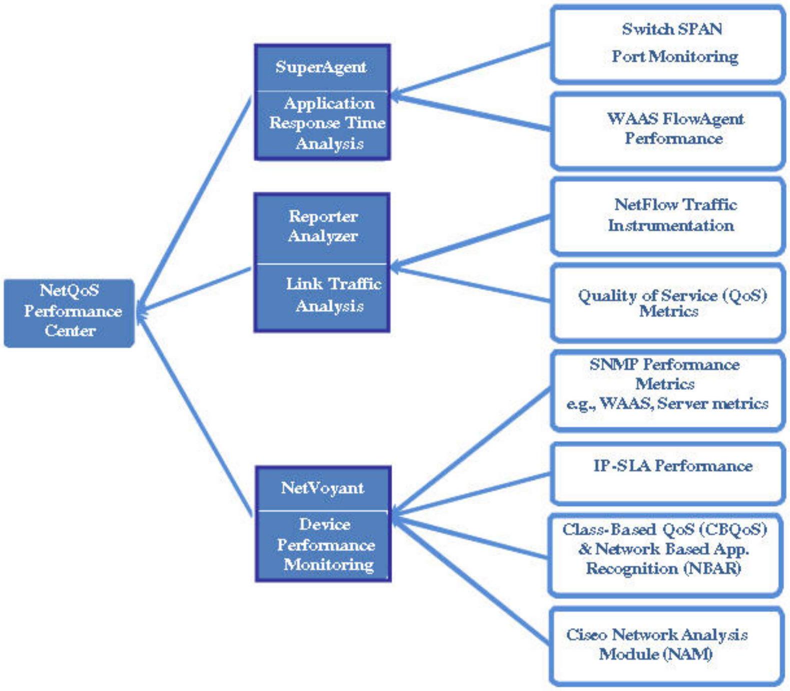
NetQoS Performance Center
SuperAgent
Application Response Time Analysis
Reporter Analyzer
Link Traffic Analysis
NetVoyant
Device Performance Monitoring
Switch SPAN Port Monitoring
WAAS FlowAgent Performance
NetFlow Traffic Instrumentation
Quality of Service (QoS) Metrics
SNMP Performance Metrics e.g., WAAS, Server metrics
IP-SLA Performance
Class-Based QoS (CBQoS) & Network Based App. Recognition (NBAR)
Ciseo Network Analysis Module (NAM)

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

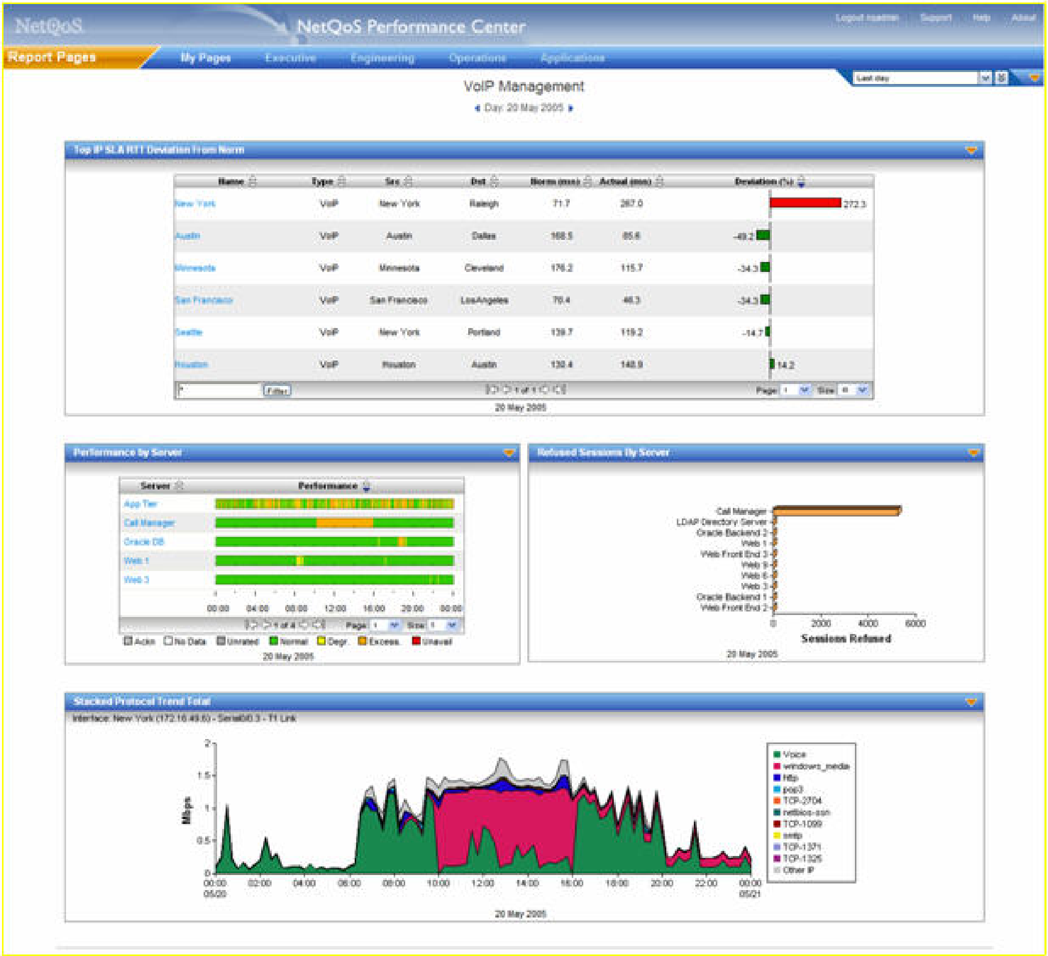

VoIP Management

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

20 May 2005

8.2.1.1 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

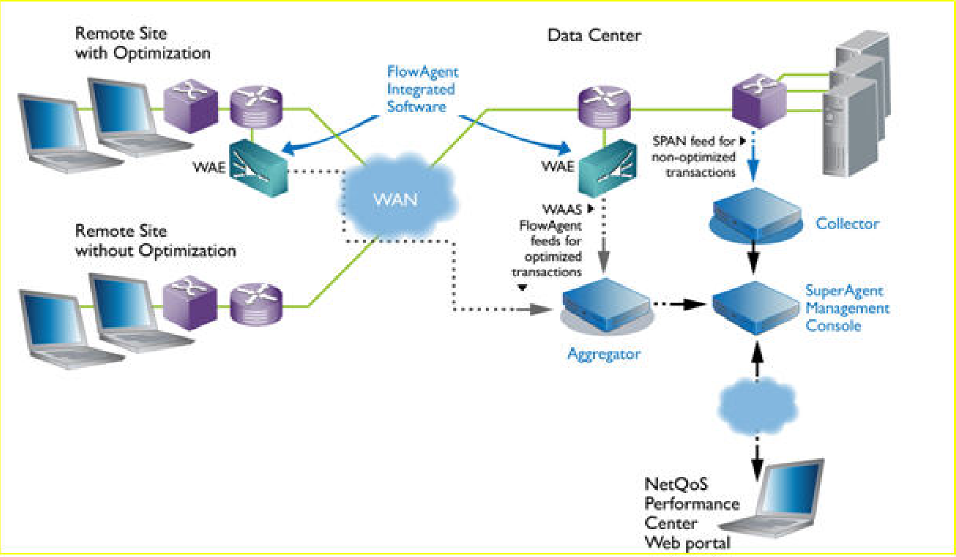
Remote Site with Optimization
Data Center
FlowAgent Integrated Software
WAE
WAE
SPAN feed for non-optimized transactions
WAN
Collector
Remote Site without Optimization
WAAS FlowAgent feeds for optimized transactions
SuperAgent Management Console
Aggregator
NetQoS Performance Center Web portal

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

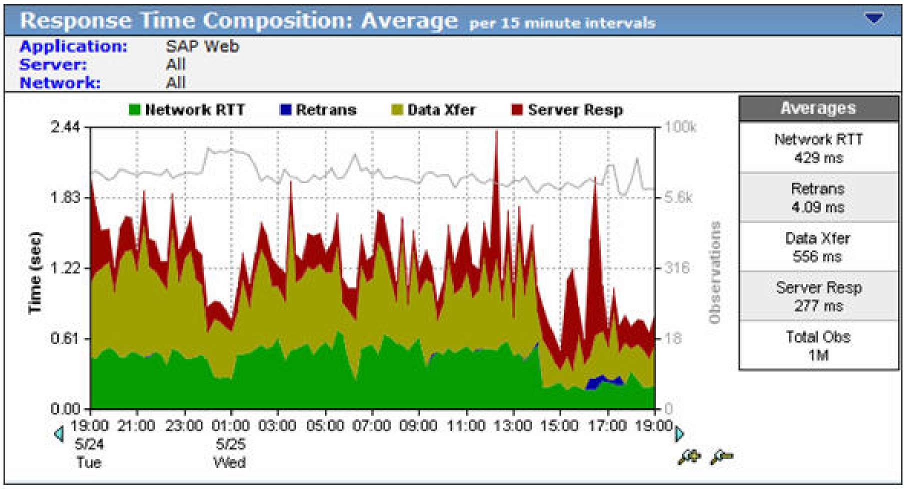

Response Time Composition: Average per 15 minute intervals

Application: SAP Web
Server: All
Network: All

Averages

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

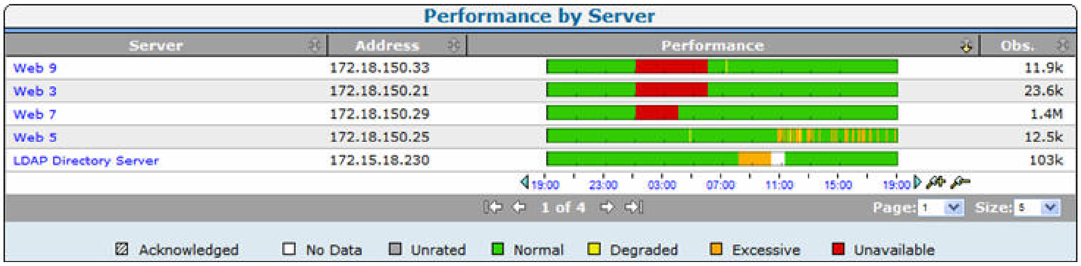

Performance by Server

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

1 of 4 Page: 1 Size: 5

▨ Acknowledged □ No Data ▩ Unrated 🟩 Normal 🟨 Degraded 🟧 Excessive 🟥 Unavailable

- 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

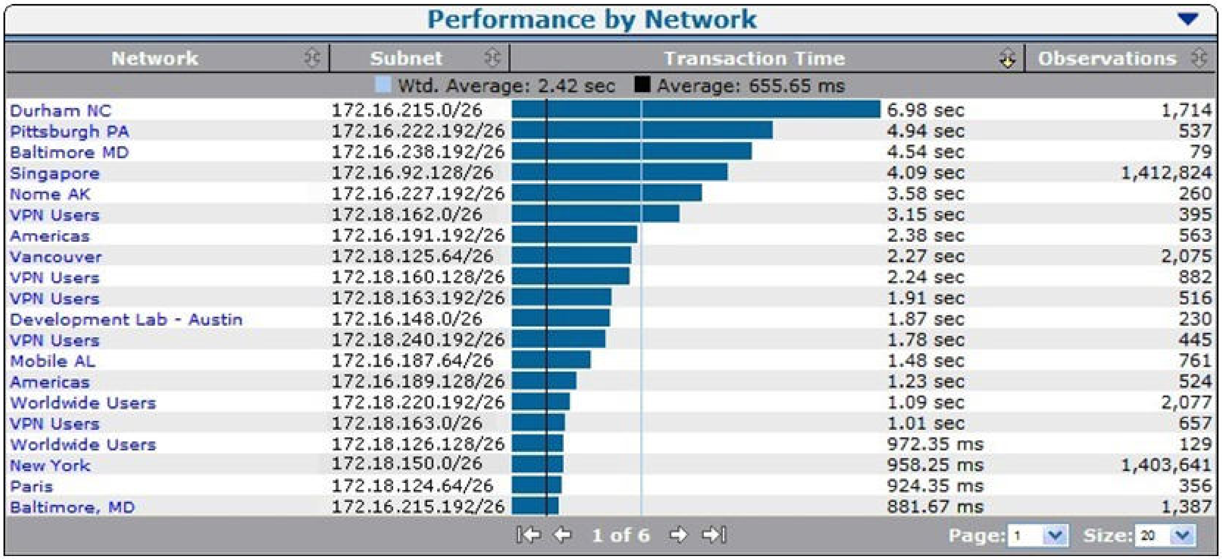

Performance by Network

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

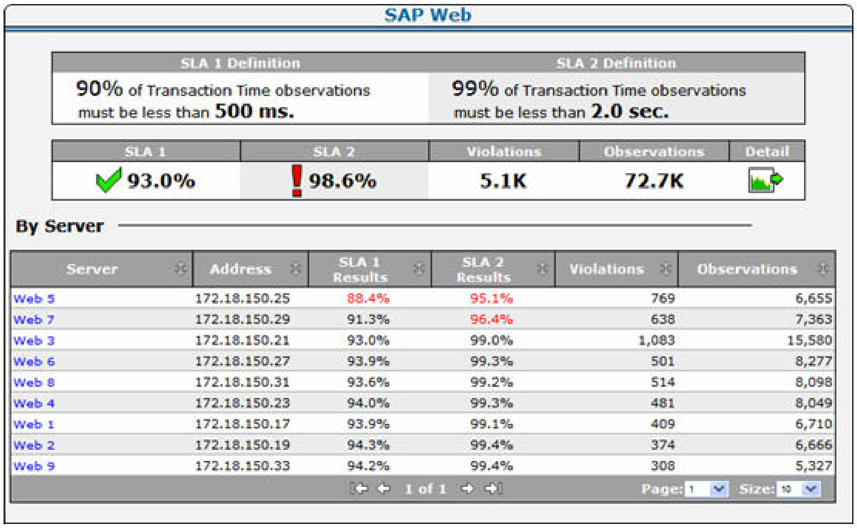

SAP Web

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

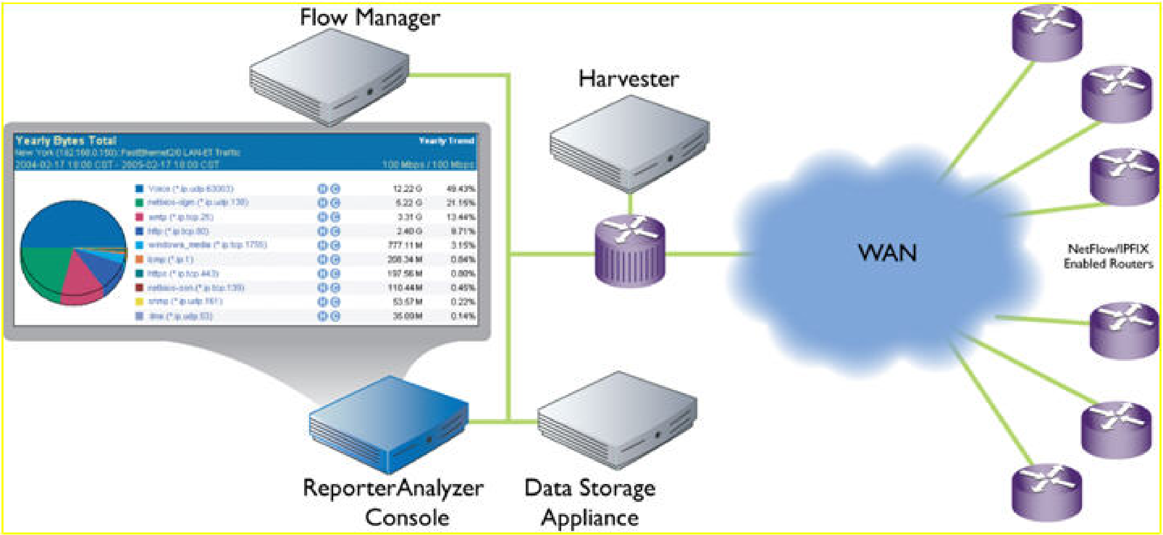

Yearly Bytes Total
New York (192.168.0.150): FastEthernet2/0 LAN-E1 Traffic
2004-02-17 18:00 CDT - 2005-02-17 18:00 CDT

Yearly Trend
100 Mbps / 100 Mbps

Voice (*ip.udp.63003)	12.22 G	49.43%
netbios-dgm (*ip.udp.138)	5.22 G	21.15%
smtp (*ip.tcp.25)	3.31 G	13.44%
http (*ip.tcp.80)	2.40 G	9.71%
windows_media (*ip.tcp.1755)	777.11 M	3.15%
icmp (*ip.1)	208.34 M	0.84%
https (*ip.tcp.443)	197.56 M	0.80%
netbios-ssn (*ip.tcp.139)	110.44 M	0.45%
snmp (*ip.udp.161)	53.57 M	0.22%
dns (*ip.udp.53)	35.09 M	0.14%

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

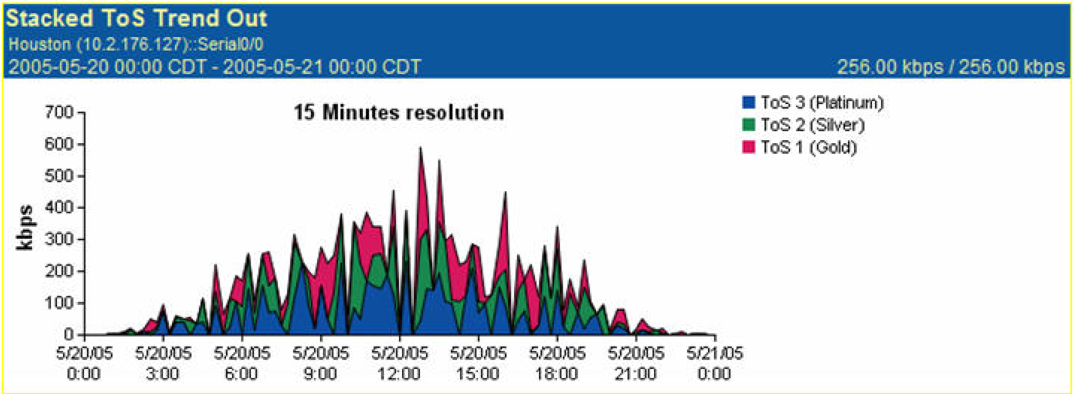

- Custom reports (see the example shown in Figure 8-11), which use a wizard to create interface, protocol, ToS, host, conversation, and combination reports for any timeframe, and that can be run at any designated time.

Figure 8-11. ReporterAnalyzer Custom Report

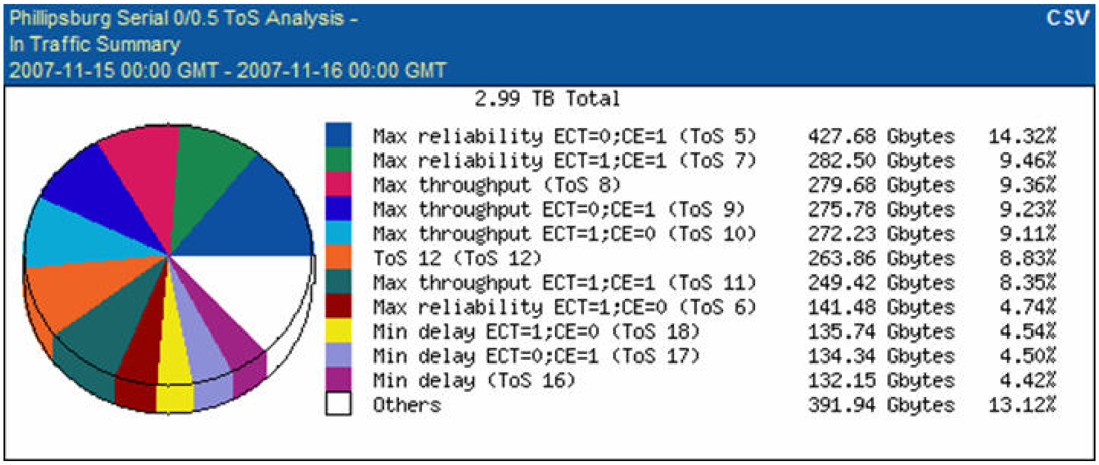
Phillipsburg Serial 0/0.5 ToS Analysis -
In Traffic Summary
2007-11-15 00:00 GMT - 2007-11-16 00:00 GMT
CSV

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

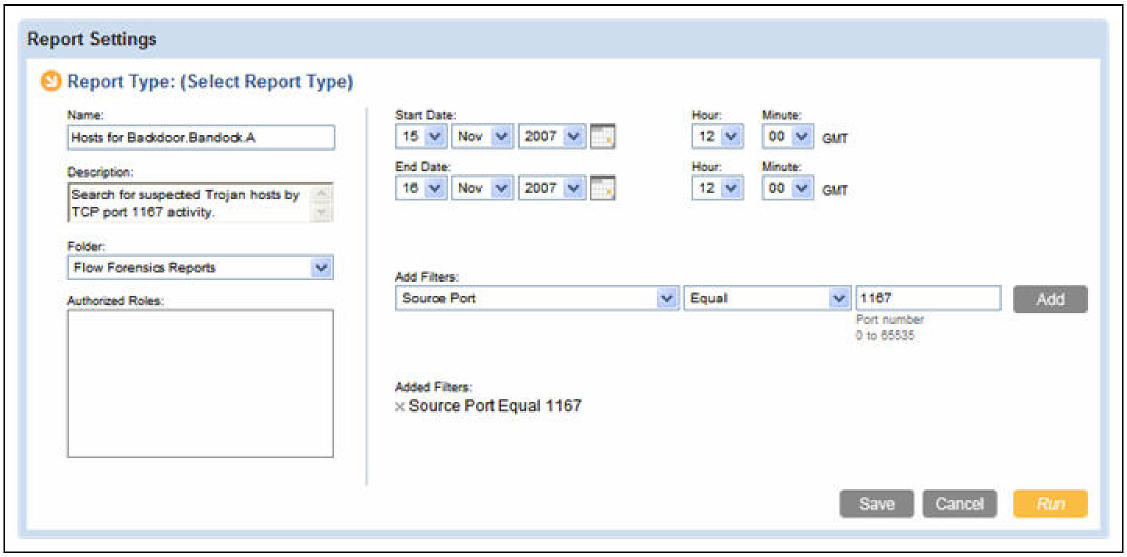

Report Settings

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

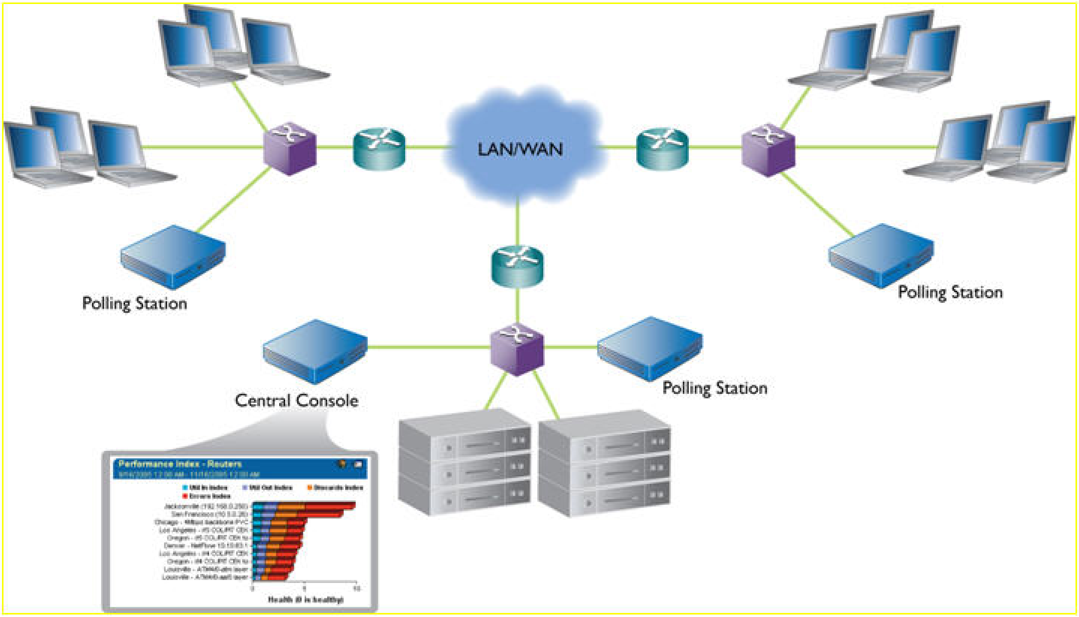

In a distributed configuration, a NetVoyant Master Console performs administration and reporting while remote Polling Stations discover and poll devices on the network. Because the volume of data collected from polling stations is much smaller than the volume of received SNMP data, on large networks 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

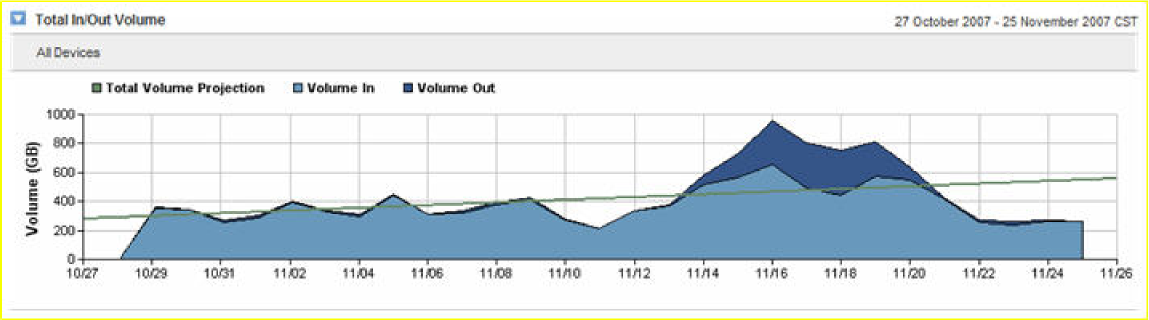

- Capacity planning reports (see the example shown in Figure 8-15) support planning for device and network link upgrades and help identify the resources with the fastest growth and characteristics approaching defined thresholds.

Figure 8-15. NetVoyant Capacity Planning

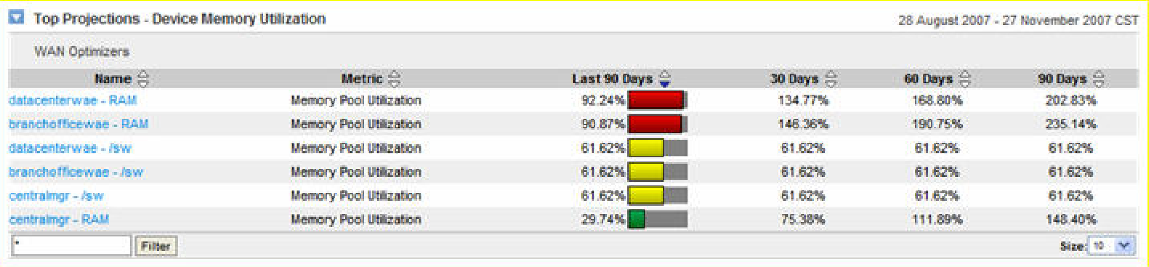

Top Projections - Device Memory Utilization

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

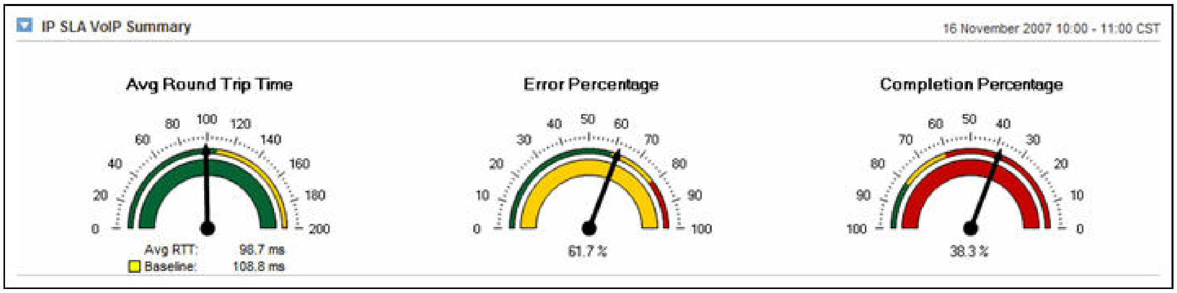

IP SLA VoIP Summary

16 November 2007 10:00 - 11:00 CST

Avg Round Trip Time

Avg RTT: 98.7 ms

🟨 Baseline: 108.8 ms

Error Percentage

61.7 %

Completion Percentage

38.3 %

- 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

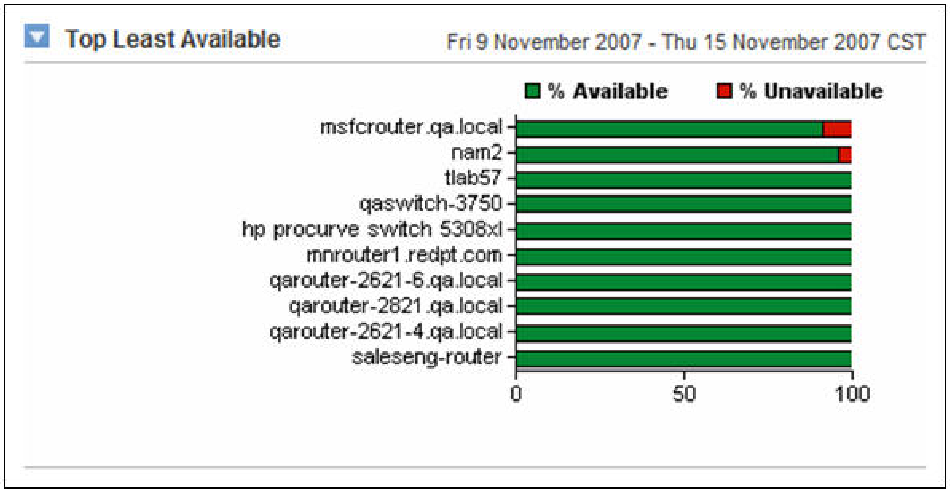

8.5 NAM: Granular Monitoring and Troubleshooting

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

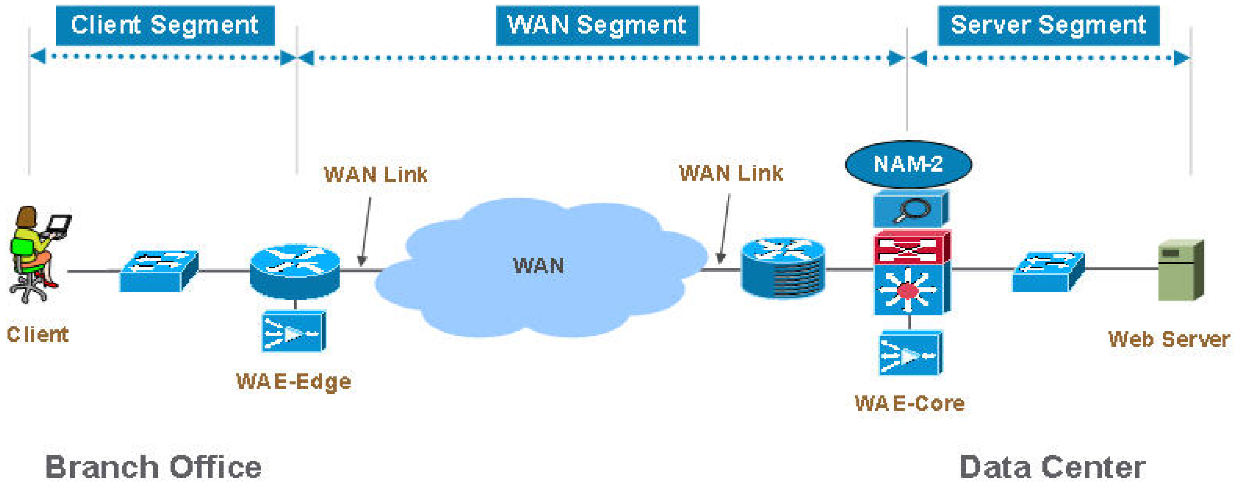

Cisco NAM is available as

- 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

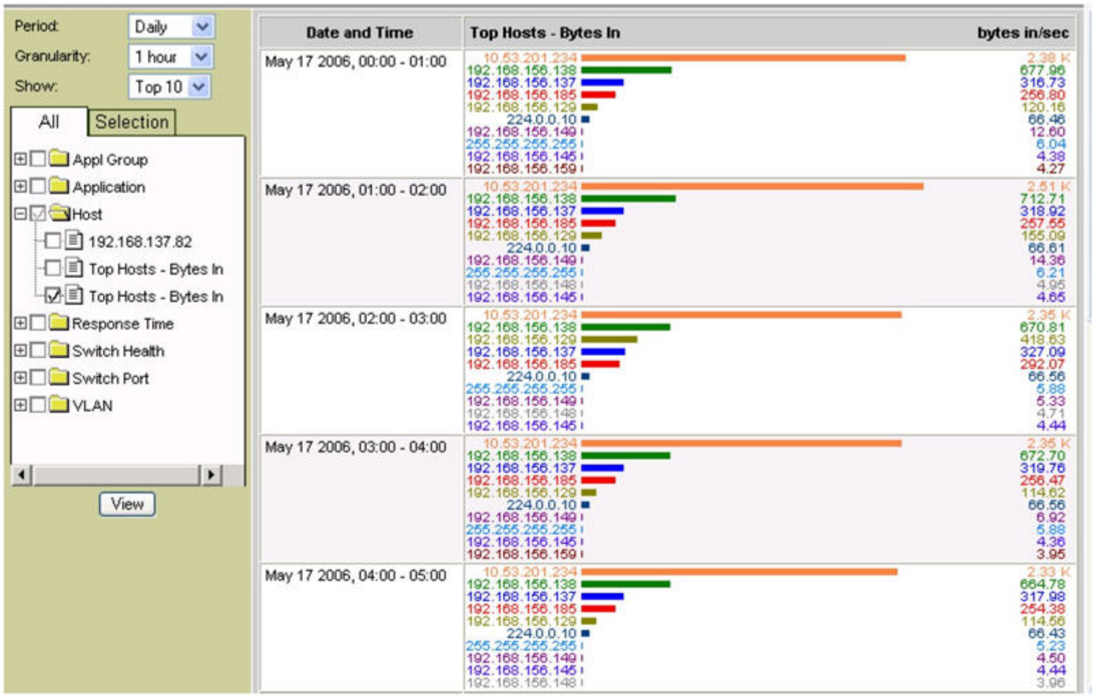

Period: Daily

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	256.80
	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

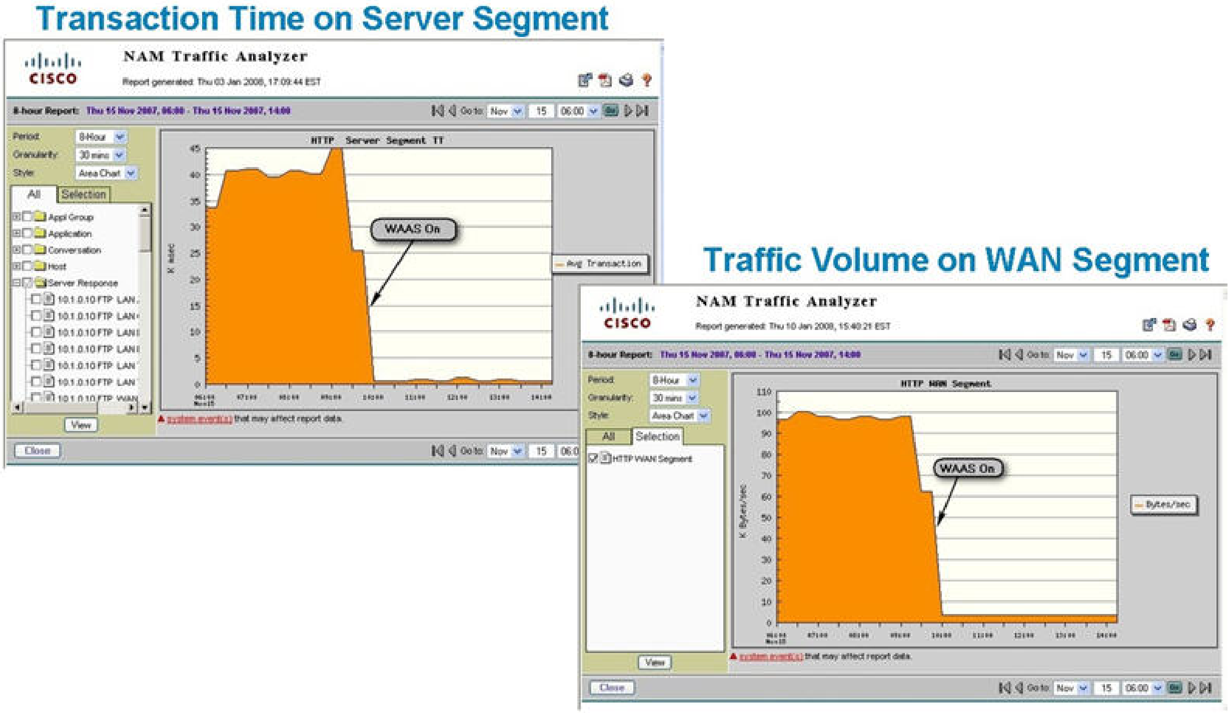

Transaction Time on Server Segment

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

All | Selection

- Appl Group
- Application
- Conversation
- Host
- Server Response
  - 10.1.0.10 FTP LAN
  - 10.1.0.10 FTP LAN
  - 10.1.0.10 FTP LAN
  - 10.1.0.10 FTP LAN
  - 10.1.0.10 FTP LAN
  - 10.1.0.10 FTP LAN
  - 10.1.0.10 FTP WAN

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

All | Selection

☑ 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

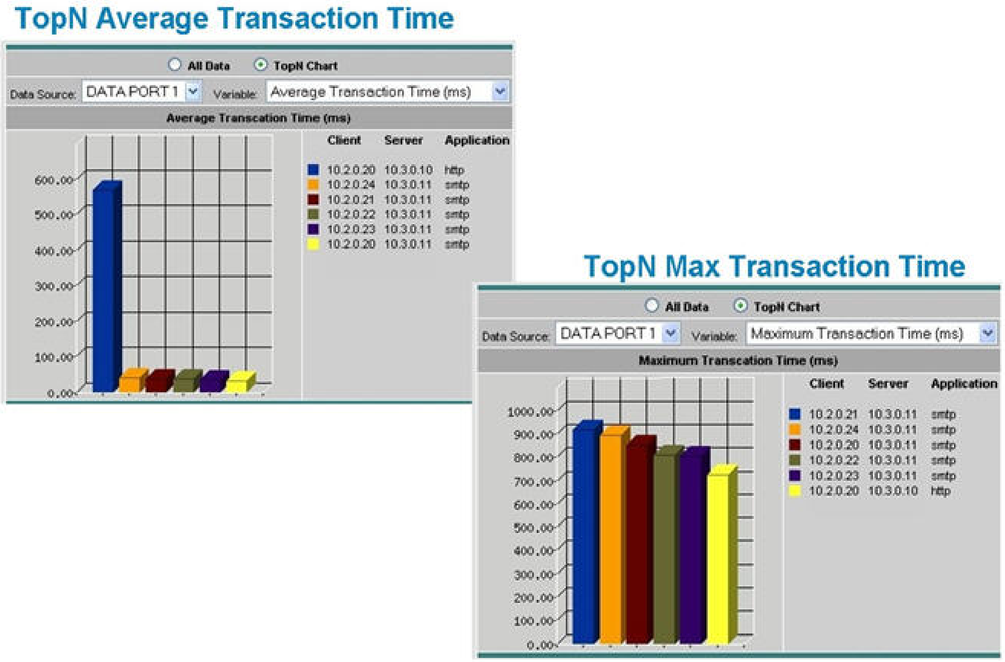

TopN Average Transaction Time

○ All Data ● TopN Chart

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

Average Transcation Time (ms)

	Client	Server	Application
🟦	10.2.0.20	10.3.0.10	http
🟧	10.2.0.24	10.3.0.11	smtp
🟥	10.2.0.21	10.3.0.11	smtp
🟩	10.2.0.22	10.3.0.11	smtp
🟪	10.2.0.23	10.3.0.11	smtp
🟨	10.2.0.20	10.3.0.11	smtp

TopN Max Transaction Time

○ All Data ● TopN Chart

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

Maximum Transcation Time (ms)

	Client	Server	Application
🟦	10.2.0.21	10.3.0.11	smtp
🟧	10.2.0.24	10.3.0.11	smtp
🟥	10.2.0.20	10.3.0.11	smtp
🟩	10.2.0.22	10.3.0.11	smtp
🟪	10.2.0.23	10.3.0.11	smtp
🟨	10.2.0.20	10.3.0.10	http

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

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

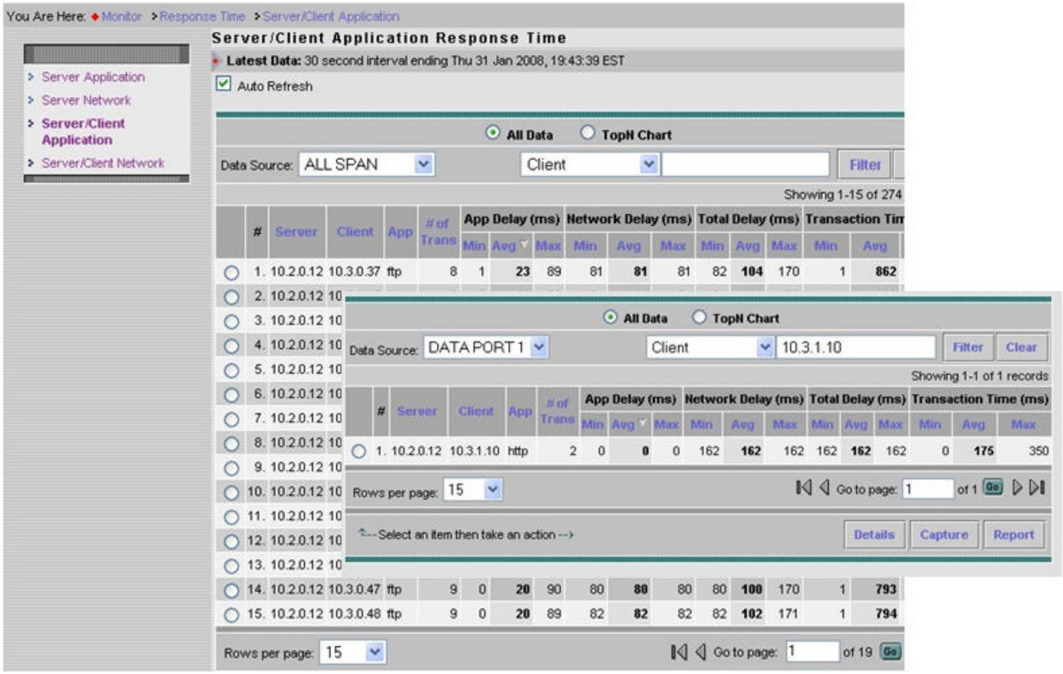

You Are Here: ◆ Monitor > Response Time > Server/Client Application
Server Application
Server Network
Server/Client Application
Server/Client Network

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 Filter

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 Go

● All Data ○ TopN Chart

Data Source: DATA PORT 1 Client 10.3.1.10 Filter Clear

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 Go

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

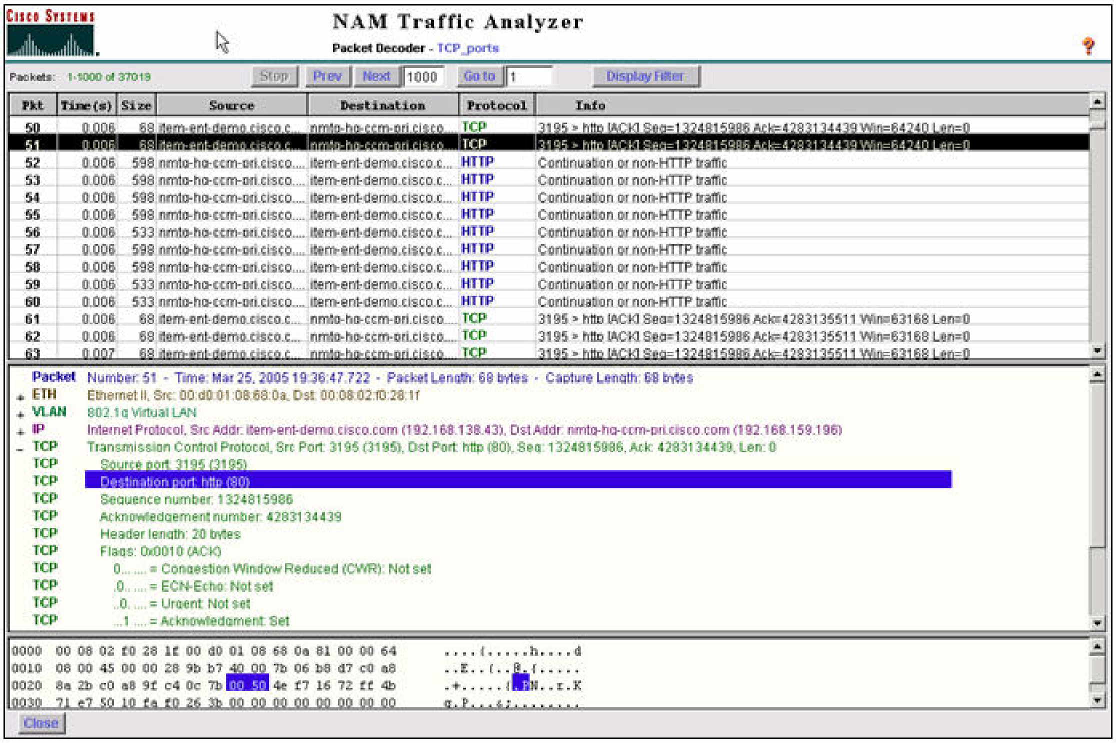

Packets: 1-1000 of 37019 | 1000 | 1

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: nmtq-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

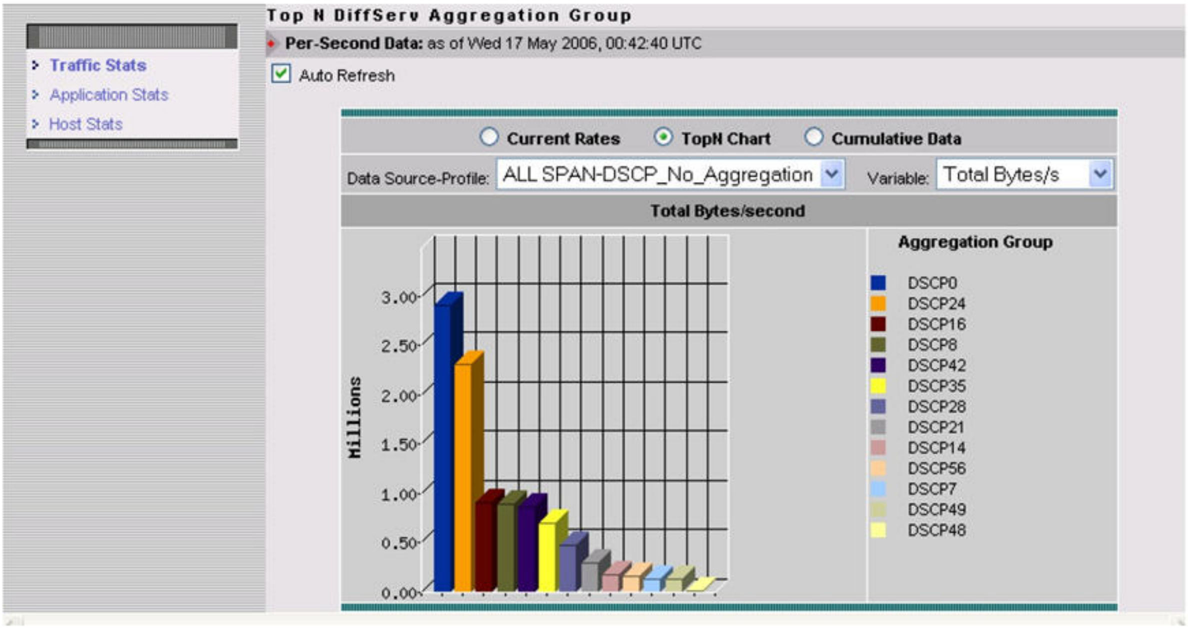

Top N DiffServ Aggregation Group

Per-Second Data: as of Wed 17 May 2006, 00:42:40 UTC

☑ Auto Refresh

○ 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 to www.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

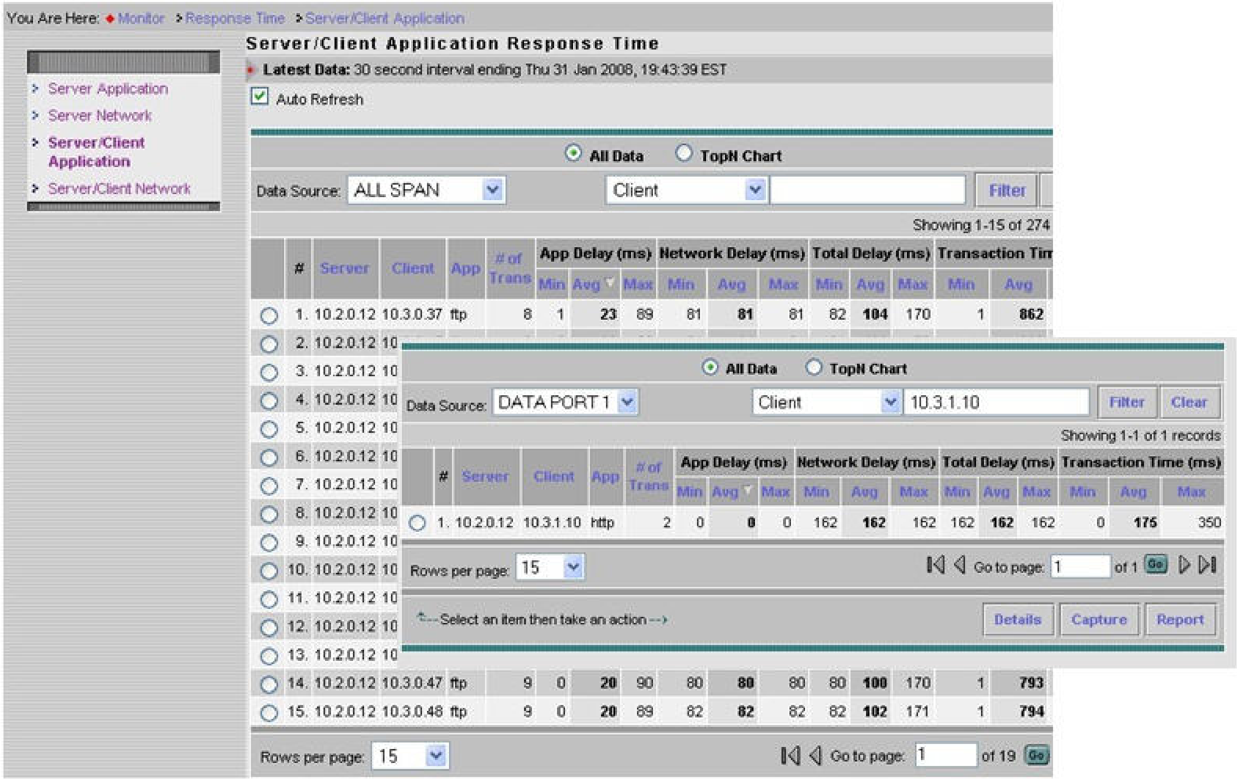
You Are Here: ◆ Monitor > Response Time > Server/Client Application
> Server Application
> Server Network
> Server/Client Application
> Server/Client Network

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 | Filter

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 Go

● All Data ○ TopN Chart

Data Source: DATA PORT 1 | Client | 10.3.1.10 | Filter | Clear

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 Go

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

9 Branch Design Considerations

In today's global economy, companies are rapidly expanding their presence throughout the world. As they grow, so do their networks. Their branch offices must have effective and reliable network connectivity to their corporate HQs (and between branch offices) in support of business applications. For any branch network, the designer must consider the following key requirements:

- Resiliency
- Security
- Network and application performance
- Load sharing

This chapter considers different technologies that can address these requirements, and introduces sample branch designs.

9.1 Resiliency/High Availability

Providing uninterrupted network connectivity between branches and headquarters, and among branches, is critical to any network design. It is essential to avoid single points of failure.

The cost of business interruption caused by network failure, the probability of component or network device failure, and other factors must be carefully analyzed Designers should consider the cost-benefit ratio of resiliency and plan for appropriate resiliency. In very small branch offices, the cost of providing resiliency may not provide adequate cost benefits.

9.2 Security

Security is another critical factor in any network design. Typically, data transfer between branches and corporate headquarters happens over a service provider (SP) network or over the Internet. Data that traverses the Internet is highly susceptible to snooping.

To protect against snooping and provide business confidentiality, some form of data encryption should be deployed. Additionally, many government regulations, such as Health Insurance Portability and Accountability Act (HIPAA) and Sarbanes-Oxley (SOX) mandate securing traffic.

9.3 Network and Application Performance

As global businesses evolve, some are deploying more and more business critical applications that can provide split-second decisions. Network performance, thus, is a key consideration for providing such applications with fast connectivity and bandwidth.

9.4 Load Sharing

A good branch design also provides resources for scaling and growth. Load sharing can provide both improved availability and scalability to the network. In a branch network, load sharing can be implemented by having multiple WAN exit routers to share the load. In addition, branch routers can have multiple exit interfaces connected to the WAN for load sharing. These methods can improve both availability and scaling.

9.5 Common Branch Topologies

Although in many cases a branch might have overlapping features from the following profiles, a typical branch with WAN optimization can be categorized into one of the most widely deployed topologies.

9.5.1 Single Tier Branches

These branches typically have one branch router with a WAE (either internal or external) for optimizing traffic. This router typically has an inbuilt switching module so that multiple same subnet end hosts and IP phones can be connected without requiring an external switch.

In a typical small-enterprise/SOHO deployment, the branch routers are usually single-homed, as shown in Figure 9-1 (a). In small enterprise branches with few users, shown in Figure 9-1 (b), the branch router is often dual homed to two SPs (or to the same SP), providing resiliency for the WAN connection. A variation of this design, shown in Figure 9-1 (c), adds another WAE to add WAE resiliency. The second WAE can be used to share bandwidth and application optimization loads.

Figure 9-1. SOHO and Single Tier Branches

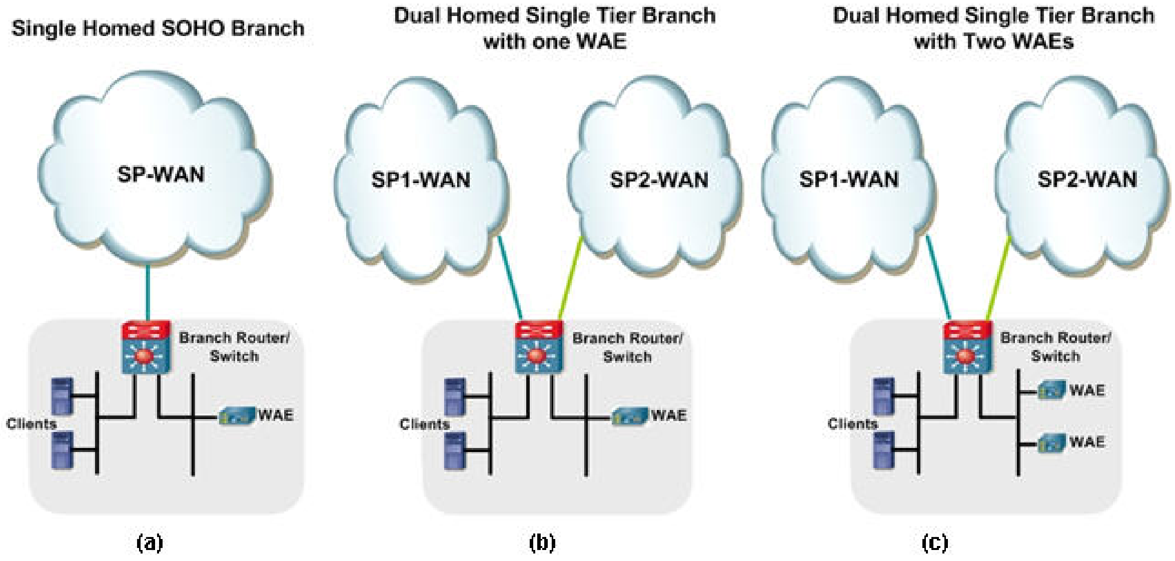

9.5.2 Dual Tier Branches

Dual tier branches employ multiple (usually two) branch routers and an external Ethernet bridge providing connectivity to the two routers. Typically, both routers are homed to different SPs to provide more resilient WAN connectivity.

A router dual homed to WAN, or a dual branch router design enables efficient link utilization with load sharing, but can also increase the possibility of asymmetric routing and associated challenges. Asymmetric routing is explained in the following section.

The routers can share one WAE for optimization, as shown in Figure 9-2 (a). Alternatively, the routers can use multiple WAEs in a load-shared and redundant manner as shown in Figure 9-2 (b). WCCP can locally correct asymmetry issues and keep optimized session on a consistent WAE.

Figure 9-2. Dual Tier Branches

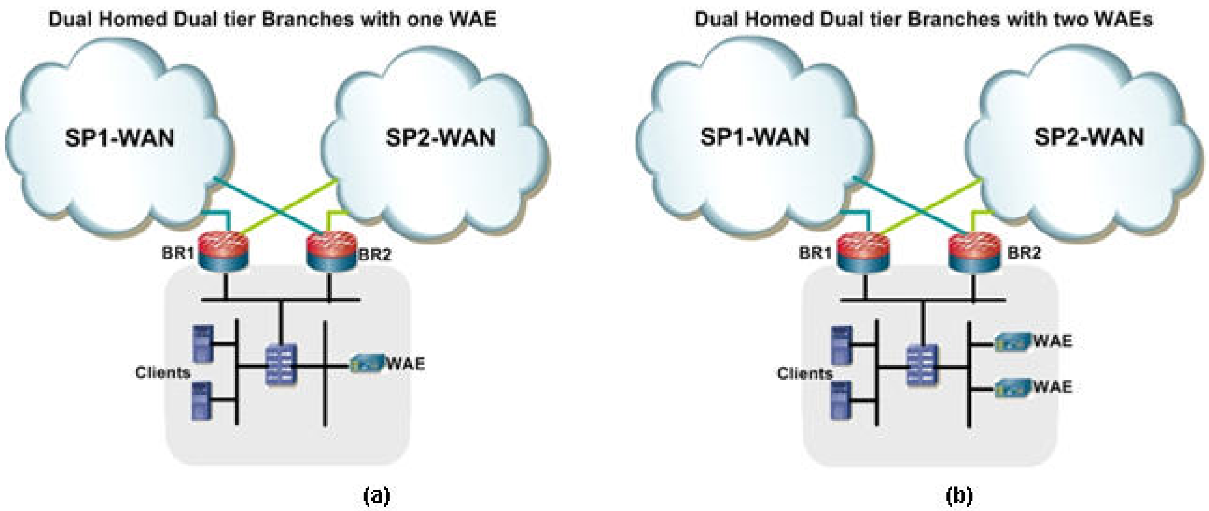
Dual Homed Dual tier Branches with one WAE
SP1-WAN
SP2-WAN
BR1
BR2
Clients
WAE
(a)
Dual Homed Dual tier Branches with two WAEs
SP1-WAN
SP2-WAN
BR1
BR2
Clients
WAE
WAE
(b)

9.5.3 Asymmetric Routing

In hop-by-hop packet routing systems (the vast majority of IP routing systems), each router independently selects the outgoing path. While routing protocols ensure that loops are avoided, the symmetry of bidirectional traffic flows cannot be guaranteed when destination prefixes are reachable by multiple paths.

Asymmetric routing occurs when traffic does not traverse the same path in both directions of a conversation. A branch site that has multiple WAN connections, with either a single router or multiple routers, is always prone to asymmetric routing. This can occur because the routing protocol on each end selects a different path, by load or session balancing, or even by path optimization mechanisms.

Figure 9-3. Asymmetric Routing

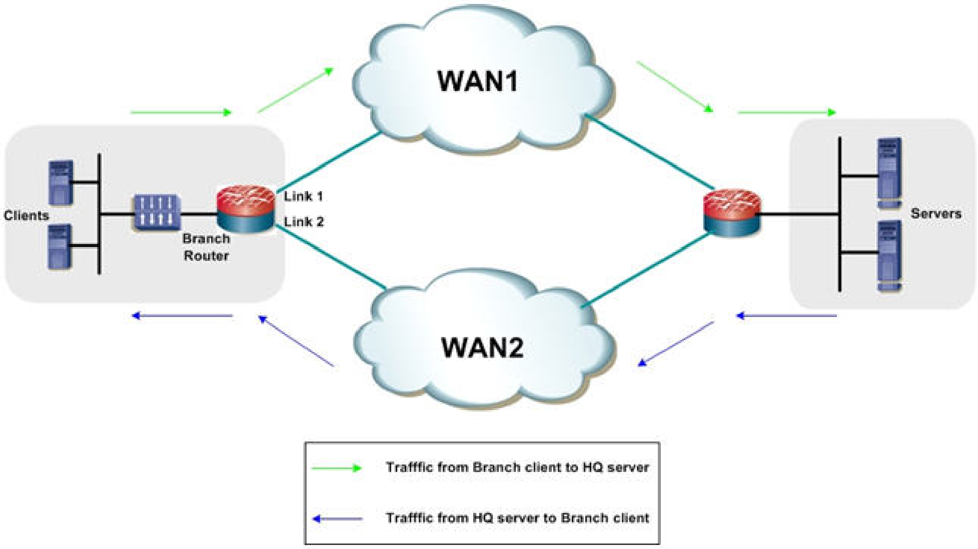
WAN1
Clients
Branch Router
Link 1
Link 2
Servers
WAN2
Trafffic from Branch client to HQ server
Trafffic from HQ server to Branch client

Consider a simple branch network with one router dual-homed to the WAN, as shown in Figure 9-3, with a bidirectional traffic session between a client on the branch LAN and a headquarters server. As Figure 9-3 shows, traffic from the branch to headquarters can exit through one WAN link on the branch router (in this case, Link l). However, the server-side routers calculated the best path for the client network to be over the WAN2 link. This is the asymmetric routing case: server-sourced traffic arrives on Link 2 of the branch router.

Without proper design and placement of network services, asymmetric routing can create challenges in networks. For example, asymmetric routing results in suboptimal TCP performance; TCP assumes that the SYN from one end and the ACK from other end traverse the same path. Because data does not traverse the same physical path in both directions, suboptimal TCP performance results.

If state is built into network services for transiting traffic, full flow information might not be available to a network device. Network services that need to act on (or simply see) both directions of a conversation include firewalls, NAT devices, and stateful identification in some applications. Some monitoring tools also rely on bidirectional traffic flows.

Such tools have potential issues when traffic exits through one router and ingresses through the other. Challenges caused by asymmetric routing in a network with a multi-homed single router are a little easier to mitigate than designs where two routers are being used, because conversation state for both directions still transits the same router in the single router case.

9.5.4 Branch LAN-Side High Availability

When using two branch routers, Gateway Load Balancing Protocol (GLBP), Hot Standby Routing Protocol (HSRP), or Virtual Router Redundancy Protocol (VRRP) can be configured on the LAN interfaces of the routers to provide high availability for the default gateway of an end host. Table 9-1 lists the advantages of HSRP and GLBP.

Table 9-1. HSRP and GLBP Advantages

Routing Protocol	Advantages

HSRP	**Higher availability:** Enhanced redundancy can eliminate single point of failure of the first-hop gateway. Enhanced object-tracking can be used with HSRP/GLBP to help ensure the redundant implementation mirrors network capabilities. Enhanced object tracking enables multiple technologies such as HSRP, GLBP, and Virtual Router Redundancy Protocol (VRRP) to track the same object and each take different actions. For more information, refer to http://www.cisco.com/en/US/docs/ios/12_4t/ip_appl/configuration/guide/taipbtrk.html **Simpler access-layer design:** More efficient use if resources is possible without configuring additional VLANs and subnets.
GLBP	**Automatic load balancing:** Off-net traffic is shared among available gateways on a per-host basis, according to the defined load-balancing algorithm. **Lower administration costs:** Because all hosts on a subnet can use a common default gateway while load balancing is still achieved, administering multiple groups and gateways is unnecessary. **Efficient use of network resources:** Multiple paths upstream from the gateways can be used simultaneously.

9.5.5 Branch WAN-Side High Availability

To ensure business continuity, it is essential to provide high availability on the WAN side. Therefore, branch routers are typically dual-homed to the SP network on the WAN side. In such cases, one path serves as the primary path, and the other is a secondary or backup path.

In other words, traffic normally takes the primary path. If the primary path fails, traffic can move to the secondary path. Routing protocols deployed over the WAN network provide reachability information that enables the branch routers to decide on the path to be taken.

However, when a path fails, it takes some time for routing protocols to converge and point to the alternate path, causing packet loss. Although routing protocol timers can be tweaked to alleviate some of these effects, this is not always fully effective. Additionally, some types of failures (brownouts, black holes, path congestion, and so on) are not captured by the routing protocol reachability information. Performance Routing (PfR), detailed in “Path Optimization Using PfR,” can be deployed to provide faster, efficient, more granular optimization per application path.

Figure 9-4 depicts a typical GLBP/HSRP implementation on the LAN side and dual homing on the WAN side on branch networks.

Figure 9-4. Typical Branch LAN/WAN High Availability

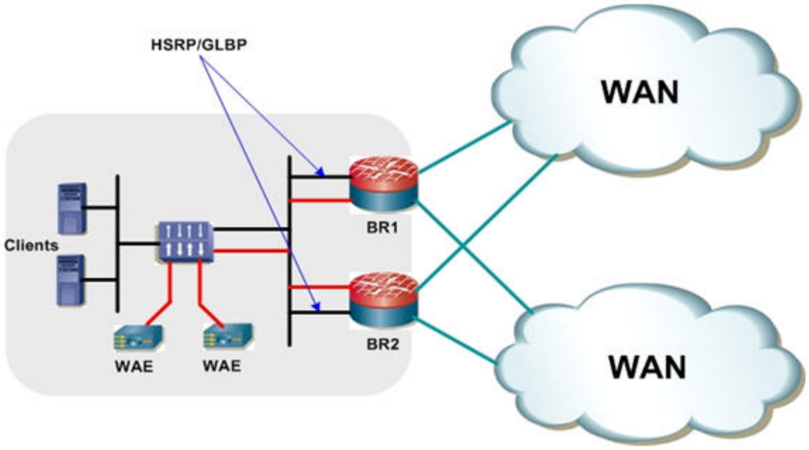

9.6 Optimization Tools

This section describes the major tools used in the Cisco WAN and application optimization solution:

- NBAR
- QoS
- NetFlow
- PfR
- Web Cache Communication Protocol (WCCP) and WAEs

9.6.1 Application Visibility Using NBAR

To use WAN bandwidth efficiently, the user should be able to implement appropriate application-aware traffic policies. When using WAEs for TCP optimization, the user faces challenges; for example, L5 through L7 details are no longer visible after packets are optimized (especially with compression). This is shown in Figure 9-5.

Figure 9-5. TCP Optimization and Application Visibility

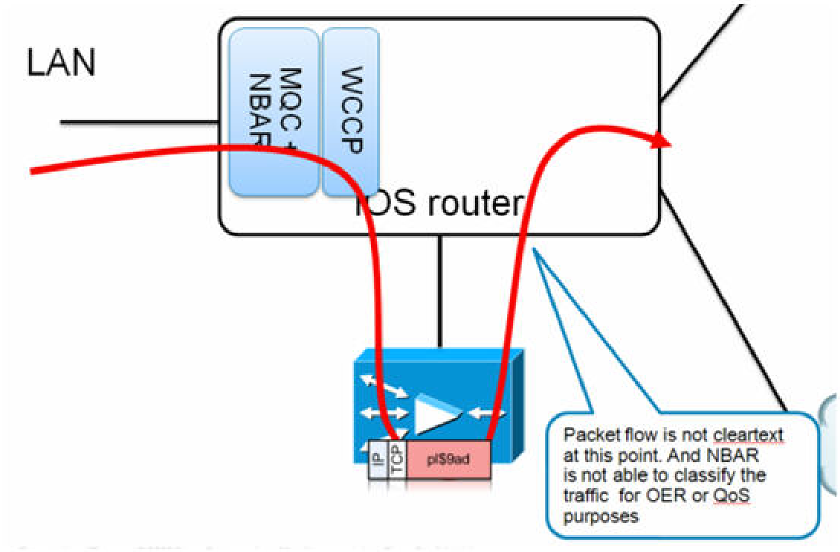

Packet flow is not cleartext at this point. And NBAR is not able to classify the traffic for OER or QoS purposes

Figure 9-6. NBAR Application Marking with TCP Optimization

This challenge makes implementing any application-aware traffic policies difficult. In such cases, using NBAR (as shown in Figure 9-6), with deep packet inspection (DPI) capabilities, and using QoS to mark traffic before WAAS, can be an ideal tool for classifying application traffic.

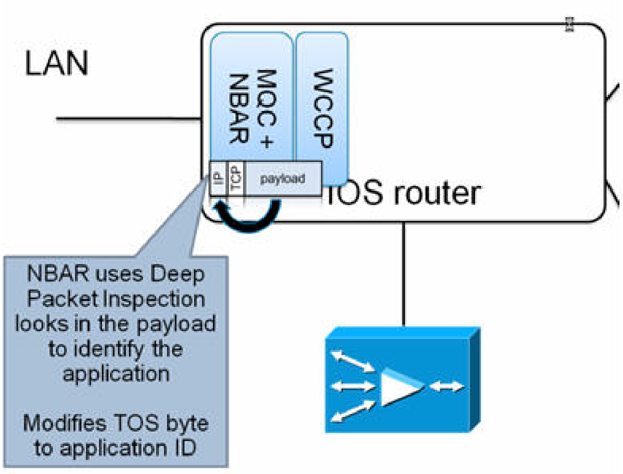
LAN
MQC + NBAR
WCCP
IP
TCP
payload
IOS router
NBAR uses Deep Packet Inspection looks in the payload to identify the application
Modifies TOS byte to application ID

NBAR has stateful packet inspection capabilities that complement NetFlow, which collects flow information only up to the TCP or UDP port level. Therefore, NetFlow is limited in identifying some kinds of traffic. NBAR can be applied on branch router LAN interfaces to inspect and mark traffic using Differentiated Services Code Point (DSCP) or type of service (ToS) values as traffic enters the branch router.

Different applications of interest can be marked with different DSCP settings. For example, the user can configure specific policies to mark HTTP traffic based on URLs. Because the WAEs do not change IP header information, appropriate QoS egress policies can be configured on the branch router WAN interfaces of the to condition optimized traffic from WAE. A sample NBAR classification configuration follows.

```
!CONFIGURE APPROPRIATE NBAR CLASS MAPS TO MATCH DIFFERENT APPLICATION/PROTOCOLS
class-map match-any NBAR_FTP
 match protocol ftp
class-map match-any NBAR_voice
 match protocol sip
 match protocol skinny
class-map match-any NBAR_HTTP
 match protocol secure-http
 match protocol http url "*cisco.htm*"
 match protocol http url "*ba_HN.jpg*"
```

9.6.2 Congestion Management Using QoS

Congestion management using QoS can play a key role in regulating WAN traffic. In a branch, WAN traffic might comprise real-time traffic, Web traffic, and other application traffic. Each traffic patterns requires different handling for delay, packet loss, and so on. For example, real-time traffic such as voice requires low latency and jitter, and appropriate low latency queues should be configured to provide priority for this traffic.

Other traffic types, such as TCP traffic, can have huge packet sizes that can delay other queued traffic, especially on slow links, such as T1. These traffic types require configuration with some link efficiency mechanism, such as fragmentation interleaving, to help minimize the effect. Even with optimization by WAE, congestion can occur and the user must configure appropriate QoS mechanisms to help ensure that different traffic classes are guaranteed appropriate bandwidth.

Similarly, traffic to and from the WAE can be prioritized. If necessary, traffic policing and shaping can be applied on the WAN/LAN interfaces to ensure adherence to traffic limits.

9.6.3 NetFlow

NetFlow supports accounting of IP flows traversing and terminating on the router. This data can be exported to a NetFlow collector and analyzer, which can generate reports, ranging from departmental usage and capacity planning to understanding, which applications are running on the network. NetFlow can also support early detection of security threats, such as denial of service (DoS) attacks. On the router itself, the NetFlow data can be a powerful troubleshooting tool.

Figure 9-7 illustrates some charting examples from a specific product (NetQoS ReporterAnalyzer). This product is described in more detail in 9.3, “Network and Application Performance.”

Seven unique fields (called key fields) identify a NetFlow flow: router source interface, source IP address, destination IP address, type of service (ToS) byte, IP protocol, L4 source port, and L4 destination port.

Additional fields are collected, such as destination interface (where the flow was forwarded), bytes sent, and number of packets sent. Because the deepest flow accounting is based on the L4 ports (in TCP, the source and destination ports), in certain cases it is impossible to identify the end application in use.

To address this problem in many cases, identify an L4 port range and IP addresses to help in identify the application. For example, voice traffic might use UDP port numbers between 16384 and 32767, and IP phones might be restricted to a certain IP address range. The identification data can be represented in the NetFlow collector, which can then use this information to name the application “voice.”

Figure 9-7. NetQoS NetFlow Analysis

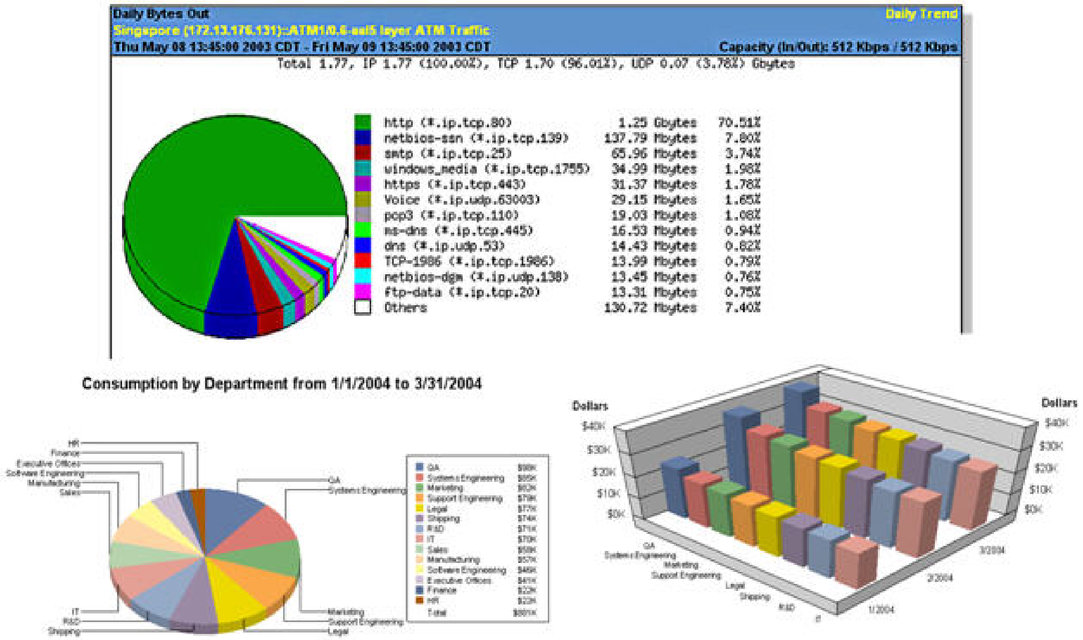

Daily Bytes Out
Singapore (172.13.176.131)::ATM1/0.6-aal5 layer ATM Traffic
Thu May 08 13:45:00 2003 CDT - Fri May 09 13:45:00 2003 CDT

Daily Trend
Capacity (In/Out): 512 Kbps / 512 Kbps

Total 1.77, IP 1.77 (100.00%), TCP 1.70 (96.01%), UDP 0.07 (3.78%) Gbytes

🟩	http (*.ip.tcp.80)	1.25 Gbytes	70.51%
🟦	netbios-ssn (*.ip.tcp.139)	137.79 Mbytes	7.80%
🟥	smtp (*.ip.tcp.25)	65.96 Mbytes	3.74%
■	windows_media (*.ip.tcp.1755)	34.99 Mbytes	1.98%
🟪	https (*.ip.tcp.443)	31.37 Mbytes	1.78%
■	Voice (*.ip.udp.63003)	29.15 Mbytes	1.65%
■	pop3 (*.ip.tcp.110)	19.03 Mbytes	1.08%
🟩	ms-dns (*.ip.tcp.445)	16.53 Mbytes	0.94%
🟦	dns (*.ip.udp.53)	14.43 Mbytes	0.82%
🟥	TCP-1986 (*.ip.tcp.1986)	13.99 Mbytes	0.79%
■	netbios-dgm (*.ip.udp.138)	13.45 Mbytes	0.76%
🟪	ftp-data (*.ip.tcp.20)	13.31 Mbytes	0.75%
□	Others	130.72 Mbytes	7.40%

Consumption by Department from 1/1/2004 to 3/31/2004

■	QA	$98K
■	Systems Engineering	$85K
■	Marketing	$82K
■	Support Engineering	$79K
■	Legal	$77K
■	Shipping	$74K
■	R&D	$71K
■	IT	$70K
■	Sales	$58K
■	Manufacturing	$57K
■	Software Engineering	$46K
■	Executive Offices	$41K
■	Finance	$22K
■	HR	$22K
	Total	$901K

Figure 9-8 shows these technologies and where they are applied.

Figure 9-8. NetFlow, NBAR, QoS at a Branch Router

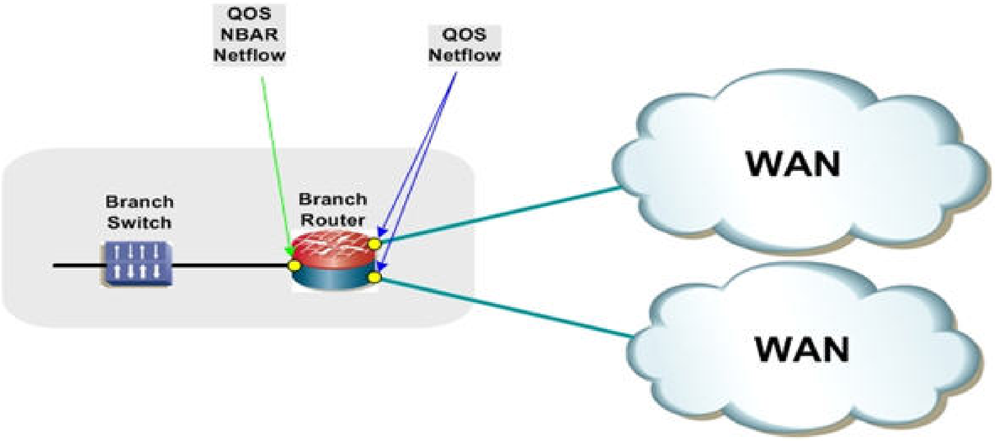

9.6.4 Path Optimization Using PfR

In any WAN design, the challenge is to provide application-specific bandwidth and resiliency while making efficient use of WAN bandwidth. While resiliency and some load sharing of WAN links can be ensured with routing protocols, traditional routing protocols do not provide best path selection based on link utilization or many of the other path characteristics. Therefore, the branch network might not be able to respond to dynamic network emergencies like intermittent congestion or a link failure upstream that does not cause a routing protocol transition. Even if the routing protocol does respond, it might take quite some time for the network to re-converge. In such cases, Cisco's PfR can be used to respond to such failures dynamically.

PfR helps to provide dynamic route optimization based on latency, link utilization, and dynamic detection of data path failure. PfR enables a user to:

- Improve network performance
- Optimally distribute load among multiple links
- Save costs through more intelligent bandwidth utilization
- Reduce operating expenses (OpEx) through automatic performance optimization
- Integrate directly with Cisco IOS® IP routing, Cisco IOS NetFlow, Cisco IOS IP Service Level Agreements (IPSLA) and other Cisco IOS Software features

Cisco PfR implementation requires at least one border router (BR) and master controller (MC) process, both of which run on IOS routers. The BR has the branch exit interfaces (attached to the WAN) and is responsible for collecting information about traffic exiting the site, and for implementing policy decisions derived by the MC. For a small office, home office (SOHO), the MC and BR can be the same device.

PfR also requires at least two exits towards the WAN. Figure 9-9 represents a typical PfR implementation with a dual-homed branch network. In this case, efficiently deploying PfR with NetFlow, NBAR, and QoS can enable the operator to maximize network performance and end user experience for different applications.

Figure 9-9. SOHO Deployment

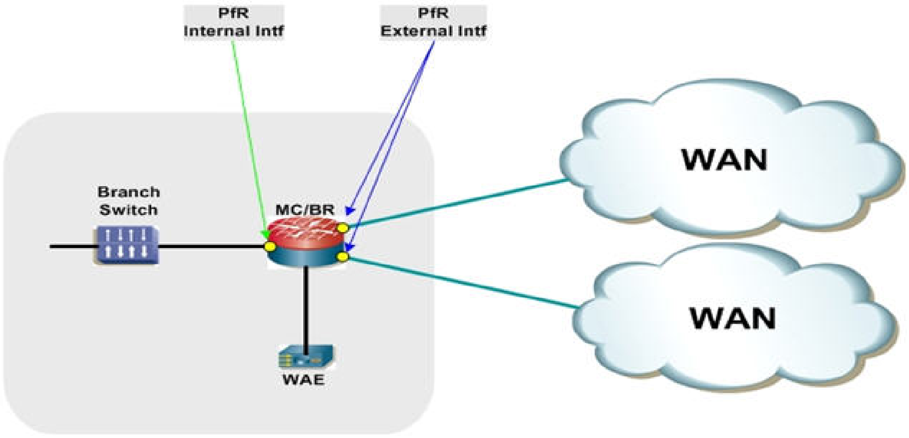

However, most deployments have two BRs and one MC. Note that Cisco Express Forwarding (CEF) must be enabled on PfR routers. Figure 9-10 depicts such a deployment.

Figure 9-10. PfR Deployment with dual Branch Routers

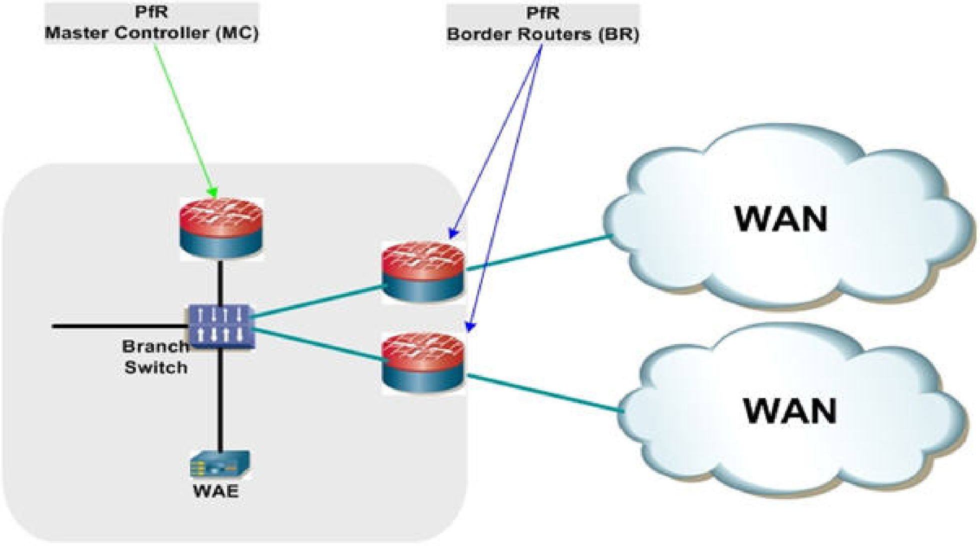

9.7 How PfR Works

The BR on a PfR-enabled network monitors the performance of the traffic going out to the WAN and relays this information to the MC. The MC verifies whether the current performance conforms to the configured policy for each traffic class. If not, the MC instructs the BRs to change the route for that particular prefix.

To perform route control, a route is injected. The injected route can be a static route, a BGP route (if BGP is used), or a Policy Based Routing (PBR) policy (if the forwarding decision is based on something more than the IP destination). Support for other routing protocols, such as Enhanced Interior Gateway Routing Protocol (EIGRP), is under development. Note that the injected routes are local and should not be advertised.

A typical PfR MC/BR configuration geared toward load balancing traffic between two serial links follows:

```
!CONFIGURE THE OER KEY TO BE USED FOR MASTER-BORDER AUTHENTICATION
key chain oer-key
 key 1
   key-string WANOPT
!
!CONFIGURE PfR MASTER
oer master
 max-range-utilization percent 10
 logging
 !
!POINT TO THE BORDER ROUTER AND OER KEY TO USE FOR AUTHENTICATION
 border 10.0.0.173 key-chain oer-key
!DEFINE THE INTERNAL AND EXTERNAL INTERFACES ON THE BORDER ROUTER
  interface FastEthernet3/1 internal
  interface Serial6/1:0 external
   max-xmit-utilization percentage 50
  interface Serial6/0:0 external
   max-xmit-utilization percentage 50
 !
 !CONFIGURE THE PfR MASTER TO LEARN TRAFFIC CLASS OR PREFIXES
learn
!CONFIGURE THE PfR TO LEARN TOP PREFIXES BASED ON THROUGHPUT
  Throughput
!CONFIGURE THE PfR TIMERS
  periodic-interval 0
  monitor-period 1
!CONFIGURE THE NUMBER OF PREFIXES TO LEARN
  prefixes 500
  traffic-class keys dscp
 !CONFIGURE THE TYPE OF PREFIXES TO AGGREGATE
 aggregation-type prefix-length 32
 no max range receive
 backoff 90 90
!CONFIGURE THE PfR MODE
 mode route control
 mode select-exit best
 periodic 180
 resolve range priority 1
 resolve utilization priority 2 variance 10
 no resolve delay
!CONFIGURE PfR BORDER
 oer border
!ENABLE LOGGING
 logging
 local Loopback0
!POINT TO THE MASTER ROUTER AND OER KEY TO USE FOR AUTHENTICATION
 master 10.0.0.173 key-chain oer-key
```

9.7.1.1 PfR-Based Load Balancing

PfR can also track flows based on destination prefixes and link utilization. By intelligently combining destination prefixes into different groups, appropriate policy maps can be defined in PfR for different kinds of traffic. These different flows can then be appropriately routed and optimized for bandwidth. For example, consider the setup shown in Figure 9-11.

Figure 9-11. Dual-Homed SOHO Branch

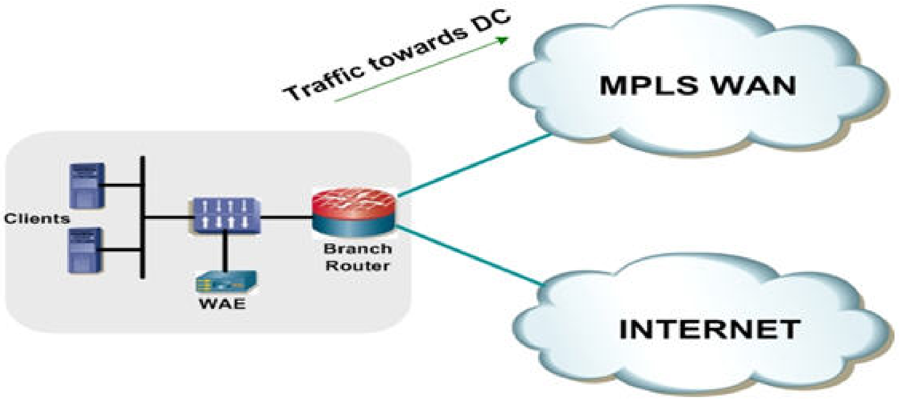

This setup has a dual homed branch router with one link over an MPLS WAN and the other over the Internet. Although the available bandwidth for each path is the same, delay and jitter are not. With normal routing protocols, the path through the MPLS WAN is the preferred exit and all traffic flows through that interface. The entire bandwidth available through the Internet exit is unused and wasted.

Using PfR to route all critical and delay sensitive traffic over the primary exit and all entertainment and non-critical traffic over Internet is probably a better solution. Thus, PfR can be used for efficient load balancing to ensure better bandwidth utilization. Section 9.7, “How PfR Works,” provides sample configurations.

9.7.1.2 PfR Link/Path Congestion/Failure

Normal routing protocols are primarily limited because they are based primarily on reachability. Although some routing protocols include cost as a metric for route calculations, cost is typically based on the theoretical link bandwidth and is a static value. This seriously impedes their routing path decisions because they do not track dynamically changing metrics such as delay or jitter.

For example, if there is congestion along a path and data packets are dropped or delay increases, normal routing protocols cannot effect a change in the routing path. A failure of this kind can affect all critical and real time user traffic. Because PfR can track metrics such as delay, jitter, and true reachability, it provides a potential solution for such situations. For example, consider the setup shown in Figure 9-12.

This setup has a dual homed branch router with two exit links. The path over Link1 provides a 30ms delay and the path over Link2 provides a delay of 60ms. The router is configured with PfR to track delay, send real time traffic over the path with the least delay, and send all other traffic over the alternate path.

Figure 9-12. Dual-Homed SOHO Branch with Multiple Exit Links

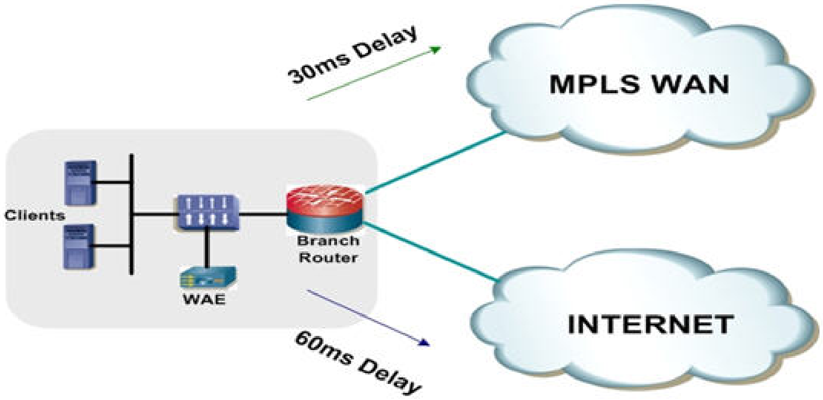

9.7.1.2.1 Case 1: Normal Conditions: No Congestion

PfR is active and is tracking delay and reachability. Because path 1 provides the lowest delay, all delay sensitive traffic is routed over path 1. As can be seen using show commands, PfR introduces a static route through path1 for delay sensitive traffic. All other traffic is routed over path 2 because of the default route. This is shown in Figure 9-13.

Figure 9-13. SOHO Branch with No Congestion

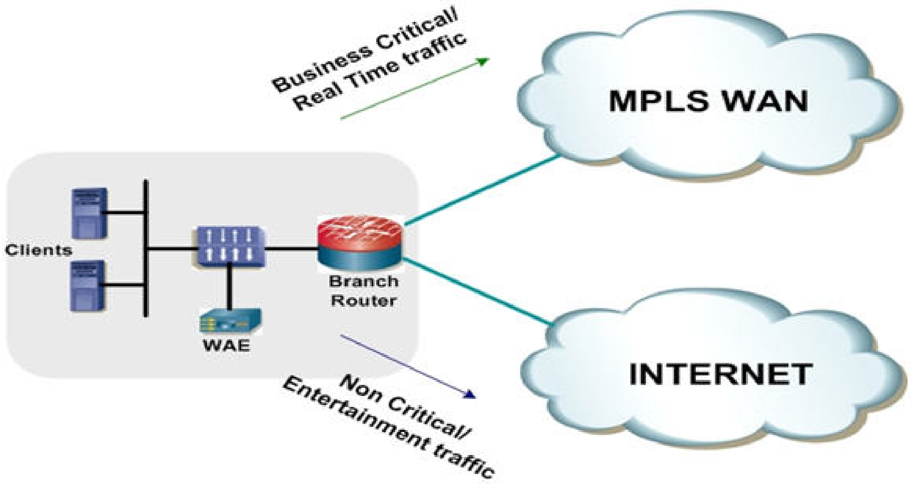

9.7.1.2.2 Case 2: Sudden Congestion in Downstream Path 1

A sudden congestion occurs downstream on path 1. Because of this, the delay over path 1 increases to more than 100ms. PfR senses this delay and effects a route change, diverting all real time traffic through path 2. Other traffic is not affected and continues to flow through path 2. See Figure 9-14 and Figure 9-15. Section 9.7, "How PfR Works," provides sample configurations.

Figure 9-14. SOHO Branch with Congestion

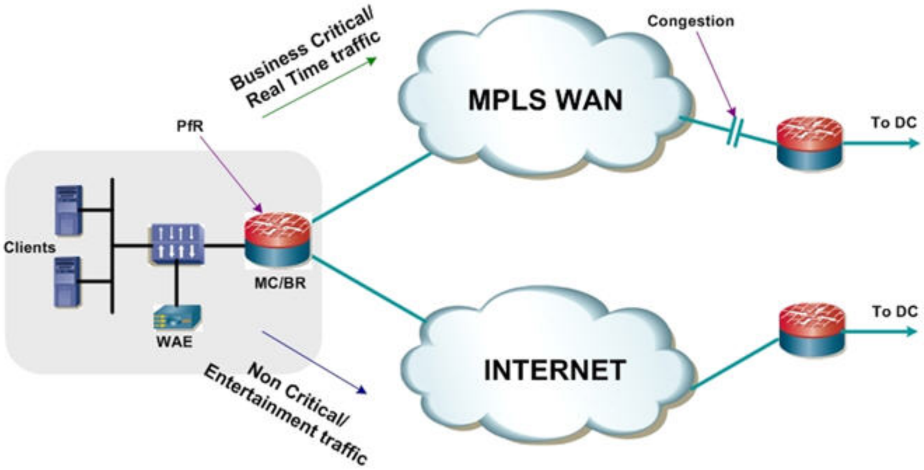

Figure 9-15. SOHO Branch Path Congestion with PfR Path Optimization

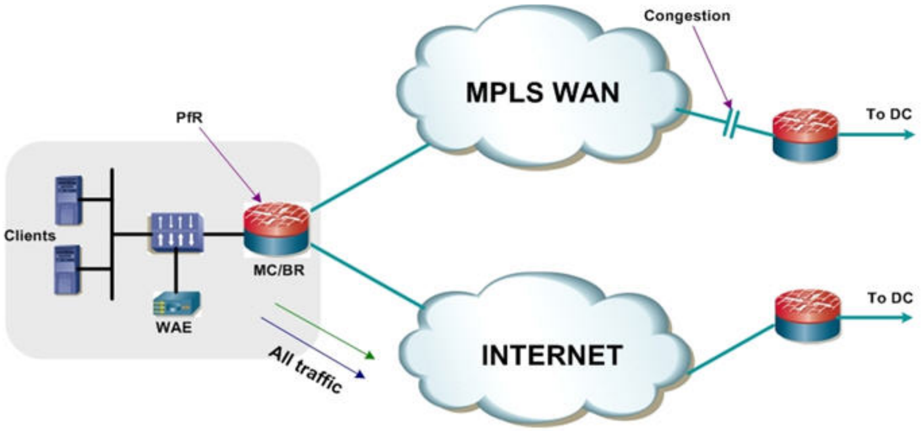

9.7.1.2.3 Case 3: Path Failure in Downstream Path 1

A link on a router within the WAN cloud fails, but the routing protocol does not relay this failure, possibly because of static routes or summarization, or possibly because data traffic is failing but the routing protocol remains up. Because of the failure, user traffic does not reach the other end of enterprise network. PfR observes increasing packet retransmissions and failures of its own probes. In response to the failure, PfR can direct all traffic to the alternate WAN path. As in the preceding case, traffic continues to flow unhindered over path 2. This is shown in Figure 9-16.

Figure 9-16. SOHO Branch Path Failure with PfR Path Optimization

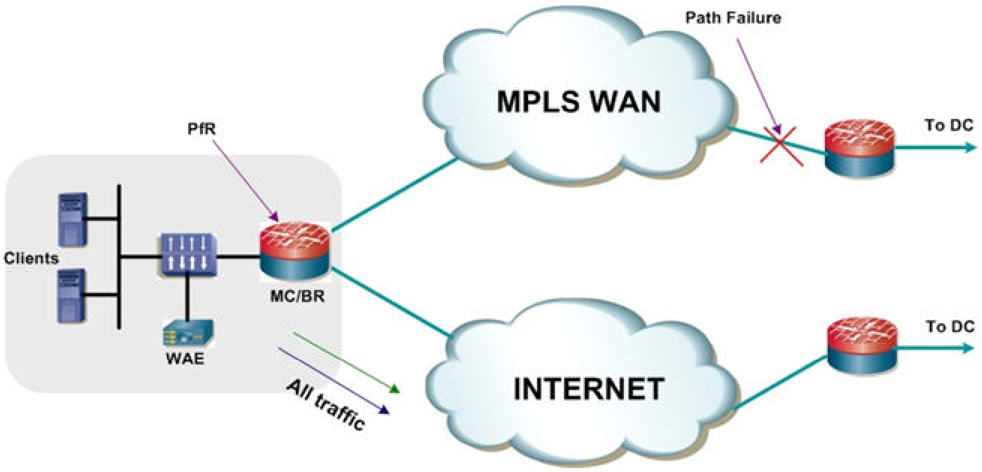

9.7.2 WCCP WAEs

WCCP was introduced in content routing to redirect web traffic to content cache engines. The idea was that with intelligent cache engines, network web traffic could be reduced; conversely, access speed could increase to improve the user experience. WCCP v1, introduced in 1997, was superseded by WCCP v2.

Cisco WAEs are devices used with Cisco IOS routers running WCCP to provide TCP optimization. WCCP is a common deployment tool used to intercept TCP traffic. WCCP redirects selected TCP traffic to the WAE, which can perform TCP optimization and compression, and which has data redundancy elimination (DRE) cache to help optimize the TCP traffic.

Figure 9-17 shows a typical WCCP and WAE deployment in a branch network.

Figure 9-17. WCCP and WAE in a Branch Network

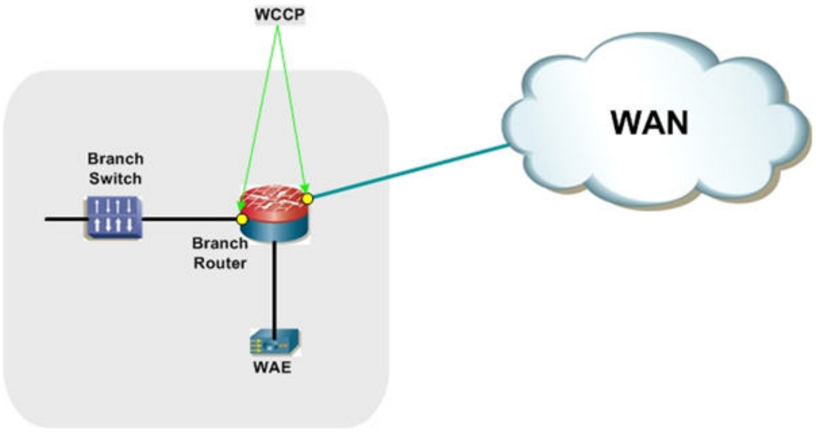

9.8 WANs

WANs are networks that cover a wide area. Typically, they interconnect LANs located in different areas. WANs imply a path of higher latency (generally 10s of milliseconds to several hundred milliseconds), relative lower bandwidth (therefore leading to more congestion), and a higher possibility of intermittent packet loss and path failures.

9.8.1 MPLS WANs

Enabling MPLS over the WANs provides multiple advantages to an Enterprise network:

- MPLS provides convergence of both L2 and L3 VPNs over the WAN.
- MPLS WANs retain extensive QoS capabilities and traffic engineering tunnels to provide real-time and mission-critical traffic with assured bandwidth and resiliency.
- Enterprises need not maintain any more specific WAN infrastructure.
- Enterprise network designs involving hub-and-spoke topologies are simpler. Hubs do not need to maintain individual adjacencies with the spokes, and interspoke traffic does not need to go through the hub. Routing becomes much simpler because spoke routers need not maintain a full routing table.
- All packet switching happens based on the labels the packets carry with less cumbersome route lookups.

Figure 9-18 shows a common deployment of branch to HQ connectivity over a MPLS WAN.

Figure 9-18. MPLS WAN

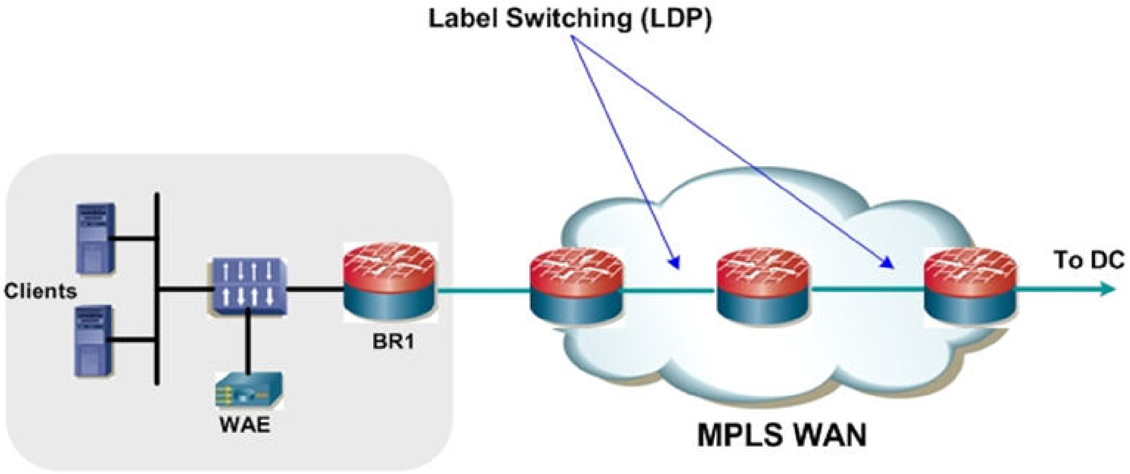

9.8.2 Internet-Based VPNs Secured using DMVPN

Even as MPLS-enabled WANs gain increased acceptance and deployment, a need remains for legacy WAN connectivity (direct leased line links, and so on), especially for SOHO and small branch offices. Additionally, Internet connectivity is useful in case the primary MPLS WAN link fails.

One major drawback for such a deployment is the lack of security for Internet traffic. Security for the traffic over Internet can be assured by securing the Internet connection with Dynamic Multicast VPNs (DMVPNs), which provide a variant of IP security (IPSec) protection with deployment friendly L3 features.

Figure 9-19 shows such a deployment.

Figure 9-19. Secure WAN over Internet

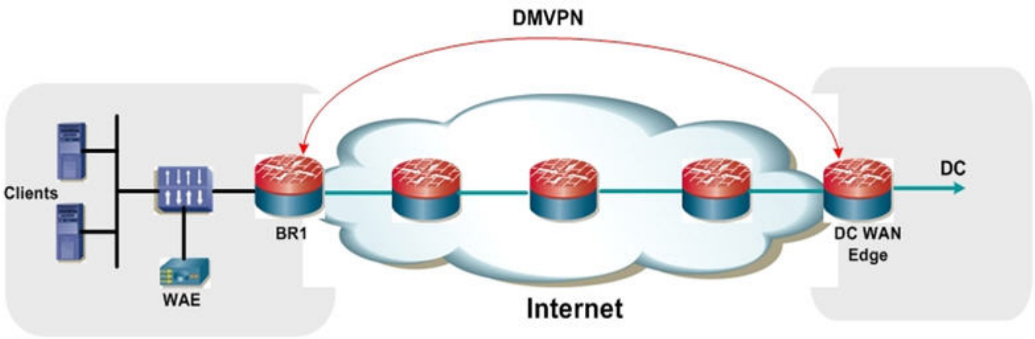

9.9 Security

The WAN and application optimization solution supports a variety of security options:

- IOS Firewall (IOS FW)
  - Content-Based Access Control (CBAC)
  - Zone-based firewall
  - Intrusion Detection System (IDS)

- DMVPN

9.9.1 IOS Firewall

IOS Firewall (IOS FW) provides users with a feature-rich and cost-effective way to implement security. IOS FW is particularly suited for branch networks where dedicated firewall boxes are not justified because of their complexity and cost. IOS FW provides users with protection along with seamless integration of other IOS features.

The Cisco IOS Firewall feature set has the following components:

- Context-Based Access Control (CBAC)
- Cisco Zone-Based Firewall
- Intrusion Detection

9.9.1.1 Content-Based Access Control (CBAC)

CBAC inspects all TCP and UDP packets and can dynamically create temporary openings for outbound traffic at the firewall interface. Return traffic is supported only for established sessions. CBAC maintains state information for every session. CBAC is interface based and is configured using ACLs.

9.9.1.2 Zone-Based Firewall

Zone-based firewall, as the name implies, is based on zones. Interfaces are usually associated with zones, and a zone can have multiple interfaces. Typically, on a Cisco router, LAN interfaces on the corporate network are configured to be in the trusted or inside zone; WAN interfaces are placed in the untrusted or outside zone. Zone-based firewalls are simpler to deploy. A class-based policy language is used to configure policies. Figure 9-20 shows a typical zone-based Firewall deployment.

Figure 9-20. Zone-Based Firewall

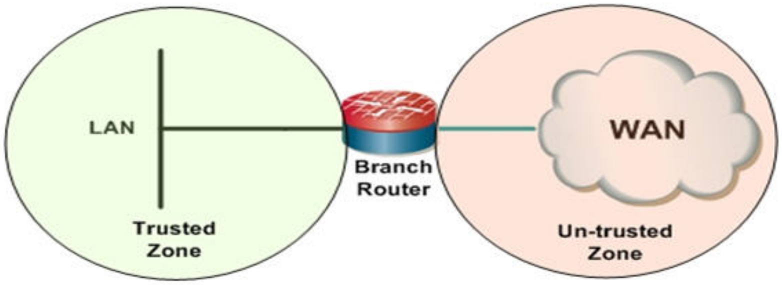

9.9.1.3 Intrusion Detection

The Firewall's Intrusion Detection system (IDS) contains most common attack signatures to detect intrusion. When IDS detects suspicious activity, IDS logs the event and can either shut down the port or send an alarm before network security is compromised.

For a more detailed explanation of these features and their use, see http://www.cisco.com/en/US/products/ps6586/products_ios_technology_home.html

9.9.2 DMVPN

DMVPN is an IPSec-based secure VPN technology that provides a high degree of security for traffic transmitted over public networks such as the Internet. DMVPN combines protocols such as Generic Routing Encapsulation (GRE), Next Hop Resolution Protocol (NHRP), and IPSec. DMVPN provides a dynamic full mesh based on spoke-to-spoke traffic.

Figure 9-21. DMVPN Hub-and-Spoke Deployment

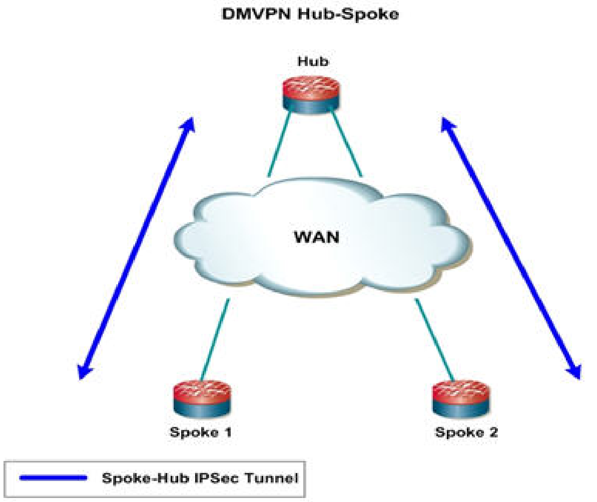

In a typical enterprise network with multiple branches, DMVPN is deployed in hub-and-spoke topology. Figure 9-21 shows such a deployment. The advantage with DMVPN is that the hub does not need to be configured for specific adjacency with each spoke. This helps scaling because hub reconfiguration is not necessary when new spokes are added. Another advantage with DMVPN is that spoke-to-spoke traffic does need not go through the hub. After spoke-to-spoke traffic is initiated, the hub sends a redirect message to the spokes. The spokes then set up a secure dynamic tunnel directly between them.

Figure 9-22 shows a typical setup.

Figure 9-22. DMVPN Spoke-to-Spoke Dynamic Tunnel

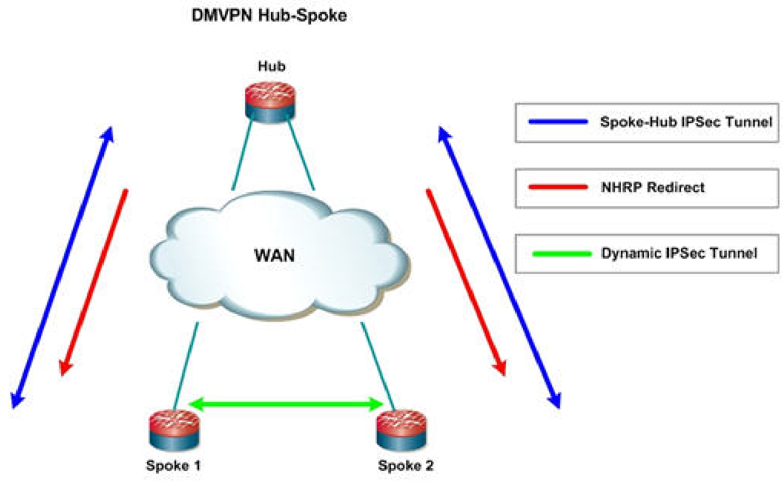

For more detailed information about these features and their use, see http://www.cisco.com/en/US/docs/ios/12_2t/12_2t13/feature/guide/ftgreips.html

9.10 Interoperability Considerations

The WAN bandwidth available to most branches is typically less than that of the LANs. The challenge for the user is to make optimal use of the available WAN bandwidth and provide LAN-like performance and experience over WAN. The technologies described in the preceding sections must be combined to realize maximum performance optimization and gains over the WAN.

9.10.1 Putting QoS and NBAR Together

In a typical branch, a myriad of applications, each with its own needs and idiosyncrasies such as response times, jitter, bandwidth, and so on, is running. In any well-designed branch, QoS plays an important role in allocating bandwidth between competing applications. However, traditional QoS cannot look very deeply into the packets and identify all applications. NBAR and its DPI capability can be used with QoS in a branch router to provide optimization. NBAR can be configured at the branch router ingress to mark flows from different applications with different ToS/DSCP. QoS can then use the marked flows on the egress to provide bandwidth optimization.

9.10.2 QoS, NBAR, NetFlow, and Path Optimization with PfR

NetFlow provides detailed tracking and statistics for each flow through the router. PfR uses NetFlow to track and monitor the prefixes that need to be controlled. Additionally, PfR can use DSCP values or L4 ports, and so on, to recognize flows based on different traffic classes.

A user can combine the DPI capability of NBAR and DSCP marking with the ability of PfR to track application classes based on those markings to provide resilience and route path optimization. Adding QoS to these provides bandwidth optimization. Figure 9-23 depicts a typical deployment scenario, and relevant sample configurations follow.

Figure 9-23. NBAR/NetFlow/PfR/QoS Interoperability

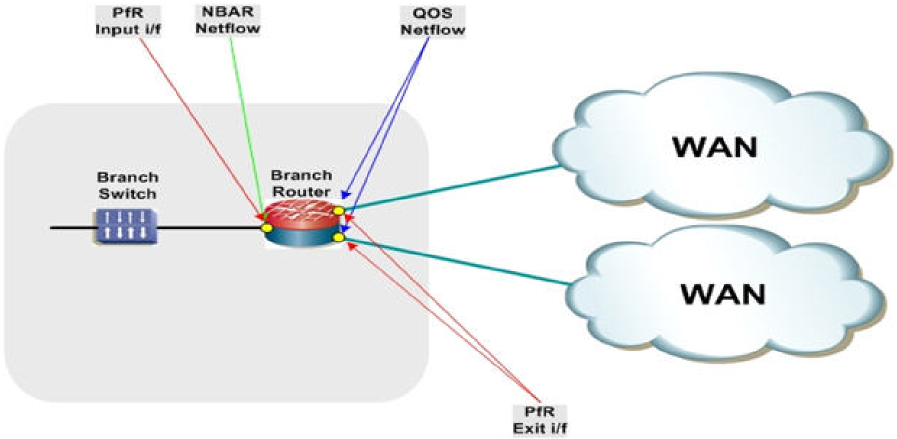

9.10.2.1 NBAR Configlets

```
!CONFIGURE APPROPRIATE NBAR CLASS MAPS TO MATCH DIFFERENT APPLICATION/PROTOCOLS
class-map match-all NBAR_UDP
 match access-group name UDP
class-map match-any NBAR_FTP
 match protocol ftp
class-map match-any NBAR_voice
 match protocol sip
 match protocol skinny
class-map match-any NBAR_HTTP
 match protocol http
 match protocol secure-http

!CONFIGURE APPROPRIATE POLICY MAPS TO SET THE IP DSCP VALUES FOR EACH
APPLICATION/PROTOCOL
policy-map NBAR
 class NBAR_HTTP
  set dscp cs2
 class NBAR_UDP
  set dscp cs5
 class NBAR_voice
  set dscp cs6
 class NBAR_FTP
  set dscp cs4
 class class-default
  set ip precedence 3
```

9.10.2.2 QoS Configlets

```
!CONFIGURE APPROPRIATE QoS CLASS MAPS TO MAP TO THE IP DSCP VALUES IN THE PACKETS
class-map match-all QoS_HTTP
 match ip precedence 2
class-map match-all QoS_FTP
 match ip precedence 4
class-map match-all QoS_UDP
 match ip precedence 5
class-map match-all QoS_voice
 match ip precedence 6

!CONFIGURE APPROPRIATE POLICY MAPS TO ALLOT BANDWIDTH TO DIFFERENT TRAFFIC IN
TIMES OF CONGESTION
policy-map QoS
 class QoS_HTTP
  bandwidth percent 25
 class QoS_UDP
  bandwidth percent 15
 class QoS_voice
  priority percent 10
 class QoS_FTP
  bandwidth percent 10
 class class-default
  bandwidth percent 15
```

9.10.2.3 PfR Configlets

```
!CONFIGURE THE OER KEY TO BE USED FOR MASTER-BORDER AUTHENTICATION
key chain oer-key
 key 1
   key-string WANOPT
!CONFIGURE PfR MASTER
oer master
!ASSIGN THE POLICY MAP TO BE USED FOR PATH-OPTIMIZATION (OPTIONAL)
 policy-rules delayPolicy
 logging
 !
!POINT TO THE BORDER ROUTER AND OER KEY TO USE FOR AUTHENTICATION
 border 10.0.0.172 key-chain oer-key
!DEFINE THE INTERNAL AND EXTERNAL INTERFACES ON THE BORDER ROUTER
  interface Serial6/1:0 external
  interface FastEthernet3/1 internal
  interface Serial6/0:0.1 external
 !
!CONFIGURE THE PfR MASTER TO LEARN TRAFFIC CLASS OR PREFIXES
 learn
!CONFIGURE THE PfR TO LEARN PREFIXES BASED ON DELAY
 delay
!CONFIGURE THE PfR TIMERS
  periodic-interval 0
  monitor-period 1
!CONFIGURE THE NUMBER OF PREFIXES TO LEARN
  prefixes 1000
!CONFIGURE THE TYPE OF PREFIXES TO AGGREGATE
  aggregation-type prefix-length 32
 no max range receive
!CONFIGURE THE PfR MODE
 mode monitor fast
 no resolve delay
 no resolve utilization
 !
!CONFIGURE AN ACTIVE PROBE TO ONE OF THE PREFIXES
active-probe echo 60.1.1.100
!
!CONFIGURE PfR BORDER
oer border
!ENABLE LOGGING
 logging
 local Loopback0
!POINT TO THE MASTER ROUTER AND OER KEY TO USE FOR AUTHENTICATION
 master 10.0.0.172 key-chain oer-key
```

Note: This configuration does not work as intended in IOS release 12.4(15)T2. See 9.11, “Caveats,” for more details. However, for TCP traffic that is intercepted by WCCP, the preceding configuration and policy does work, as packets would then exit and reenter the router over the interface connected to WAE. NetFlow can then recognize those markings and PfR can act on them. Applications can still be recognized based on L4 information. As a workaround, an upstream router can mark packets with the required DSCP rather than the WAN edge branch router.

9.10.3 WAAS Interoperability

WCCP can be deployed with all of the preceding IOS features to further enhance performance and optimization. Figure 9-24 depicts such a deployment. WAE provides TCP optimization only so other traffic, such as voice (over UDP) flows, is not affected by WAE. The following sections detail interoperability considerations when using WCCP with other features.

Figure 9-24. WCCP/NBAR/NetFlow/PfR/QoS Interoperability

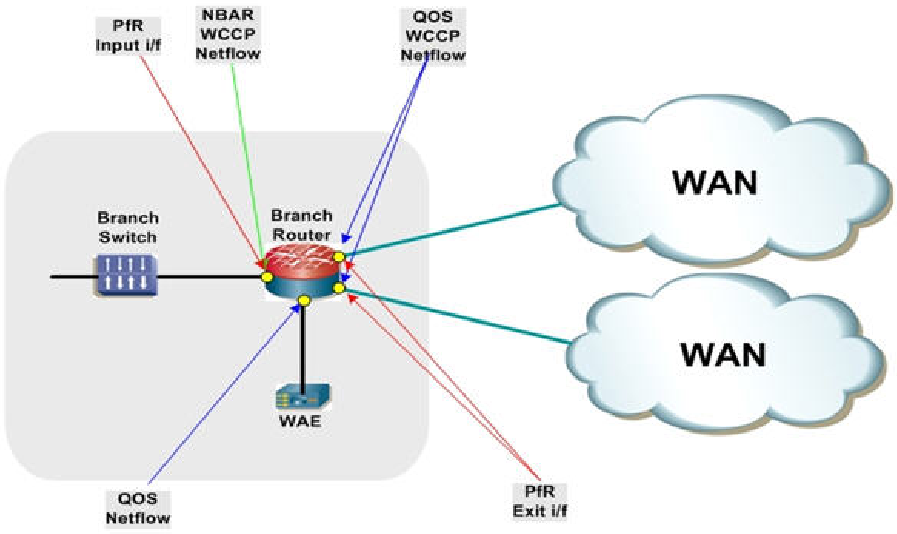

9.10.3.1 WAAS and Firewalls

Introduction of WAAS sets up three different TCP segments, as shown in Figure 9-25, along the data path. They include sessions between:

1. The client on the branch and the branch WAE
2. The branch WAE and the HQ/DC WAE, and
3. The HQ/DC WAE and the server

Figure 9-25. TCP Optimization with WAAS

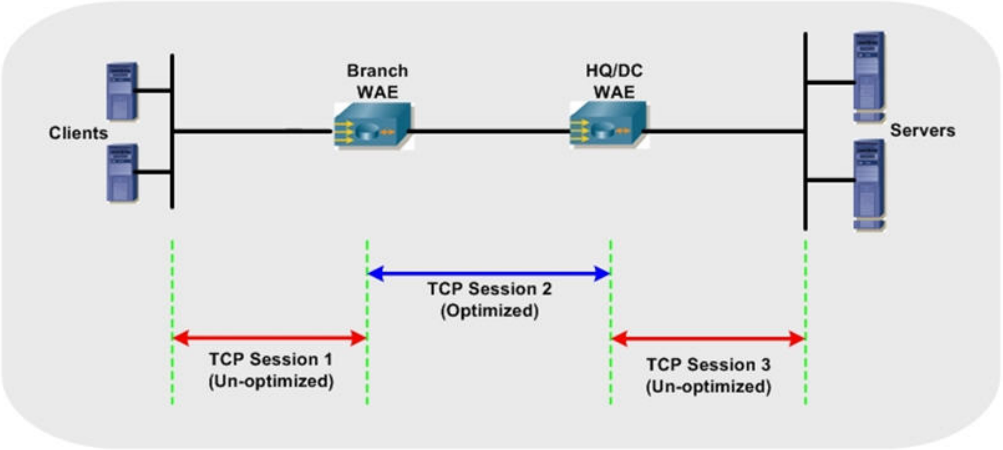
Branch WAE
HQ/DC WAE
Clients
Servers
TCP Session 2 (Optimized)
TCP Session 1 (Un-optimized)
TCP Session 3 (Un-optimized)

In the second session, the segment between the two WAEs initially uses a TCP sequence to match the original client, but on confirmation of a remote WAE willing to perform optimization, the TCP sequence number jumps to greater than 2 million for the leg between the WAEs. This jump provides a different TCP window for the optimized leg to help differentiate it from the unoptimized leg that the client and server actually see.

TCP options are used in the WAE-WAE TCP segment. This behavior causes the firewall, if it is deployed anywhere along the path in the middle segment, to treat the traffic as suspicious and drop it. Table 9-2 lists the various operating systems and the versions in which firewall drop behavior was fixed.

Table 9-2. Firewall Fixes

Operating System	Version	Command
IOS	12.4(11)T2	inspect waas enable
FWSM	3.2.1	inspect waas
PIX	7.2(3)	inspect waas

9.10.3.2 WCCP and NetFlow

NetFlow, which helps track different flows as they traverse routers and networks, has become indispensable for network operators. PfR depends on NetFlow to track and work with flows. However, when enabling NetFlow along with WCCP on the branch router, note the following:

- With both WCCP and NetFlow ingress enabled on the ingress interface of the branch router, a redirected flow cache entry has the output interface as “unknown,” as shown in Figure 9-26.
- When `61 WCCP in` and `62 WCCP out` are both configured on the branch router LAN interface, the router looks up some packets twice. This causes duplicate counts in the NetFlow statistics.

Figure 9-26. NetFlow and WCCP (NetFlow, WCCP, IP return (12.4T))

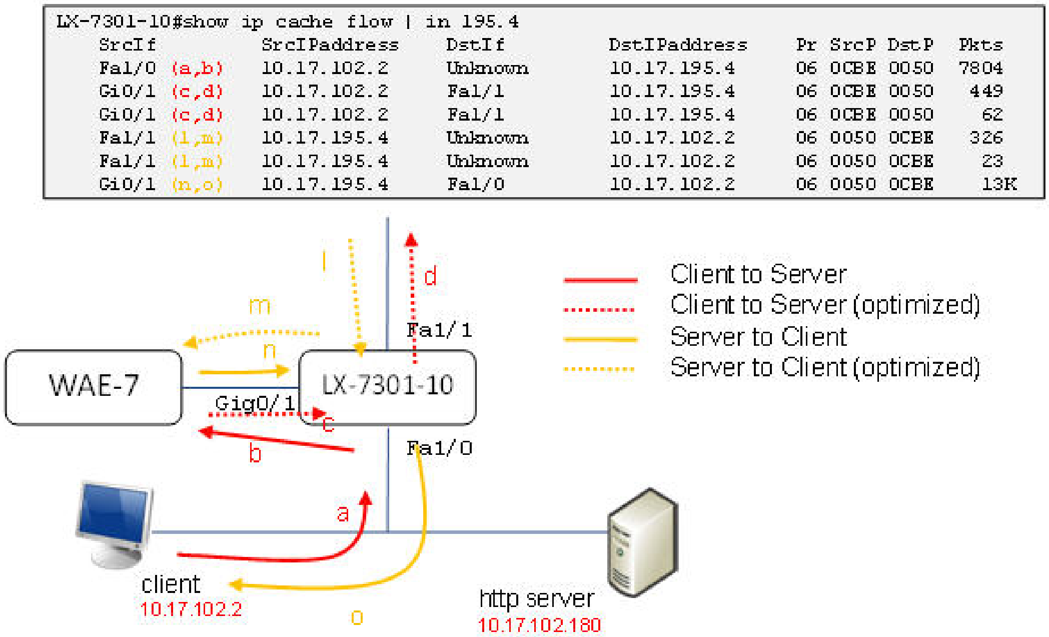

```
LX-7301-10#show ip cache flow | in 195.4
    SrcIf           SrcIPaddress    DstIf          DstIPaddress     Pr SrcP DstP  Pkts
    Fa1/0  (a,b)    10.17.102.2     Unknown        10.17.195.4      06 0CBE 0050  7804
    Gi0/1  (c,d)    10.17.102.2     Fa1/1          10.17.195.4      06 0CBE 0050   449
    Gi0/1  (c,d)    10.17.102.2     Fa1/1          10.17.195.4      06 0CBE 0050    62
    Fa1/1  (l,m)    10.17.195.4     Unknown        10.17.102.2      06 0050 0CBE   326
    Fa1/1  (l,m)    10.17.195.4     Unknown        10.17.102.2      06 0050 0CBE    23
    Gi0/1  (n,o)    10.17.195.4     Fa1/0          10.17.102.2      06 0050 0CBE   13K
```

9.10.3.3 WCCP and PfR

When configuring PfR, interfaces must be explicitly labeled as ”internal” and ”external” so that outgoing WAN traffic can be identified. Because of the NetFlow-WCCP “Unknown” issue, the GRE return feature in WAAS must be enabled so that an explicit internal to external NetFlow cache entry is seen. Additionally, because there are cases where WAAS itself originates traffic (for example, CIFS tunnel traffic), the WAAS-facing interface should labeled as PfR “internal.”

9.10.3.4 High Availability and WCCP

9.10.3.4.1 Branch WAE High Availability and Load Sharing

Consider a typical branch network with one branch router and one WAE. The WAE can become a single point of failure, disrupting branch activities. A WAE failure can be a complete failure (for example, power outage, software error, and so on) or a case in which the WAE single box performance limit is reached.

Although these failure possibilities are remote, deploying multiple WAEs at the branch can help minimize the effects of such failures. Using WCCP automatic load distribution features, the WAEs can work as a cluster and can share the load, and potentially protect against a single point of failure by providing backup if one WAE fails.

To help protect against branch router failure and improve availability, a user can set up two branch routers in the network. This also provides load sharing by the routers. In such cases, it is customary to implement HSRP or GLBP on the branch LANs. These two branch routers can share either one WAE or multiple WAEs.

9.10.3.4.2 Branch LAN High Availability with One Shared WAE

A typical network is shown in Figure 9-27.

Figure 9-27. Branch LAN High Availability - One WAN

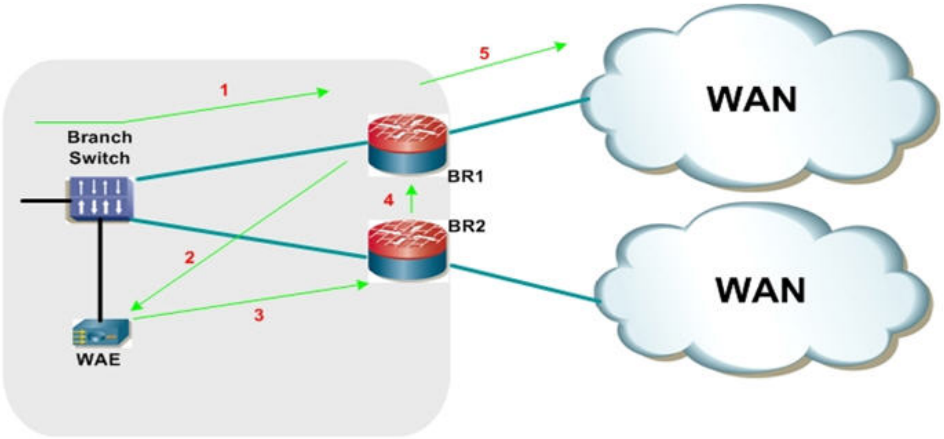

Configuring GLBP on the Branch LAN is more advantageous in such situations because GLBP provides load sharing. On the WAE side, the routers can be configured with HSRP and the WAE default gateway can be pointed toward the HSRP address. This, however, can result in inefficient routing.

Assume that branch routers BR1 and BR2 run a routing protocol between them and with the WAN routers. Let BR1 be the preferred exit to network X over the WAN. Let BR2 be the higher priority router for HSRP on the WAE VLAN. With this setup, let us look at a packet flow from branch to network X:

1. The client on the branch and the branch WAE.
2. WCCP intercepts the packet on BR1 and sends them to the WAE.
3. The WAE processes the packet and returns it to its default gateway, which is BR2.
4. BR2 does a lookup for the packet, realizes that BR1 is the preferred exit, and sends the packet to BR1.
5. BR1 forwards the packet through the WAN.

As can be easily seen, if the WAE could have returned the packet to BR1 instead of to BR2, the extra hop and the consequent delay could have been avoided. A workaround for such issues is to add static routes to pointing to the standby address of the preferred branch router on the WAE. The same problem can happen on the reverse path with asymmetric routing. The problem might be worse in that case, because it might involve unoptimized traffic flowing between the two branch routers. Therefore, there should be adequate bandwidth between the branch routers.

Starting with version 4.0.13, WAEs support negotiated GRE return for optimized packets. With GRE return, the WAE returns the packet directly to the router that sent it. Note that on the routers, HSRP on the WAE side and on the WAE pointing the default gateway to the HSRP address is still recommended. This provides higher availability for WAE connectivity to the network and Central Manager.

9.10.3.4.3 Branch LAN High Availability with Two Shared WAEs

Figure 9-28 shows this implementation.

Figure 9-28. Branch LAN High Availability with Two WAE

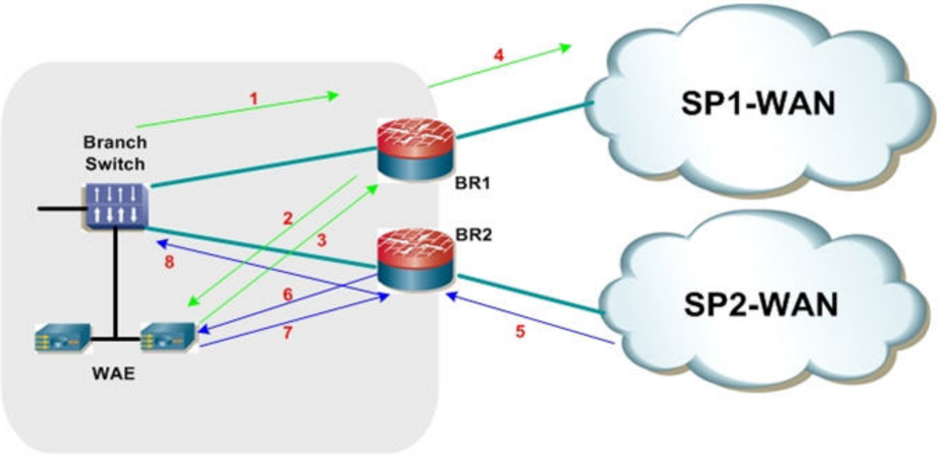

As in the previous case, on the routers, GLBP/HSRP is configured on the Branch LAN side and HSRP on the WAE side. The WAE default gateways are pointed to the HSRP address. Configuring GRE return is preferred on the WAEs.

Any multiple link/router environments might give rise to asymmetric routing issues, where packets exit the branch through one router (or one link), such as BR1, and return through another link on the same router (or, possibly, on a different router at the site). On the branch routers, WCCP handles intercepting and directing returning packets to the correct WAE. Figure 9-28 shows a packet flow in the asymmetric routing case.

1. Packets from Branch LAN to network X are forwarded to BR1.
2. WCCP intercepts these packets on BR1 and sends them to the WAE.
3. The WAE processes it and returns the packet to its BR1 (GRE return).
4. BR1 forwards the packet to DC/HQ over SP1-WAN.
5. DC/HQ forwards return packet using SP2-WAN to BR2.
6. BR2 WCCP intercepts this packet and sends to WAE.
7. WAE processes and returns the packet to BR2 (GRE return).
8. BR2 now forwards the packet to the client through the LAN.

9.11 Caveats

This section lists caveats for the various components of the WAN and application optimization solution.

9.11.1 PfR Supports Only One Next Hop per interface

PfR, by design has traditionally supported only one Next Hop per interface. This means PfR can be deployed only in topologies using logical point-to-point technologies, such as:

- Single-peer Ethernet and Packet over SONET/SDH (POS)
- Frame-Relay (P2P), Point-to-Point Protocol (PPP), High-Level Data Link Control (HLDC), GRE
- L3VPN Services

PfR cannot be deployed in topologies using:

- ISP peering exchanges (generally, common Ethernet VLAN, or SONET ring)
- VPLS
- DMVPN (mGRE)

However, PfR can work with DMVPN solution using point-to-point GRE tunnels. PfR support for multipoint interfaces is under development.

9.11.2 PfR Supports only BGP or Static Routes for Path Optimization

For path optimization in a router, PfR uses static routes or BGP routes as parent routes. PfR cannot control routes learned using other protocols. However, many WAN deployments use EIGRP or Open Shortest Path First (OSPF). In such cases, the workaround is to use static summary routes and use PfR for route unreachability mitigation.

Support for EIGRP in PfR is under development.

9.11.3 PfR Might Break WAAS TCP Optimization if the WAAS Network Path is Changed

PfR tracks various traffic parameters and, based on those parameters, shifts the network path of an application from one link to another to perform path optimization. While doing so, if WAAS is inside the PfR domain as shown in Figure 9-29, the WAAS TCP optimization will be broken. To avoid this, always place outer WAEs outside the PfR domain.

Figure 9-29. PfR-WAAS Network Path

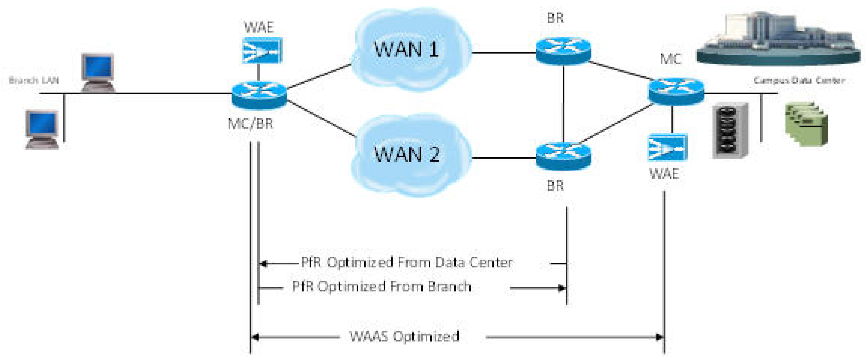

9.11.4 PfR Interface Mapping and WAAS

PfR tracks traffic that enters and exits the router. PfR tags interfaces through which traffic flows as external interfaces or internal interfaces. Typically, the interface through which the traffic enters into the router from the LAN network is tagged as internal and the interfaces through which the traffic exits the router (the WAN) are tagged as external. PfR terms traffic that flows from internal to external interfaces as interesting. PfR needs the traffic in the reverse direction, from external to internal, for verification. Figure 9-30 shows this tagging.

If a router is doing TCP optimization with WAAS, the interface connecting the router to the WAAS cannot be clearly distinguished as internal or external. This is because it acts as the exit interface for traffic that is intercepted by WCCP on the router and as an ingress interface for optimized or un-optimized traffic. The workaround is to configure the WAE to do GRE return.

Additionally, if there is CIFS traffic, the CIFS traffic is not returned using GRE. Therefore, for CIFS traffic, the WAE interface on the router should be tagged as PfR internal.

Another option is to disable the return traffic verification of PfR by using the command `no mode verify bidirectional`. This, however, exposes PfR to self-created black holes.

9.11.5 PfR Cannot Recognize MQC Marking Done by the Same Router

If the branch router running PfR is doing DSCP or ToS marking on ingress traffic, PfR cannot recognize these remarkings. PfR uses ingress NetFlow to learn and verify traffic. Ingress NetFlow records the original received ToS marking when the packet entered the router and does not record the remarked value. Therefore, PfR cannot match locally remarked flows based on DSCP or ToS. Figure 9-30 illustrates this issue.

Figure 9-30. PfR and Modular QoS CLI (MQC) Mappings

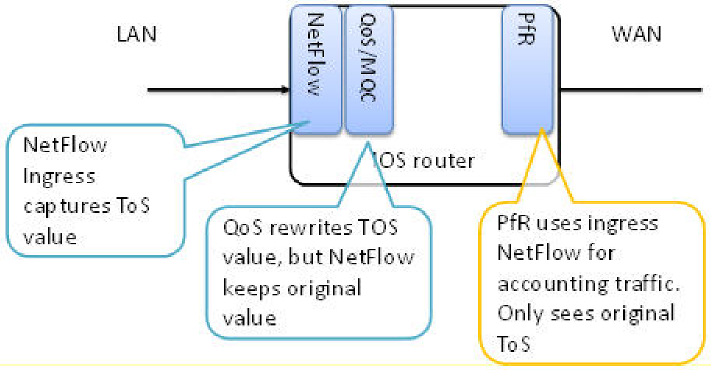
LAN
NetFlow
QoS/MQC
PfR
WAN
IOS router
NetFlow Ingress captures ToS value
QoS rewrites TOS value, but NetFlow keeps original value
PfR uses ingress NetFlow for accounting traffic. Only sees original ToS

This, however, does not affect TCP traffic redirected to WAAS because the markings on the packets on reentry are recorded by NetFlow, which PfR can then use. The workaround is to use upstream devices to mark the traffic. DDTS CSCsk99096 tracks this issue.

9.11.6 PFR Interface Mapping and NetFlow Sampling

When NetFlow sampling is enabled on an interface, that interface cannot be used as an internal or external PfR interface. A solution for this issue is under development.

9.11.7 CIFS tunneling on WAE and Network visibility

In Windows operating system environments, CIFS (also known as Windows file sharing) is predominantly used for file sharing. With CIFS application acceleration enabled on such networks with WAAS, WAEs tunnel CIFS traffic across the network over TCP port 4050 and the IP addresses of the WAEs involved.

This affects the visibility of technologies such as NetFlow and QoS, because they are no longer aware of the actual endpoints. Tunneled traffic will affect the ability of PfR to granularly path optimize the optimized CIFS traffic because PfR relies on NetFlow. This issue will be resolved in the upcoming release of WAE software in late 2008.

Figure 9-31 illustrates CIFS tunneling with WAE.

Figure 9-31. WAE CIFS Tunneling

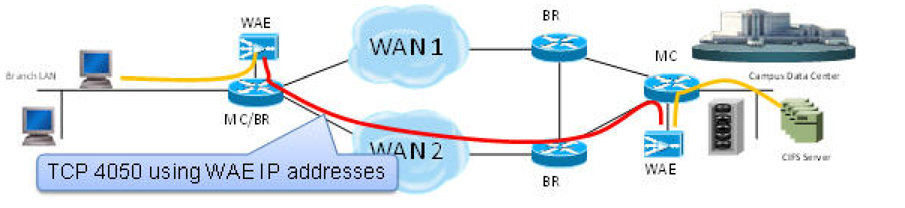

9.11.8 WAAS and Firewall

Introducing WAAS creates three different TCP segments: one between the client and the client side WAE, the second between the client side WAE and server side WAE, and the third between the server side WAE and the server.

The WAEs use TCP options (0x21) for autodiscovery and jump up the sequence number to more than 2 million in the WAE-WAE TCP session. Any firewall deployed in between the WAEs might view these activities as suspicious and might drop the packets. This behavior is fixed in the IOS versions detailed in 9.10.3.1, “WAAS and Firewalls.”

9.11.9 WCCP and NHRP Redirect

Next Hop Routing Protocol (NHRP) is part of DMVPN technology. In a typical hub-and-spoke DMVPN deployment, when spoke-to-spoke traffic hits the hub, it sends a NHRP redirect to the spokes (as shown in Figure 9-32) to enable the spokes to establish a direct secure tunnel between them. For NHRP on the hub to send the redirect, it should see the same packet enter and exit the same tunnel interface. This breaks when WCCP is enabled on the DMVPN hub.

Figure 9-32. DMVPN-NHRP Redirect

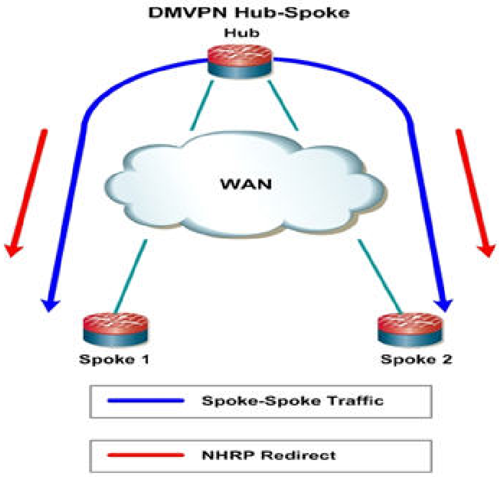

In DMVPN deployments with WCCP, WCCP intercept is configured on the tunnels. Therefore, any packet traveling from spoke-to-spoke, on reaching the tunnel, is intercepted by WCCP and sent to the WAE. Assuming that the WAE is doing L3 return, the return packet from the WAE is then sent out the hub. This breaks the NHRP condition to send the redirect. Therefore, no redirect is sent to the spokes and no direct tunnels are established between them. With WAE doing GRE return, NHRP incorrectly sends the redirect to the WAE instead of the spokes, and no spoke-to-spoke tunnels can be established.

Note that this affects only spoke-to-spoke traffic and optimization, and does not affect spoke-to-hub or hub-to-spoke traffic. The workaround is to remove the WCCP intercept on the tunnel interface on the hub and configure it on its LAN interface. The LAN interface on the hub will have WCCP intercept for both in and out directions. This has some performance implications on the hub, as the hub router must do two route lookups for spoke-to-hub traffic.

9.11.10 WAAS Might Not Intercept IP SLA Probes Configured on the Branch Router

In any network, IP SLAs are an essential component for monitoring network performance parameters such as round-trip-time (RTT). On a branch network, it is typical to configure IP SLAs on the branch router connected to the WAN to measure performance. This is shown in Figure 9-33.

Current IOS code does not perform WCCP intercept of any IP SLA probes configured on the branch router. However, WCCP intercepts all regular TCP traffic that passes through the branch router. This can create monitoring issues because parameters, such as network RTT, reported by such SLAs on the branch router do not represent the actual user experience. WCCP does not intercept SLA probes, so WAAS does not optimize them. The workaround is to configure such SLA probes on an upstream router, where WCCP will intercept them on the branch router.

Figure 9-33. IP SLA and WCCP

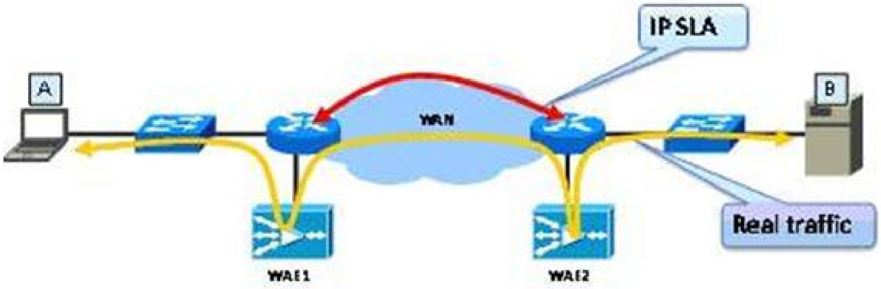

9.11.11 NBAR Cannot Perform DPI if WAE TCP Optimization Occurs before NBAR Discovery

One of the key abilities of NBAR is its ability to look deep inside packets and classify protocols and applications. In a network using NBAR and inline WAE, an inline WAE located before the branch router, as shown in Figure 9-34, causes optimization to occur first on the TCP flows. Similarly, with WCCP interception along with egress NBAR on branch routers, optimization occurs before NBAR protocol discovery.

Figure 9-34. WAAS Inline and NBAR

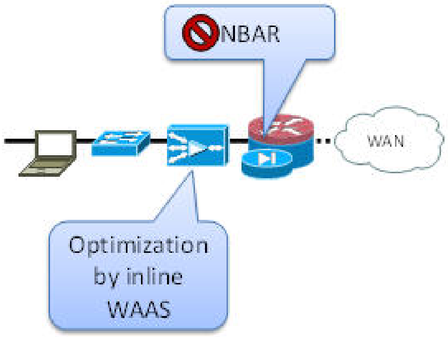

For flows that WAE optimizes using LZ and DRE, packets that exit the WAE can be obfuscated and NBAR can no longer rely on DPI to identify such flows. NBAR can still identify flows based on L4 ports, but as discussed previously, such a classification is not always dependable.

Therefore, WAE and NBAR should be placed so that NBAR can operate on the flows before WAE can. For example, in the WCCP interception case, NBAR should be configured on the ingress of the branch router, as shown in Figure 9-35.

Figure 9-35. WCCP and Egress NBAR

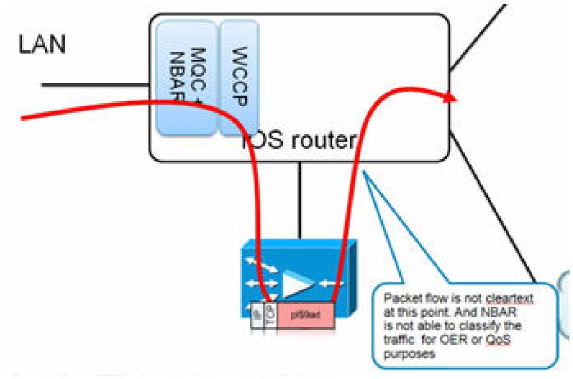
LAN
MQC + NBAR
WCCP
IOS router
pt50ad
Packet flow is not cleartext at this point. And NBAR is not able to classify the traffic for OER or QoS purposes

9.12 Example Deployment Models

This section describes a variety of branch office deployment models, varying in size and optimization technologies.

9.12.1 Small Branch Office with Single-Homed SOHO Branch Router

This model, shown in Figure 9-36, is commonly deployed in SOHO branches. It is simple to configure and maintain.

Figure 9-36. Small Branch Office with Single-Homed Branch Router

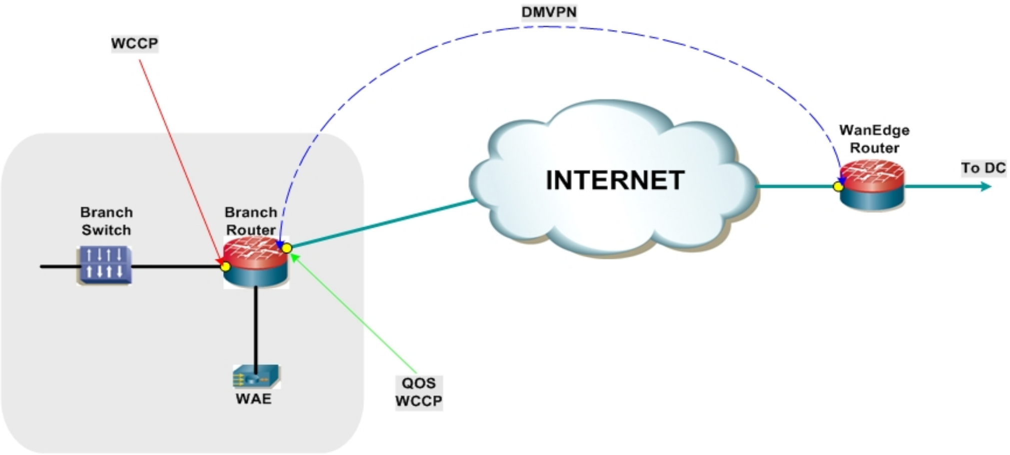

In this deployment, the branch router is single homed, and has one interface connected to Corporate HQ over Internet. The branch router is configured with the following:

1. Packets arriving on the LAN interface on the branch router are classified using DPI and the DSCP field is marked accordingly.

2. TCP packets are redirected to the WAE using WCCP.

3. The TCP flows are optimized by WAE and returned to the branch IOS router.

4. The WAE preserves the DSCP markings originally done by NBAR (using DPI) on the branch IOS router ingress interface.

5. Low latency traffic such as voice is not sent to the WAE, but directly to the WAN interface, which is connected through a DMVPN tunnel to the corporate HQ.

6. Return traffic from the WAE is mixed with the non-WAE optimized traffic and outbound QoS gives relative priority to the application that needs it, based on DSCP.

9.12.1.1 Sample Branch Router Configuration

A sample branch router configuration for this deployment model follows:

```
Current configuration : 8884 bytes
```

```
!
! Last configuration change at 19:37:45 EST Tue Jan 29 2008
! NVRAM config last updated at 19:27:30 EST Tue Jan 29 2008
!
version 12.4
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
service internal
!
hostname WANOPT-3845-BR1
!
boot-start-marker
boot system flash:c3845-adventerprisek9-mz.124-15.T2.fc3
boot-end-marker
!
logging buffered 16000000
enable password lab
!
no aaa new-model
clock timezone EST -5
clock summer-time EDT recurring
no network-clock-participate wic 1
!
!
crypto isakmp policy 10
 encr 3des
 authentication pre-share
 group 2
crypto isakmp key cisco address 0.0.0.0 0.0.0.0
crypto isakmp keepalive 30
!
!
crypto ipsec transform-set gre_set esp-3des esp-sha-hmac
 mode transport
!
crypto ipsec profile gre_prof
 set transform-set gre_set
!
!
ip wccp 61
ip wccp 62
ip cef
!
ip domain list wanopt4.cisco.com
ip domain name wanopt4.cisco.com
ip host www.cisco.com 60.1.1.100
ip name-server 52.1.1.100
ip inspect WAAS enable
!
multilink bundle-name authenticated
!
voice-card 0
 no dspfarm
!
username cisco privilege 15 secret 5 $1$D.pS$pRcbUNacHbYzs9.BnwDeP0
username lab privilege 15 password 0 lab
archive
 log config
  hidekeys
!
!
controller T1 0/1/0
 framing esf
 linecode b8zs
!
controller T1 0/1/1
 framing esf
 linecode b8zs
 channel-group 0 timeslots 1-24
vlan internal allocation policy ascending
!
ip ftp source-interface FastEthernet1/15
ip ftp username anonymous
ip ftp password nobody@cisco.com
!
class-map match-all NBAR_UDP
 match access-group name UDP
class-map match-any NBAR_FTP
 match protocol ftp
class-map match-all QoS_HTTP
 match ip precedence 2
class-map type inspect match-any wanopt
 match protocol tcp
 match protocol udp
 match protocol icmp
class-map match-any NBAR_voice
 match protocol sip
 match protocol skinny
class-map match-all QoS_FTP
 match ip precedence 4
class-map match-all QoS_UDP
 match ip precedence 5
class-map match-any NBAR_HTTP
 match protocol http
 match protocol secure-http
class-map match-all QoS_voice
 match ip precedence 6
!
!
policy-map NBAR
 class NBAR_HTTP
  set dscp cs2
 class NBAR_UDP
  set dscp cs5
 class NBAR_voice
  set dscp cs6
 class NBAR_FTP
  set dscp cs4
 class class-default
  set ip precedence 3
policy-map type inspect fromOutside
 class type inspect wanopt
  inspect
 class class-default
policy-map QoS
 class QoS_HTTP
  bandwidth percent 25
 class QoS_UDP
  bandwidth percent 15
 class QoS_voice
  priority percent 10
 class QoS_FTP
  bandwidth percent 10
 class class-default
  bandwidth percent 15
policy-map type inspect toOutside
 class type inspect wanopt
  inspect
 class class-default
!
!
interface Loopback0
 ip address 10.0.0.161 255.255.255.255
!
interface Tunnel1
```

```
 ip address 172.20.1.2 255.255.255.0
 no ip redirects
 ip mtu 1400
 ip wccp 62 redirect in
 ip flow ingress
 ip nhrp authentication nsite
 ip nhrp map 172.20.1.254 10.18.101.2
 ip nhrp map multicast 10.18.101.2
 ip nhrp network-id 101
 ip nhrp holdtime 300
 ip nhrp nhs 172.20.1.254
 ip nhrp cache non-authoritative
 ip nhrp shortcut
 zone-member security outside
 ip tcp adjust-mss 1360
 tunnel source Serial0/0/0
 tunnel mode gre multipoint
 tunnel protection ipsec profile gre_prof
!
interface GigabitEthernet0/0
 description To AGILENT via WANOPT-6500-MAIN1:G1/25
 ip address 10.19.1.1 255.255.255.0
 ip wccp 61 redirect in
ip flow ingress
load-interval 30
 duplex full
 speed 1000
 media-type rj45
 service-policy input NBAR
!
interface GigabitEthernet0/1
 description To CAT35k-BrMAN-1:F0/1
 ip address 100.1.1.161 255.255.255.0
 duplex full
 speed 100
 media-type rj45
 ntp broadcast client
!
interface Serial0/0/0
 description To WANOPT-7206-INTERNET:S5/0:0
 ip address 10.18.1.2 255.255.255.0
 ip flow ingress
 load-interval 30
 service-policy output QoS
!
interface Serial0/1/1:0
 no ip address
 shutdown
!
interface FastEthernet1/0
 duplex full
 speed 100
!
interface FastEthernet1/1
 duplex full
 speed 100
!
interface FastEthernet1/2
 duplex full
 speed 100
!
interface FastEthernet1/3
 duplex full
 speed 100
!
interface FastEthernet1/4
 duplex full
 speed 100
```

```
!
interface FastEthernet1/5
 duplex full
 speed 100
!
interface FastEthernet1/6
 duplex full
 speed 100
!
interface FastEthernet1/7
 duplex full
 speed 100
!
interface FastEthernet1/8
 duplex full
 speed 100
!
interface FastEthernet1/9
 duplex full
 speed 100
!
interface FastEthernet1/10
 duplex full
 speed 100
!
interface FastEthernet1/11
 duplex full
 speed 100
!
interface FastEthernet1/12
 description To CAT6K-MAIN:G1/20
 no switchport
 no ip address
 shutdown
 duplex full
 speed 100
!
interface FastEthernet1/13
 description To CAT6K-MAIN:G1/18
 no switchport
 no ip address
 shutdown
 duplex full
 speed 100
!
interface FastEthernet1/14
 description To CAT6K-MAIN:G1/16
 no switchport
 no ip address
 duplex full
 speed 100
!
interface FastEthernet1/15
 no switchport
 no ip address
 ip tcp adjust-mss 1260
 load-interval 30
 duplex full
 speed 100
!
interface GigabitEthernet1/0
!
interface Integrated-Service-Engine4/0
 ip address 10.18.51.1 255.255.255.0
 ip flow ingress
 service-module ip address 10.18.51.2 255.255.255.0
 service-module ip default-gateway 10.18.51.1
 no keepalive
```

```
!
interface Vlan1
 no ip address
!
router eigrp 10
 passive-interface default
 no passive-interface Tunnel1
 network 10.0.0.161 0.0.0.0
 network 10.19.1.0 0.0.0.255
 network 172.20.1.0 0.0.0.255
 no auto-summary
!
router bgp 171
 no synchronization
 bgp log-neighbor-changes
 network 10.0.0.161 mask 255.255.255.255
 network 10.17.1.0 mask 255.255.255.0
 network 10.18.51.0 mask 255.255.255.0
 network 10.19.1.0 mask 255.255.255.0
 neighbor 10.17.1.1 remote-as 103
 no auto-summary
!
ip route 0.0.0.0 0.0.0.0 10.18.1.1
ip route 100.1.1.22 255.255.255.255 10.18.1.1
!
ip flow-cache timeout active 1
ip flow-export source GigabitEthernet0/0
ip flow-export version 9
ip flow-export destination 52.1.1.22 9995
ip flow-top-talkers
 top 10
 sort-by bytes
!
ip http server
ip http authentication local
ip http secure-server
ip http timeout-policy idle 60 life 86400 requests 10000
!
ip sla 1
 http get http://www.cisco.com
 timeout 5000
 owner HTTP - 100.1.1.161
 tag WANOPT HTTP ECHO
ip sla schedule 1 life forever start-time now ageout 3600
ip sla 3
 icmp-echo 60.1.1.100
 owner ICMP Echo - 100.1.1.161 - 60.1.1.100
 tag WANOPT ICMP ECHO
ip sla schedule 3 life forever start-time now ageout 3600
ip sla 4
 dns www.cisco.com name-server 52.1.1.100
 timeout 5000
 owner DNS - 100.1.1.161
 tag WANOPT DNS SLA
ip sla schedule 4 life forever start-time now ageout 3600
logging 100.1.1.104
snmp-server community closed RW
snmp-server community open RO
snmp-server community public RO
snmp-server community private RW
snmp-server ifindex persist
snmp-server trap-source GigabitEthernet0/1
snmp-server enable traps cnpd
snmp-server host 100.1.1.21 version 2c public
!
tftp-server flash:c3845-adventerprisek9-mz.124-11.T2
!
control-plane
```

```
!
!
banner login ^C
-----------------------------------------------------------------------
Cisco Router and Security Device Manager (SDM) is installed on this device.
This feature requires the one-time use of the username "cisco"
with the password "cisco". The default username and password have a privilege
level

Please change these publicly known initial credentials using SDM or the IOS CLI.
Here are the Cisco IOS commands.

username <myuser>  privilege 15 secret 0 <mypassword>
no username cisco

Replace <myuser> and <mypassword> with the username and password you want to use.

For more information about SDM please follow the instructions in the QUICK START
GUIDE for your router or go to http://www.cisco.com/go/sdm
-----------------------------------------------------------------------
^C
!
line con 0
 login local
line aux 0
line 258
 no activation-character
 no exec
 transport preferred none
 transport input all
 transport output pad telnet rlogin lapb-ta mop udptn v120 ssh
line vty 0 4
 exec-timeout 0 0
 privilege level 15
 password lab
 login
 transport input telnet ssh
line vty 5 15
 exec-timeout 0 0
 privilege level 15
 password lab
 login
 transport input telnet ssh
line vty 16 20
 password lab
 login
line vty 21 30
 password cisco
 login
!
scheduler allocate 20000 1000
ntp clock-period 17179694

!
webvpn cef
!
end
```

9.12.2 Small Branch Office with Dual-Homed, Single-Tier Branch Router

This deployment model (shown in Figure 9-37) is similar to the previous one. In this model, the branch router is dual-homed. One branch router exit interfaces is connected to the SP WAN, and the other is connected to the Internet.

Figure 9-37. Small Branch Office with Dual-Homed Router

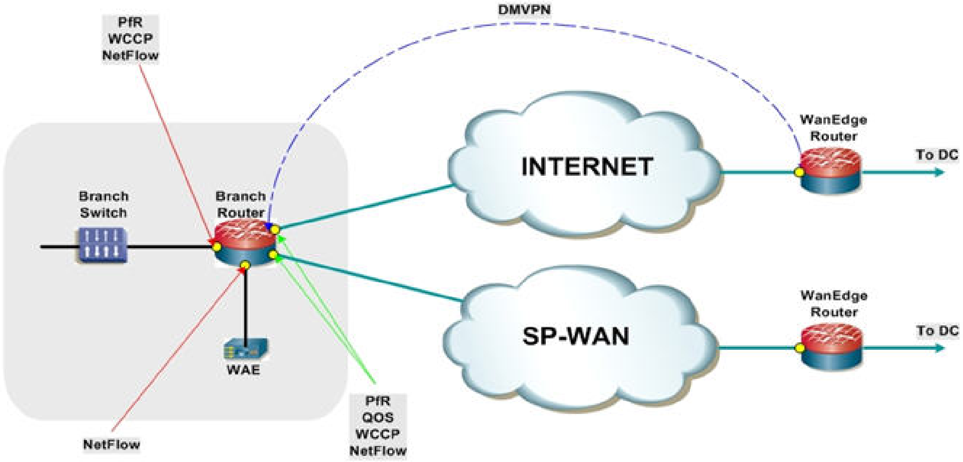

The following protocols are deployed on the branch router:

1. QoS on the exit interface to do congestion avoidance and management
2. WCCP on the router to do TCP optimization
3. NetFlow on all LAN/WAN links
4. PfR on the router to do path optimization
5. DMVPN tunnel over Internet to the Corporate HQ to provide security

9.12.2.1 Sample Branch Router Configuration

```
Current configuration : 8234 bytes
!
! Last configuration change at 10:49:16 EST Wed Jan 30 2008
! NVRAM config last updated at 10:49:38 EST Wed Jan 30 2008
!
upgrade fpd auto
version 12.4
no service pad
service timestamps debug datetime msec localtime
service timestamps log datetime localtime
no service password-encryption
!
hostname WANOPT-7206-BR6
!
boot-start-marker
boot system flash disk0:c7200-adventerprisek9-mz.124-15.T2.fc3
boot-end-marker
!
logging buffered 16000000
no logging console
enable secret 5 $1$qmlj$FYkMGt1ksVs.eWxvzphXM1
enable password cisco
!
no aaa new-model
clock timezone EST -5
clock summer-time EDT recurring
ip wccp 61
ip wccp 62
ip cef
!
!
ip domain list wanopt4.cisco.com
ip domain name wanopt4.cisco.com
ip name-server 52.1.1.100
!
multilink bundle-name authenticated
!
!
key chain oer-key
 key 1
   key-string WANOPT
!
!
oer master
 policy-rules delayPolicy
 logging
 !
 border 10.0.0.172 key-chain oer-key
  interface Tunnel1 external
  interface Serial6/1:0 external
  interface FastEthernet3/1 internal
 no max range receive
 mode monitor fast
 no resolve delay
 no resolve utilization
 !
 active-probe echo 60.1.1.101
!
oer border
 logging
 local Loopback0
 master 10.0.0.172 key-chain oer-key
!
!
!
!
crypto isakmp policy 10
 encr 3des
 authentication pre-share
 group 2
crypto isakmp key cisco address 0.0.0.0 0.0.0.0
crypto isakmp keepalive 30
!
!
crypto ipsec transform-set gre_set esp-3des esp-sha-hmac
 mode transport
!
crypto ipsec profile gre_prof
 set transform-set gre_set
!
!
username lab privilege 15 password 0 lab
archive
 log config
  hidekeys
!
!
controller T1 6/0
 framing esf
 linecode b8zs
 channel-group 0 timeslots 1-24
!
controller T1 6/1
 framing esf
 linecode b8zs
 channel-group 0 timeslots 1-24
```

```
!
ip ftp username anonymous
ip ftp password nobody@cisco.com
!
class-map match-all NBAR_UDP
 match access-group name UDP
class-map match-any NBAR_FTP
 match protocol ftp
class-map match-all QoS_HTTP
 match ip precedence 2
class-map match-any NBAR_voice
 match protocol sip
 match protocol skinny
class-map match-all QoS_FTP
 match ip precedence 4
class-map match-all QoS_UDP
 match ip precedence 5
class-map match-any NBAR_HTTP
 match protocol http
 match protocol secure-http
class-map match-all QoS_voice
 match ip precedence 6
!
!
policy-map QoS
 class QoS_HTTP
  bandwidth percent 25
 class QoS_UDP
  bandwidth percent 15
 class QoS_voice
  priority percent 10
 class QoS_FTP
  bandwidth percent 10
 class class-default
  bandwidth percent 15
policy-map tunnel
 class class-default
  shape average percent 75
  service-policy QoS
policy-map NBAR
 class NBAR_HTTP
  set dscp cs2
 class NBAR_UDP
  set dscp cs5
 class NBAR_voice
  set dscp cs6
 class NBAR_FTP
  set dscp cs4
 class class-default
  set dscp cs3
policy-map fromWAE
 class QoS_HTTP
 class QoS_UDP
 class class-default
!
!
!
interface Loopback0
 ip address 10.0.0.172 255.255.255.255
!
interface Tunnel1
 ip address 172.20.1.1 255.255.255.0
 no ip redirects
 ip mtu 1400
 ip flow ingress
 ip nhrp authentication nsite
 ip nhrp map 172.20.1.254 10.18.101.2
 ip nhrp map multicast 10.18.101.2
 ip nhrp network-id 101
 ip nhrp holdtime 300
 ip nhrp nhs 172.20.1.254
 ip nhrp cache non-authoritative
 ip nhrp shortcut
 ip tcp adjust-mss 1360
 load-interval 30
 QoS pre-classify
 tunnel source Serial6/0:0.1
 tunnel mode gre multipoint
 tunnel protection ipsec profile gre_prof
!
interface FastEthernet0/0
 description To CAT6K-MAIN:G1/31
 no ip address
 no ip route-cache cef
 shutdown
 duplex full
!
interface FastEthernet2/0
 no ip address
 shutdown
 duplex half
!
interface FastEthernet3/0
 description To WANOPT-WAE512-BR6 via WANOPT-6500-MAIN1:G1/5
 ip address 10.18.56.1 255.255.255.0
 ip flow ingress
 load-interval 30
 duplex full
 speed 100
!
interface FastEthernet3/1
 description To WANOPT-6500-MAIN1:G1/31
 ip address 10.19.6.1 255.255.255.0
 ip wccp 61 redirect in
 ip nbar protocol-discovery
 ip flow ingress
 load-interval 30
 duplex full
 speed 100
 service-policy input NBAR
!
interface Ethernet5/0
 description To CAT35k-BrMAN-1:F0/7
 ip address 100.1.1.172 255.255.255.0
 duplex full
 ntp broadcast client
!
interface Ethernet5/1
 no ip address
 no ip route-cache cef
 no ip route-cache
 shutdown
 duplex half
!
interface Ethernet5/2
 no ip address
 no ip route-cache cef
 no ip route-cache
 shutdown
 duplex half
!
interface Ethernet5/3
 no ip address
 no ip route-cache cef
 no ip route-cache
 shutdown
 duplex half
!
interface Ethernet5/4
 no ip address
 no ip route-cache cef
 no ip route-cache
 shutdown
 duplex half
!
interface Ethernet5/5
 no ip address
 no ip route-cache cef
 no ip route-cache
 shutdown
 duplex half
!
interface Ethernet5/6
 no ip address
 no ip route-cache cef
 no ip route-cache
 shutdown
 duplex half
!
interface Ethernet5/7
 no ip address
 no ip route-cache cef
 no ip route-cache
 shutdown
 duplex half
!
interface Serial6/0:0
 description To WANOPT-7206-INTERNET-PE:S3/1:0
 no ip address
 encapsulation frame-relay
 load-interval 30
 no keepalive
!
interface Serial6/0:0.1 point-to-point
 ip address 10.18.6.2 255.255.255.0
 ip wccp 62 redirect in
 ip flow ingress
 snmp trap link-status
 frame-relay interface-dlci 101
 service-policy output tunnel
!
interface Serial6/1:0
 description To WANOPT-7206-MPLS-PE:S3/1:0
 ip address 10.17.6.2 255.255.255.0
 ip wccp 62 redirect in
 ip flow ingress
 load-interval 30
 service-policy output QoS
!
router eigrp 10
 passive-interface default
 no passive-interface Tunnel1
 network 10.0.0.172 0.0.0.0
 network 10.18.56.0 0.0.0.255
 network 10.19.6.0 0.0.0.255
 network 172.20.1.0 0.0.0.255
 no auto-summary
!
router bgp 176
 no synchronization
 bgp log-neighbor-changes
 network 10.17.6.0 mask 255.255.255.0
 network 10.18.56.0 mask 255.255.255.0
 network 10.19.6.0 mask 255.255.255.0
 neighbor 10.17.6.1 remote-as 103
 no auto-summary
!
ip route 0.0.0.0 0.0.0.0 10.18.6.1
ip route 0.0.0.0 0.0.0.0 10.17.6.1 100
no ip http server
no ip http secure-server
!
ip flow-cache timeout active 1
ip flow-export source FastEthernet3/1
ip flow-export version 5
ip flow-export destination 52.1.1.22 9995
!
!
ip access-list extended UDP
 permit udp any any
 deny   ip any any
ip access-list extended noHTTP
 permit ip any any dscp cs3
 permit ip any any dscp cs5
 deny   ip any any
ip access-list extended onlyHTTP_ip
 permit ip any host 60.1.1.100 log
 permit ip host 60.1.1.100 any log
 deny   ip any any
ip access-list extended onlyRT
 permit ip 10.19.0.0 0.0.255.255 60.1.1.0 0.0.0.255 dscp cs5
ip access-list extended onlyRealTime
 permit ip any any dscp cs5
ip access-list extended onlyTCP
 permit ip any any dscp cs2
ip access-list extended others
 permit ip any any dscp cs3
 permit ip any any dscp cs5
ip access-list extended permitList1
 permit ip any any dscp cs3
!
!
ip prefix-list onlyHTTP seq 5 permit 60.1.1.100/32
ip sla logging traps
ip sla 275
 http get http://www.cisco.com
 timeout 5000
 owner HTTP - 100.1.1.172
 tag WANOPT HTTP ECHO
ip sla schedule 275 life forever start-time now ageout 3600
ip sla 277
 icmp-echo 60.1.1.100
 owner ICMP Echo - 100.1.1.172 - 60.1.1.100
 tag WANOPT ICMP ECHO
ip sla schedule 277 life forever start-time now ageout 3600
ip sla 278
 dns www.cisco.com name-server 52.1.1.100
 timeout 5000
 owner DNS - 100.1.1.172
 tag WANOPT DNS Echo
ip sla schedule 278 life forever start-time now ageout 3600
logging alarm informational
logging 100.1.1.104
snmp-server community closed RW
snmp-server community open RO
snmp-server community public RO
snmp-server community private RW
snmp-server ifindex persist
snmp-server host 100.1.1.21 version 2c public
!
!
!
```

```
!
!
ttycap detected \
at '^' marker.
!
oer-map delayPolicy 10
 match traffic-class prefix-list onlyHTTP
 set mode select-exit best
 set delay threshold 55
 set mode route control
 set mode monitor fast
 set resolve delay priority 1 variance 5
 no set resolve utilization
 set probe frequency 2
!
control-plane
!
!
!
!
gatekeeper
 shutdown
!
!
line con 0
 exec-timeout 0 0
 stopbits 1
line aux 0
 stopbits 1
line vty 0 4
 exec-timeout 0 0
 password lab
 login
line vty 5 20
 password lab
 login
line vty 21 30
 password cisco
 login
!
exception data-corruption buffer truncate
ntp clock-period 17179708

!
webvpn cef
!
end
```

9.12.3 Medium Branch Office with Dual-Homed, Dual-Tiered Branch Routers

A typical medium Branch Office has a number of users on the LAN. To allow for resiliency and scaling, the branch network is designed with two branch routers. Each router is dual homed, with one link connected to SP-WAN and the other is connected to Internet. Such a deployment is shown in Figure 9-38.

Figure 9-38. Typical Medium Branch Office

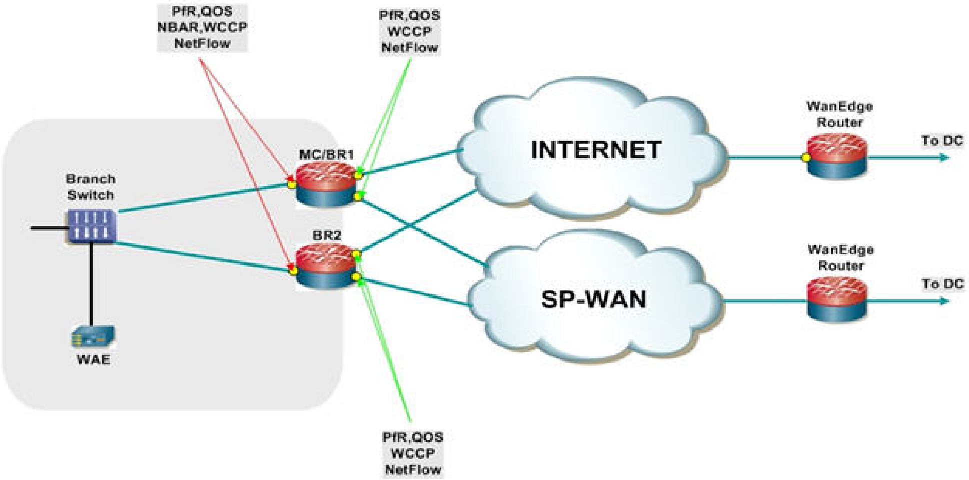

The branch routers have the following protocols deployed:

1. NetFlow on all LAN/WAN links
2. NBAR to do protocol discovery and QoS to do marking
3. QoS on the exit interface to do congestion avoidance and management
4. WCCP on the router to do TCP optimization
5. PfR on the routers to do path optimization. One branch router doubles as both Master Controller and Border Router
6. (Optional) DMVPN tunnel over Internet to the corporate HQ to provide security

9.12.3.1 Sample Branch Router Configuration (PfR-Master Controller/Border Router)

```
Current configuration : 7977 bytes
!
! Last configuration change at 13:14:35 EST Wed Jan 30 2008
! NVRAM config last updated at 13:17:03 EST Wed Jan 30 2008
!
upgrade fpd auto
version 12.4
no service pad
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname WANOPT-7206-BR5
!
boot-start-marker
boot system flash disk0:c7200-adventerprisek9-mz.124-15.T2.fc3
boot-end-marker
!
logging buffered 16000000
no logging console
enable password lab
!
no aaa new-model
clock timezone EST -5
clock summer-time EDT recurring
ip wccp 61
ip wccp 62
ip cef
!
ip domain list wanopt4.cisco.com
ip domain name wanopt4.cisco.com
ip name-server 52.1.1.100
!
multilink bundle-name authenticated
!
!
key chain oer-key
 key 1
   key-string WANOPT
!
!
oer master
 policy-rules delayPolicy
 logging
 !
 border 10.0.0.171 key-chain oer-key
  interface Serial6/1:0 external
  interface Serial6/0:0 external
  interface FastEthernet3/1.3051 internal
  interface FastEthernet3/1.4051 internal
 !
 border 10.0.0.165 key-chain oer-key
  interface Serial0/2/1:0 external
  interface Serial0/2/0:0 external
  interface GigabitEthernet0/0.4051 internal
  interface GigabitEthernet0/0.3051 internal
 !
 learn
  delay
  periodic-interval 0
  monitor-period 1
  prefixes 1000
  aggregation-type prefix-length 32
 no max range receive
 mode monitor fast
 no resolve delay
 no resolve utilization
 !
 active-probe echo 60.1.1.100
!
oer border
 logging
 local Loopback0
 master 10.0.0.171 key-chain oer-key
!
!
!
```

```
!
username lab privilege 15 password 0 lab
archive
 log config
  hidekeys
!
!
controller T1 6/0
 framing esf
 linecode b8zs
 channel-group 0 timeslots 1-24
!
controller T1 6/1
 framing esf
 linecode b8zs
 channel-group 0 timeslots 1-24
!
controller T1 6/2
 framing esf
 linecode b8zs
 channel-group 0 timeslots 1-24
!
controller T1 6/3
 framing esf
 linecode b8zs
 channel-group 0 timeslots 1-24
!
ip ftp source-interface Ethernet5/2
ip ftp username anonymous
ip ftp password nobody@cisco.com
!
class-map match-all NBAR_UDP
 match access-group name UDP
class-map match-any NBAR_FTP
 match protocol ftp
class-map match-all QoS_HTTP
 match ip precedence 2
class-map match-any NBAR_voice
 match protocol sip
 match protocol skinny
class-map match-all QoS_FTP
 match ip precedence 4
class-map match-all QoS_UDP
 match ip precedence 5
class-map match-any NBAR_HTTP
 match protocol http
 match protocol secure-http
class-map match-all QoS_voice
 match ip precedence 6
class-map match-all ramki
!
!
policy-map NBAR
 class NBAR_HTTP
  set dscp cs2
 class NBAR_UDP
  set dscp cs5
 class NBAR_voice
  set dscp cs6
 class NBAR_FTP
  set dscp cs4
 class class-default
  set dscp cs3
policy-map QoS
 class QoS_HTTP
  bandwidth percent 25
 class QoS_UDP
  bandwidth percent 15
 class QoS_voice
  priority percent 10
 class QoS_FTP
  bandwidth percent 10
 class class-default
  bandwidth percent 15
!
!
!
interface Loopback0
 ip address 10.0.0.171 255.255.255.255
 no ip route-cache cef
 no ip route-cache
!
interface FastEthernet0/0
 description To CAT6K-MAIN:g1/43
 no ip address
 no ip route-cache cef
 no ip route-cache
 shutdown
 duplex full
!
interface FastEthernet3/0
 description To WANOPT-WAE512-BR5 via WANOPT-6500-MAIN1:G1/3
 ip address 10.18.55.65 255.255.255.192
 ip flow ingress
 duplex full
 speed 100
!
interface FastEthernet3/1
 description To AGILENT via WANOPT-6500-MAIN1:G1/30
 no ip address
 load-interval 30
 duplex full
 speed 100
!
interface FastEthernet3/1.3051
 encapsulation dot1Q 3051
 ip address 10.19.5.2 255.255.255.0
 ip wccp 61 redirect in
 ip flow ingress
 glbp 1 ip 10.19.5.1
 glbp 1 weighting 80
 glbp 1 load-balancing weighted
!
interface FastEthernet3/1.4051
 encapsulation dot1Q 4051
 ip address 110.19.19.1 255.255.255.0
 ip wccp 61 redirect in
 ip flow ingress
 shutdown
 service-policy input NBAR
!
interface Ethernet5/0
 description To CAT35k-BrMAN-1:F0/6
 ip address 100.1.1.171 255.255.255.0
 no ip route-cache cef
 no ip route-cache
 duplex full
 ntp broadcast client
!
interface Ethernet5/1
 description To WANOPT-6500-MAIN1:G1/34
 ip address 10.55.55.2 255.255.255.252
 duplex full
!
interface Ethernet5/2
 description To CAT6K-MAIN:g1/39
 no ip address
 no ip route-cache cef
 shutdown
 duplex full
!
interface Ethernet5/3
 description To CAT6K-MAIN:g1/41
 no ip address
 no ip route-cache cef
 no ip route-cache
 shutdown
 duplex full
!
interface Serial6/0:0
 description To WANOPT-7206-INTERNET-PE:S3/0:0
 ip address 10.18.5.66 255.255.255.192
 ip wccp 62 redirect in
 ip flow ingress
 load-interval 30
 service-policy output QoS
!
interface Serial6/1:0
 description To WANOPT-7206-MPLS-PE:S3/0:0
 ip address 10.17.5.66 255.255.255.192
 ip wccp 62 redirect in
 ip flow ingress
 load-interval 30
 service-policy output QoS
!
interface Serial6/2:0
 no ip address
 no ip route-cache cef
 no ip route-cache
!
interface Serial6/3:0
 no ip address
 no ip route-cache cef
 no ip route-cache
!
!
router bgp 175
 no synchronization
 bgp log-neighbor-changes
 network 10.0.0.171 mask 255.255.255.255
 network 10.17.5.64 mask 255.255.255.192
 network 10.18.55.64 mask 255.255.255.192
 network 10.19.5.0 mask 255.255.255.0
 network 10.55.55.0 mask 255.255.255.0
 network 110.19.19.0 mask 255.255.255.0
 neighbor 10.17.5.65 remote-as 103
 neighbor 10.55.55.1 remote-as 175
 no auto-summary
!
ip route 0.0.0.0 0.0.0.0 10.18.5.65
ip route 0.0.0.0 0.0.0.0 10.17.5.65 50
no ip http server
no ip http secure-server
!
ip flow-cache timeout active 1
ip flow-export source FastEthernet3/1.3051
ip flow-export version 5
ip flow-export destination 52.1.1.22 9995
!
!
ip access-list extended UDP
 permit udp any any
 deny   ip any any
ip access-list extended noHTTP
```

```
 permit ip any any dscp cs3
 permit ip any any dscp cs5
 deny   ip any any
ip access-list extended only35
 permit ip any any precedence critical log
 permit ip any any precedence flash log
 permit ip any any precedence immediate log
 permit ip any any log
ip access-list extended onlyRealTime
 permit ip any any dscp cs5
ip access-list extended ramki
!
!
ip prefix-list onlyHTTP seq 5 permit 60.1.1.100/32
ip sla 1
 http get http://www.cisco.com
 timeout 5000
 owner HTTP - 100.1.1.171
 tag WANOPT HTTP ECHO
ip sla schedule 1 life forever start-time now ageout 3600
ip sla 3
 icmp-echo 60.1.1.100
 owner ICMP Echo - 100.1.1.171 - 60.1.1.100
 tag WANOPT ICMP ECHO
ip sla schedule 3 life forever start-time now ageout 3600
ip sla 4
 dns www.cisco.com name-server 52.1.1.100
 timeout 5000
 owner DNS - 100.1.1.171
 tag WANOPT DNS
ip sla schedule 4 life forever start-time now ageout 3600
logging alarm informational
logging 100.1.1.104
access-list 1 permit 10.19.5.0 0.0.0.255 log
access-list 1 permit 110.19.5.0 0.0.0.255 log
access-list 1 permit any log
access-list 2 permit 10.19.5.0 0.0.0.255
access-list 2 deny   any
access-list 10 permit 110.19.5.0 0.0.0.255
access-list 10 permit 10.18.55.0 0.0.0.255
access-list 101 permit ip any 10.19.5.0 0.0.0.255 log
access-list 101 permit ip any 110.19.5.0 0.0.0.255 log
access-list 101 permit ip any any log
snmp-server community closed RW
snmp-server community open RO
snmp-server community public RO
snmp-server community private RW
snmp-server ifindex persist
snmp-server host 100.1.1.21 version 2c public
!
!
!
!
!
!
oer-map delayPolicy 10
 match traffic-class prefix-list onlyHTTP
 set mode select-exit best
 set delay threshold 55
 set mode route control
 set mode monitor fast
 set resolve delay priority 1 variance 20
 set probe frequency 10
!
control-plane
!
!
!
```

```
gatekeeper
 shutdown
!
!
line con 0
 exec-timeout 0 0
 stopbits 1
line aux 0
 stopbits 1
line vty 0 4
 exec-timeout 0 0
 password lab
 login
line vty 5 20
 password lab
 login
line vty 21 30
 password cisco
 login
!
ntp clock-period 17180110

!
webvpn cef
!
end
```

9.12.3.2 Sample Configuration of Branch Router (PfR-Border Router)

```
Current configuration : 5321 bytes
!
! Last configuration change at 13:14:52 EST Wed Jan 30 2008
! NVRAM config last updated at 13:12:12 EST Wed Jan 30 2008
!
version 12.4
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
no service password-encryption
!
hostname WANOPT-3825-BR5
!
boot-start-marker
boot system flash:c3825-adventerprisek9-mz.124-15.T2.fc3
boot-end-marker
!
enable password lab
!
no aaa new-model
clock timezone EST -5
clock summer-time edt recurring
no network-clock-participate wic 2
!
!
ip wccp 61
ip wccp 62
ip cef
!
!
no ip domain lookup
ip host aswan 24.1.1.1
!
multilink bundle-name authenticated
!
voice-card 0
 no dspfarm
!
!
```

```
!
key chain oer-key
 key 1
   key-string WANOPT
!
!
oer border
 logging
 local Loopback0
 master 10.0.0.171 key-chain oer-key
!
!
archive
 log config
  hidekeys
!
!
controller T1 0/2/0
 framing esf
 linecode b8zs
 channel-group 0 timeslots 1-24
!
controller T1 0/2/1
 framing esf
 linecode b8zs
 channel-group 0 timeslots 1-24
!
!
class-map match-all NBAR_UDP
 match access-group name UDP
class-map match-any NBAR_FTP
 match protocol ftp
class-map match-all QoS_HTTP
 match ip precedence 2
class-map match-any QCLASS_COS1
 match  dscp 47
 match  dscp ef
class-map match-any NBAR_voice
 match protocol sip
 match protocol skinny
class-map match-all QoS_FTP
 match ip precedence 4
class-map match-all QoS_UDP
 match ip precedence 5
class-map match-any NBAR_HTTP
 match protocol http
 match protocol secure-http
class-map match-all QoS_voice
 match ip precedence 6
!
!
policy-map NBAR
 class NBAR_HTTP
  set dscp cs2
 class NBAR_UDP
  set dscp cs5
 class NBAR_voice
  set dscp cs6
 class NBAR_FTP
  set dscp cs4
 class class-default
  set dscp cs3
policy-map QoS
 class QoS_HTTP
  bandwidth percent 25
 class QoS_UDP
  bandwidth percent 15
 class QoS_voice
```

```
  priority percent 10
 class QoS_FTP
  bandwidth percent 10
 class class-default
  bandwidth percent 15
!
!
interface Loopback0
 ip address 10.0.0.165 255.255.255.255
!
interface GigabitEthernet0/0
 description To AGILENT via WANOPT-6500-MAIN1:G1/29
 no ip address
 load-interval 30
 duplex full
 speed 1000
 media-type rj45
!
interface GigabitEthernet0/0.3051
 encapsulation dot1Q 3051
 ip address 10.19.5.3 255.255.255.0
 ip flow ingress
 glbp 1 ip 10.19.5.1
 glbp 1 weighting 20
 glbp 1 load-balancing weighted
!
interface GigabitEthernet0/0.3053
 encapsulation dot1Q 3053
 ip address 10.18.55.68 255.255.255.192
!
interface GigabitEthernet0/0.3055
 encapsulation dot1Q 3055
 ip address 10.55.55.1 255.255.255.252
!
interface GigabitEthernet0/0.4051
 encapsulation dot1Q 4051
 ip address 110.19.5.3 255.255.255.0
 ip wccp 61 redirect in
 ip flow ingress
 shutdown
 glbp 2 ip 110.19.5.1
 glbp 2 load-balancing weighted
 service-policy input NBAR
!
interface GigabitEthernet0/1
 description To CAT35k-BrMAN-1:F0/5
 ip address 100.1.1.165 255.255.255.0
 load-interval 30
 duplex full
 speed 100
 media-type rj45
 ntp broadcast client
 no keepalive
!
interface Serial0/2/0:0
 description To WANOPT-7206-INTERNET-PE:S5/4:0
 ip address 10.18.5.2 255.255.255.192
 ip wccp 62 redirect in
 ip flow ingress
 encapsulation frame-relay
 frame-relay map ip 10.18.5.1 101 broadcast
 service-policy output QoS
!
interface Serial0/2/1:0
 description To WANOPT-7206-MPLS-PE:S5/4:0
 ip address 10.17.5.2 255.255.255.192
 ip wccp 62 redirect in
 ip flow ingress
 service-policy output QoS
!
interface Integrated-Service-Engine2/0
 ip address 10.18.55.1 255.255.255.192
 shutdown
 service-module external ip address 10.18.55.67 255.255.255.192
 service-module ip address 10.18.55.2 255.255.255.192
 service-module ip default-gateway 10.18.55.65
 no keepalive
!
router bgp 175
 no synchronization
 bgp log-neighbor-changes
 network 10.0.0.165 mask 255.255.255.255
 network 10.17.5.0 mask 255.255.255.192
 network 10.18.55.64 mask 255.255.255.192
 network 10.19.5.0 mask 255.255.255.0
 network 10.55.55.0 mask 255.255.255.0
 neighbor 10.17.5.1 remote-as 103
 neighbor 10.55.55.2 remote-as 175
 no auto-summary
!
ip route 0.0.0.0 0.0.0.0 10.18.5.1
ip route 0.0.0.0 0.0.0.0 10.17.5.1 50
!
!
no ip http server
no ip http secure-server
!
ip access-list extended UDP
 permit udp any any
 deny   ip any any
!
!
ip prefix-list only110 seq 5 permit 110.19.19.2/32
ip sla 1
 http get http://www.cisco.com
 timeout 5000
 owner http - 60.1.1.100
 tag Cisco
ip sla schedule 1 life forever start-time now ageout 3600
logging 100.1.1.104
access-list 102 permit ip 10.19.5.0 0.0.0.255 60.1.1.0 0.0.0.255
access-list 102 deny   ip any any
snmp-server community closed RW
snmp-server community open RO
snmp-server community public RO
snmp-server community private RW
snmp-server ifindex persist
snmp-server host 100.1.1.21 version 2c public
!
!
!
control-plane
!
!
!
alias exec sib show ip interface brief
privilege exec level 0 terminal monitor
privilege exec level 0 terminal
!
line con 0
 exec-timeout 0 0
line aux 0
line 130
 no activation-character
 no exec
 transport preferred none
 transport input all
 transport output lat pad telnet rlogin lapb-ta mop udptn v120 ssh
line vty 0 4
 exec-timeout 0 0
 password lab
 login
 autocommand  term mon
 autocommand-options nohangup
!
scheduler allocate 20000 1000
ntp clock-period 17179645

!
webvpn cef
!
end
```

9.12.4 Large Branch Offices with Dual-Homed, Dual-Tiered Branch Routers

Larger branch offices need more resiliency and scaling than medium sized branches. In addition to using a higher capacity/platform router and WAEs, typical larger branch designs also use WAE clusters to provide WAE redundancy. Such branches have two branch routers, each dual-homed to two SP WANs (or one SP WAN) and two WAEs shared by the branch routers. Along with these branch routers, these branches also have a separate router to act as MC for PfR. Figure 9-39 shows such a deployment.

Figure 9-39. Typical Large Branch Office

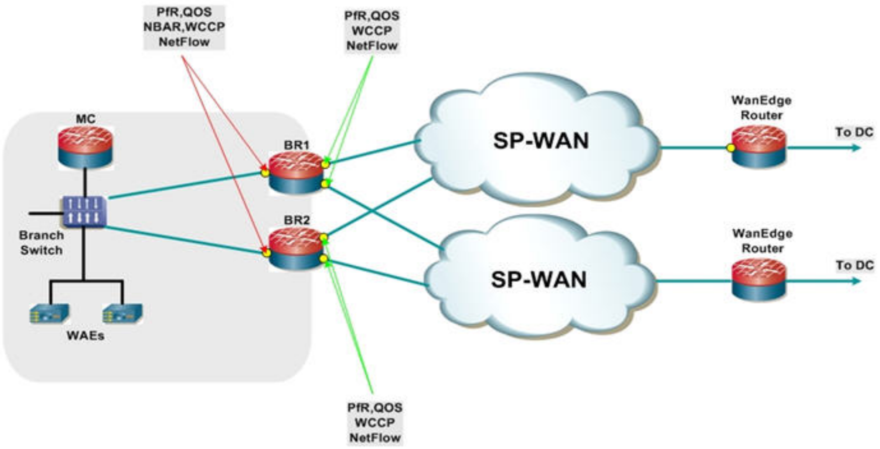

The branch routers have the following protocols deployed:

1. NetFlow on all LAN and WAN links
2. NBAR to do protocol discovery and QoS to do marking
3. QoS on the exit interface to do congestion avoidance and management
4. WCCP on the router to do TCP optimization
5. PfR on the routers to do path optimization

9.13 Suggested Code Versions

Table 9-3. Recommended Software Versions

Platform	Version
Cisco 28xx/38xx ISRs	12.4(15)T3
Cisco 72xx-VXRs	12.4(15)T3
Cisco WAE -511/611/7326 appliances/	
NM WAEs	4.0.13

9.14 Data Center Design

Content is consolidated for user access in the data center. A typical data center design, shown in Figure 9-40) comprises a core, aggregation, and access layer.

- The Core/WAN focuses on bringing the packets into the data center reliably and quickly.
- The aggregation layer, where the network services are located, often includes firewall, load balancing, and SSL offload.
- The access layer typically includes web servers, databases, and middleware.

Figure 9-40. Typical Data Center Design

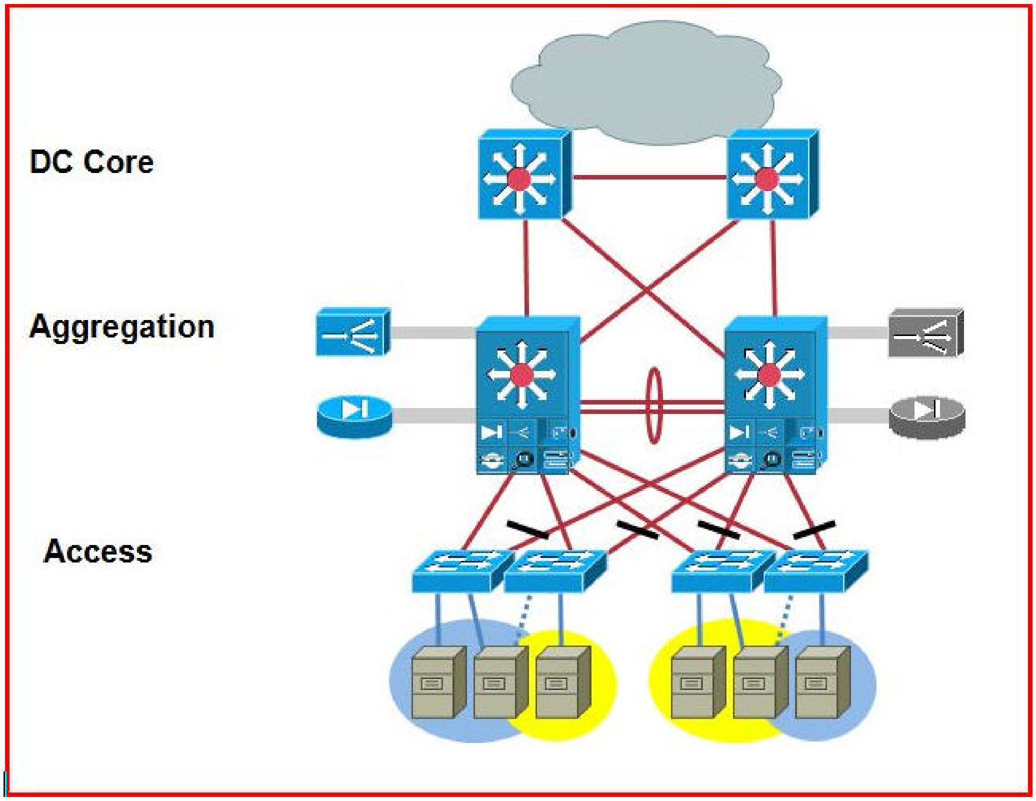

When it comes to the specific insertion of WAN optimization components, the WAAS cluster is placed at the WAN edge/core or the aggregation layer. This WAAS cluster placement can determine which servers traffic can be optimized for or affect the network services.

If the WAAS cluster is located closer to the WAN, the cluster can offer optimized content from many servers. On the other hand, if the WAAS cluster is located at the aggregation layer, it can only optimize content from the servers in its aggregation module.

As far as network services, the DC traffic could be optimized. This could affect how firewalls or other network services at the edge of the DC inspect traffic. These services might require code changes or configuration changes. Network services that only see nonoptimized traffic would not require changes.

9.14.1 FWSM

The Firewall Services Module (FWSM) is a high-speed, integrated firewall module for Cisco Catalyst 6500 switches. FWSM can be deployed in pairs to provide intrachassis or interchassis stateful failover services to help ensure resilient network protection for the data center.

9.14.1.1 FWSM Deployment Options

The two FWSM deployment options are:

1. Downstream of the DC WAAS
    - Nonoptimized traffic - inspects L4 and L7
2. Upstream of the DC WAAS
    - Inspects optimized and non-optimized traffic
    - Optimized traffic - inspects only TCP at L4
    - Nonoptimized traffic - inspects L4 and L7

WAAS alters the original TCP session through Transport Flow Optimization (TFO) and application data though DRE. DRE helps to improve efficiency and reduces unnecessary bandwidth consumption. On the other hand, TFO optimizes TCP to enable better performance and efficiency in WAN environments. FWSM introduced enhancements in 3.2(1) to handle TCP traffic optimized by WAAS. With these enhancements, L4 inspection can now be performed on a WAAS optimized TCP session. However, FWSM features that proxy the TCP connection stop WAAS from optimizing flows.

9.14.1.2 TCP Options

The WAAS implementation can decide which TCP connection to optimize on a TCP port basis. WAAS typically does not optimize TCP control channels (FTP control channel, for example). Any session it does optimize has the 0x21 option. When the option is seen, FWSM inspects at L4. Connections that are not optimized do not have option 0x21. Therefore, FWSM can apply all inspections to the connections.

9.14.1.3 Sequence Number Change

The initial SYN packet from the client is forwarded to the server with no changes in sequence number. The branch WAAS device intercepts this session and attaches TCP option 0x21 to the SYN. The DC WAAS device later intercepts this session and responds to the branch WAAS after the DC WAAS sees a SYN-ACK from the server side.

The branch WAAS device bumps up the TCP sequence number by 2^31 when it sends the final ACK of the TCP handshake to the server side. This ACK also contains option 0x21. FWSM then knows that the session will be optimized.

9.14.1.4 Configuration

FWSM uses Modular Policy Framework (MPF) to configure which class of traffic must be inspected. MPF uses class maps to identify traffic and policy maps are used to apply actions. MPF policy map actions include TCP connection limits/timeouts and application inspection.

The following configuration includes a default class map that the FWSM uses in the default global policy. It is called `inspection_default` and matches the default inspection traffic. To enable WAAS application inspection, use the `inspect waas` command.

```
class-map inspection_default
  match default-inspection-traffic
!
policy-map global_policy
  class inspection_default
    inspect waas
!
service-policy global_policy global
```

The `inspect waas` command, when applied to the default inspection class, checks all TCP traffic for option 0x21 and handles the 2G Sequence Space Bump. TCP option SACK, timestamp and window scale are also allowed even if TCP SYN/SYN-ACK did not contain those options. For WAAS optimized flows, FWSM uses a window scale of 7.

9.14.1.5 Show Commands

The `show conn` command displays the connection state for the designated connection type. A “W” indicates that the TCP flow is a WAAS Session.

```
wanopt-fwsm1/NSITE# show conn
79 in use, 130 most used
 Network Processor 1 connections
TCP out 10.19.7.140:62645 in 60.1.1.100:80 idle 0:00:00 Bytes 2647 FLAGS - UBfOIW
 Network Processor 2 connections
TCP out 10.19.7.140:62646 in 60.1.1.100:80 idle 0:00:00 Bytes 4623 FLAGS - UBOIW
TCP out 10.19.7.140:62647 in 60.1.1.100:80 idle 0:00:00 Bytes 2340 FLAGS -
FRdUBfrOIW
TCP out 10.19.7.140:62648 in 60.1.1.100:80 idle 0:00:00 Bytes 5844 FLAGS – UBOIW
```

9.14.2 WAAS Catalyst 6500 Load Balancing

WCCPv2 is a protocol used in the data center to send relevant TCP traffic to the Core WAAS cluster. The WAAS cluster is positioned on a common VLAN at either the WAN edge or the at the distribution switch layer.

WCCPv2 introduced several concepts to improve traffic forwarding:

- WCCP routers and clients were bundled together in service groups with multiple routers per service group being supported
- Protocol messages were allowed to use multicast

- Return traffic by clients was supported as well as the negotiation of forwarding, assignment, and return methods between the router and WCCP
- Mask assignment with L2 forwarding and return were introduced allowing hardware forwarding of traffic packets

WAAS service groups 61 and 62:

- Service group 61—Redirects TCP packets to a WAE device and distributes load based on the source IP address of the flow
- Service group 62—Redirects packets to a WAE device and distributes load based on the IP destination address of the flow.

WCCP WAE devices leverage a combination of both groups so all traffic is redirected to the WAEs.

9.14.2.1 WCCP Components

WCCP has three components: the Assignment method, the Redirection method, and the Return method. The Assignment method determines which WCCP appliance receives redirected traffic; the Redirection method refers to how traffic is sent to the WCCP appliance; the Return method determines how WCCP bypass traffic is handled.

9.14.2.1.1 Assignment Method

The Assignment method determines how traffic is load balanced across multiple WCCP appliances. The introduction of WCCPv2 enabled the negotiation of Mask Assignment per service group.

Mask-based assignment, as an ingress feature, can use an ACL redirect-adjacency entry in the ACL table. The mask, which is programmed into the ACL Ternary Content Addressable Memory (TCAM) before packet forwarding, does not require the use of the NetFlow table or software processing. The WCCP-designated appliance chooses a number of hash-buckets, assigns an address mask and WCCP appliance to each bucket.

After these assignments are done, the Supervisor programs one TCAM entry and one adjacency for each bucket. This enables the redirecting of packets that match the address mask to the associated WCCP appliance through an L2 rewrite. If WCCP is configured as an egress feature, ACL redirect-adjacencies are not hardware supported, and the first packet of a flow is sent to software for processing. Once the proper redirect-adjacency is determined, it is programmed into the NetFlow hardware, where the entry points to an adjacency that performs either an L2 rewrite or GRE encapsulation.

9.14.2.1.2 Redirection Method

Redirection handles how traffic is sent to the WCCP appliance. Using L3 redirection, each WCCP packet is encapsulated in a GRE header marked with a protocol type of 0x883E followed by a four-octet WCCP redirect header. The WCCP packets are subsequently sent to the WCCP appliance.

With WCCP v2, “accelerated WCCP” or L2 redirection was added to take advantage of hardware switching platforms. With L2 redirection, the WCCP appliance must be L2 adjacent (on the same L2 VLAN); redirected traffic is forwarded through normal switching with a rewritten MAC destination address.

9.14.2.1.3 L2 Forwarding Method Detail

With L2 forwarding, WCCP appliances within a service group are part of the same subnet and L2 adjacent to the supporting switch. This supports higher throughput and low latency deployment. The packet is rewritten with the source MAC set to the router and the destination MAC set to the WCCP device.

L2 Forwarding Breakdown

Ingress – L2 redirection + Mask Assignment (Full Hardware Processing)

Mask assignments further enhance the performance of L2 redirection. When configured on ingress, L2 Redirection and Mask assignment is the most efficient WCCP method on the Catalyst 6500. All packet traffic, including the first packet, is switched in hardware and no software processing is needed.

The current Catalyst platform supports a 7-bit mask, with default mask of 0x1741 on the source IP address. Fine-tuning of the mask can yield better traffic distribution to the WAEs. For example, if a network uses only 191.x.x.x address space, the most significant bit can be reused on the last three octets, such as 0x0751, because the leading octet (191) is always the same.

The following example is a show output from `show ip wccp 61 detail` with a mask of 0x1000.

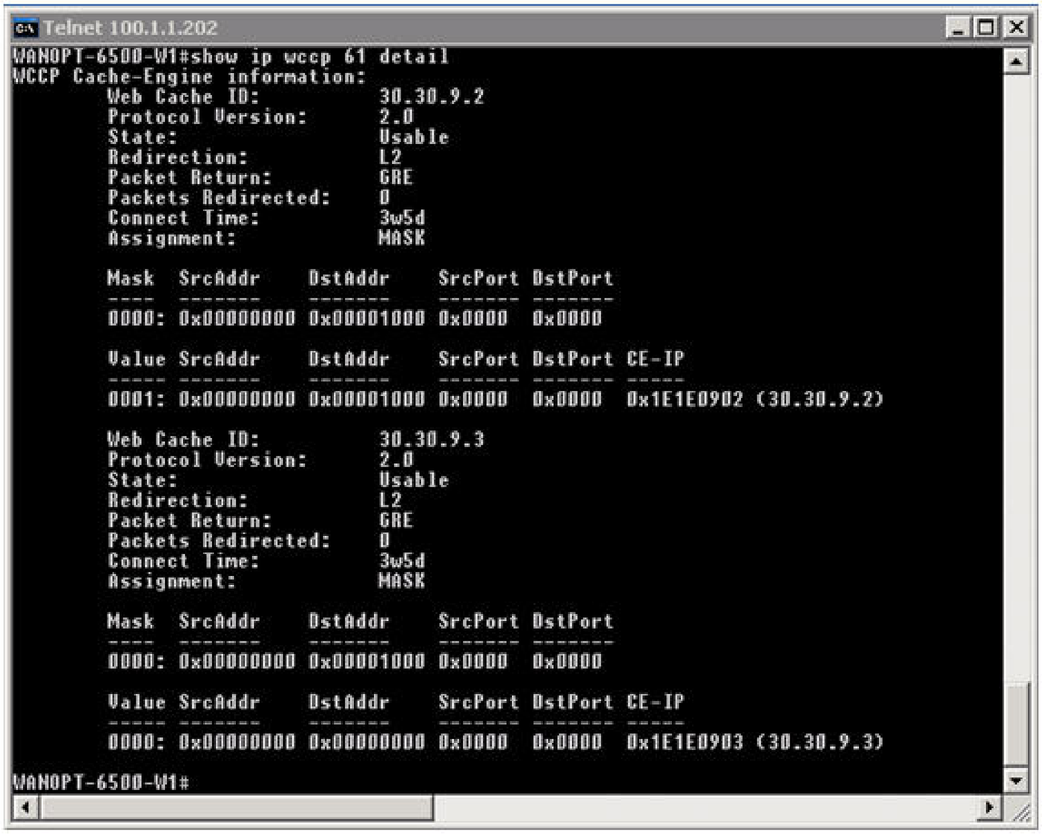

```
WANOPT-6500-W1#show ip wccp 61 detail
WCCP Cache-Engine information:
        Web Cache ID:           30.30.9.2
        Protocol Version:       2.0
        State:                  Usable
        Redirection:            L2
        Packet Return:          GRE
        Packets Redirected:     0
        Connect Time:           3w5d
        Assignment:             MASK

        Mask  SrcAddr    DstAddr    SrcPort DstPort
        ----  -------    -------    ------- -------
        0000: 0x00000000 0x00001000 0x0000  0x0000

        Value SrcAddr    DstAddr    SrcPort DstPort CE-IP
        ----- -------    -------    ------- ------- -----
        0001: 0x00000000 0x00001000 0x0000  0x0000  0x1E1E0902 (30.30.9.2)

        Web Cache ID:           30.30.9.3
        Protocol Version:       2.0
        State:                  Usable
        Redirection:            L2
        Packet Return:          GRE
        Packets Redirected:     0
        Connect Time:           3w5d
        Assignment:             MASK

        Mask  SrcAddr    DstAddr    SrcPort DstPort
        ----  -------    -------    ------- -------
        0000: 0x00000000 0x00001000 0x0000  0x0000

        Value SrcAddr    DstAddr    SrcPort DstPort CE-IP
        ----- -------    -------    ------- ------- -----
        0000: 0x00000000 0x00000000 0x0000  0x0000  0x1E1E0903 (30.30.9.3)

WANOPT-6500-W1#
```

The WAE and Catalyst negotiate which redirect and return method to use when the service group is formed. There can be many access VLANs on the aggregation switches and redirection is configured on all VLANs that need optimization. L2 switching ports, including trunk ports, are not eligible for redirection.

L3 Forwarding Method Detail

WCCP L3 operation involves the use of GRE as the encapsulation method. Redirected packets are encapsulated in a GRE header with a protocol type of 0x883e, along with a 4-byte WCCP Redirection header that includes a service ID and hash bucket matched. Using GRE enables the WCCP client to be separated from the WCCP switch over multiple L3 hops.

Figure 9-41. L3 Forwarding Method Detail

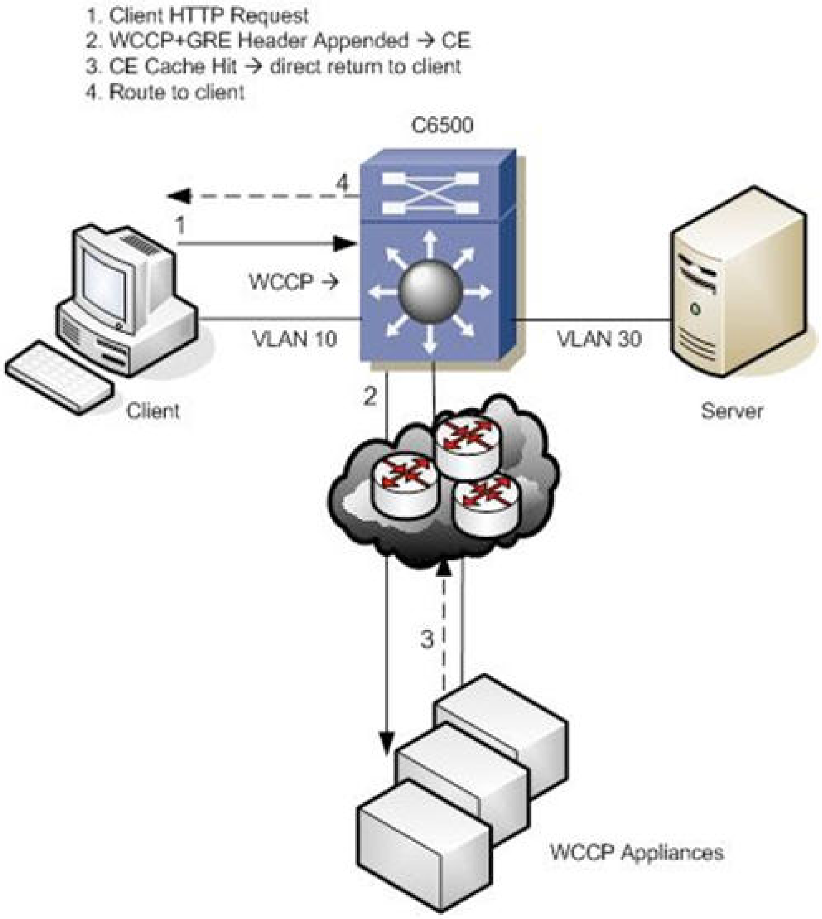

L3 Forwarding Method Breakdown

Ingress – L3 (GRE) redirection + Mask Assignment (Full Hardware Processing -Sup32/Sup720 only)

When using the mask assignment method on the ingress, the initial and subsequent packets are both forwarded in hardware. Normally, the WCCP appliance returns traffic directly to the client, but if the WCCP device is unable to process the request, then it is encapsulated in GRE and returned to the Catalyst 6500. PFC 3 cannot process GRE protocol type 0x833E in hardware causing it to be processed in software instead.

Method Summary

While WCCP on the Catalyst 6500 offers many options, hardware acceleration is available only with ingress L2 redirection with mask assignment and ingress L3 (GRE) redirection with mask assignment. From a hardware perspective, all egress WCCP configurations require some software processing and have CPU utilization impact. Software processing is also required on the ingress when using the hash assignment method. Note that L2 redirection with mask assignments is the recommended deployment for Catalyst 6500 WCCP.

9.14.2.2 Configuring L2 Redirection with Catalyst 6500

On Catalyst 6500:

```
switch(config)# interface Vlan76      ! Dedicated VLAN for WAAS cluster
switch(config)# ip address 30.30.9.1 255.255.255.0
switch(config)# load-interval 120

switch(config)# int GigabitEthernet 0/1
switch(config-if)# switchport
switch(config-if)# switchport mode access
switch(config-if)# switchport access vlan 76 (shared VLAN of WCCP devices)
switch(config-if)# spanning-tree portfast
switch(config-if)# spanning-tree bpdugaurd enable
```

On WAE device:

```
Configure the interface:
wae(config)# int GigabitEthernet 0/1
wae(config-if)# ip address 30.30.9.3
wae(config)# primary-interface GigabitEthernet 0/1
wae(config)# wccp tcp-promiscuous mask src-ip-mask 0x0 dst-ip-mask 0x1000
wae(config)# wccp tcp-promiscuous router-list-num 1 l2-redirect mask-assign
```

Configuration Show Commands:

```
wae# sh wccp masks tcp-promiscuous  view currently configured masks
wae# sh wccp routers  show current routers the WCCP device is communicating with
```

Output of show WCCP mask command:

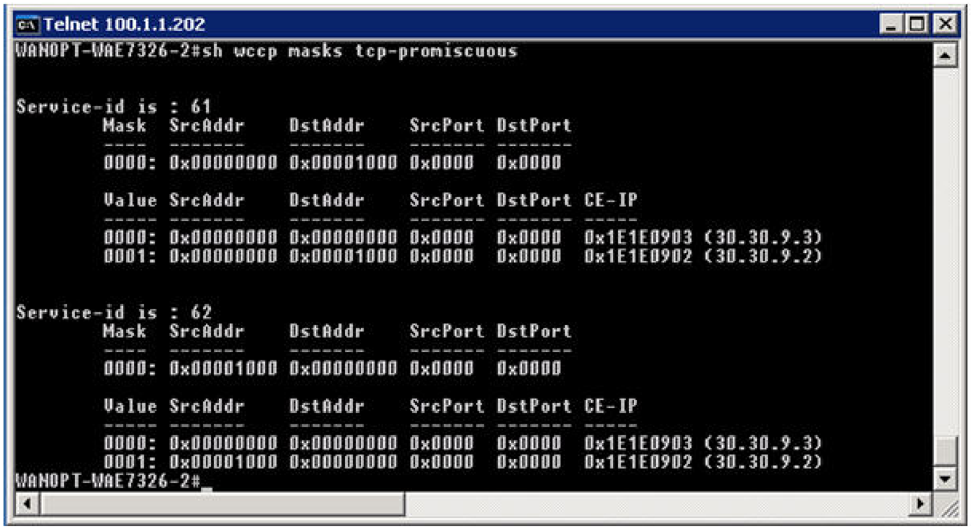

```
WANOPT-WAE7326-2#sh wccp masks tcp-promiscuous


Service-id is : 61
        Mask  SrcAddr    DstAddr    SrcPort DstPort
        ----  -------    -------    ------- -------
        0000: 0x00000000 0x00001000 0x0000  0x0000

        Value SrcAddr    DstAddr    SrcPort DstPort CE-IP
        ----- -------    -------    ------- ------- -----
        0000: 0x00000000 0x00000000 0x0000  0x0000  0x1E1E0903 (30.30.9.3)
        0001: 0x00000000 0x00001000 0x0000  0x0000  0x1E1E0902 (30.30.9.2)


Service-id is : 62
        Mask  SrcAddr    DstAddr    SrcPort DstPort
        ----  -------    -------    ------- -------
        0000: 0x00001000 0x00000000 0x0000  0x0000

        Value SrcAddr    DstAddr    SrcPort DstPort CE-IP
        ----- -------    -------    ------- ------- -----
        0000: 0x00000000 0x00000000 0x0000  0x0000  0x1E1E0903 (30.30.9.3)
        0001: 0x00001000 0x00000000 0x0000  0x0000  0x1E1E0902 (30.30.9.2)
WANOPT-WAE7326-2#_
```

Output of show WCCP router command:

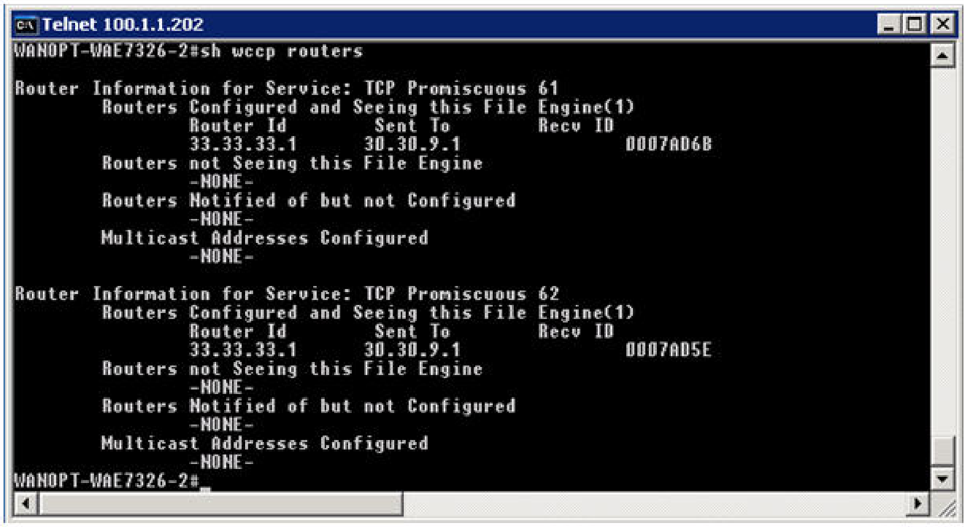

```
WANOPT-WAE7326-2#sh wccp routers

Router Information for Service: TCP Promiscuous 61
        Routers Configured and Seeing this File Engine(1)
                Router Id        Sent To         Recv ID
                33.33.33.1      30.30.9.1               0007AD6B
        Routers not Seeing this File Engine
                -NONE-
        Routers Notified of but not Configured
                -NONE-
        Multicast Addresses Configured
                -NONE-

Router Information for Service: TCP Promiscuous 62
        Routers Configured and Seeing this File Engine(1)
                Router Id        Sent To         Recv ID
                33.33.33.1      30.30.9.1               0007AD5E
        Routers not Seeing this File Engine
                -NONE-
        Routers Notified of but not Configured
                -NONE-
        Multicast Addresses Configured
                -NONE-
WANOPT-WAE7326-2#_
```

9.14.3 ACE SSL

Secure Sockets Layer (SSL) provides a secure transport for HTTP applications. The SSL layer is logically situated between the application layer and TCP. The purpose of an SSL connection is to securely exchange information between two endpoints.

SSL provides the following security elements:

- Confidentiality: encryption of traffic using ciphers
- Authentication: verification of identity using certificates
- Integrity: An optional Message Authentication Code (MAC)prevents data tampering

9.14.3.1 SSL Offload

SSL is used to secure many web transactions. The Application Control Engine (ACE) can encrypt and decrypt SSL traffic. For SSL offload, the ACE operates as a virtual SSL server where all inbound SSL traffic from a client terminates at the ACE. After the connection is terminated, the ACE decrypts the ciphertext from the client and sends the data as cleartext to an HTTP server.

The ACE module supports SSL version 3 and Transport Layer Security (TLS) version 1. A typical SSL session with the ACE requires encryption ciphers to establish and maintain the secure connection. Cipher suites provide the cryptographic algorithms required by the ACE to perform key exchange, authentication, and MAC.

The client and server use the SSL handshake protocol to establish an SSL session between the two devices. During the handshake process, the client and server negotiate the SSL parameters that they will use during the secure session. These sessions are fully proxied throughout the life of the connection, as every packet must be processed by the full TCP and SSL stacks.

- Maximum number of SSL connections: 100 K
- Maximum number of SSL TPS: 1000 with default license.
- Maximum amount of SSL bandwidth: 2Gb/s

9.14.3.2 SSL PKI

SSL features on the ACE support both certificates and public key infrastructure (PKI). Digital certificates contain the following identification attributes: name of the CA, CA digital signature, serial number, name of the server, the subject's public key, and expiration date. The key pair refers to a public key and its corresponding private (secret) key. During the handshake, the RSA key pairs are used to encrypt the session key that both devices will use to encrypt the data that follows the handshake.

The ACE supports the creation of the public/private keys using the `crypto generate key` command. While the ACE does not support certificate generation, it does support configuring Certificate Signing Request (CSR) parameters as shown:

```
WANOPT-ACE1/NSITE-WAAS(config)# crypto csr-params CISCO-WANOPT
WANOPT-ACE1/NSITE-WAAS(config-csr-params)# ?
```

Configure CSR parameters:

```
common-name        Configure organization's common name
country            Configure country name
email              Configure email address
locality           Configure locality name
organization-name  Configure organization name
organization-unit  Configure organization unit's name
serial-number      Configure serial number
state              Configure state name
```

The CSR generated using the `crypto generate csr` command is sent to another system for processing into a certificate. The ACE can import PEM, DER, and PKCS12 certificates and the ACE supports 512, 1024, and 1536 length keys. In a HA environment, the same keys and certificates should be configured on both ACE modules.

```
WANOPT-ACE1/NSITE-WAAS#   show crypto certificate all
All Certificate Files Loaded:
ciscocert:
Subject: /C=US/ST=North
Carolina/L=RTP/O=Cisco/OU=NSITE/CN=www.cisco.com/emailAddress=wanopt@cisco.com
Issuer:
/C=RM/ST=RMSTATE/L=RMCITY/O=RMCOMPANY/OU=RMUNIT/CN=RMNAME/emailAddress=RM@RM.COM
Not Before: Sep 14 19:46:04 2007 GMT
Not After: Sep 13 19:46:04 2008 GMT
CA Cert: FALSE
```

9.14.3.3 SSL Configuration

Before configuring your ACE for SSL operation, it must be configured for server load balancing (SLB). After configuring SLB, the SSL proxy server service is added to the existing SLB policy maps and class maps. The SSL proxy server service defines the handshake parameters that the ACE.

```
serverfarm host HTTPS-FARM
  probe CHECK
  rserver APACHE1 80
    inservice
  rserver APACHE3 80
    inservice
ssl-proxy service SSLPROXY
  key ciscokey
  cert ciscocert
policy-map multi-match L4_LB_WAAS
  class L4_HTTPS_VIP_ADDRESS
    loadbalance vip inservice
    loadbalance policy WAAS_HTTPS_POLICY
    loadbalance vip icmp-reply
    ssl-proxy server SSLPROXY
```

9.15 Network Performance Management

This section describes how the network performance management tools introduced in Chapter 8 can be deployed and utilized to support the WAN and application optimization solution. After a brief overview of the needs for performance monitoring in this environment, each network performance management tool is presented.

Section 9.16, “Performance Monitoring for WAN and Application Optimization” presents the NetQoS performance monitoring products and 9.19, “Cisco NAM Use Cases for WAN and Application Optimization,” presents the Cisco Network Analysis Module (NAM). Each section contains an overview of the product support features relevant to each phase of the WAN and application optimization solution, deployment considerations, and use cases. The use cases illustrate how enterprise users can use these tools to support the WAN and application optimization solution. Section 9.25.5, “Deployment Caveats” presents common deployment caveats for both products.

9.16 Performance Monitoring for WAN and Application Optimization

Performance monitoring tools should be used in all phases of a WAN and application optimization initiative to help ensure a successful deployment.

- **During predeployment testing and baselining**, network performance monitoring validates that the selected WAN and application optimization strategies will have the intended impact on the business critical applications and protocols. This enables proper prioritization of sites for WAN and application optimization based on usage and performance trends.
- **Durin**g **deployment**, monitoring validates the effectiveness of the WAN and application optimization strategy for the selected sites and applications.
- **After deployment**, performance monitoring ensures the ongoing effectiveness of the WAN and application optimization strategy as network, datacenter, and usage conditions change and by maintaining the visibility necessary to enable efficient troubleshooting. Performance monitoring also helps in identifying new opportunities for WAN and application optimization.

Section 9.16.1 describes how to use NetQoS to support deployed WAN and application initiatives. Section 9.19 describes how to use Cisco NAM 3.6 to support such initiatives; in this case, the focus is on troubleshooting and conversation specific WAN and application optimization validation.

9.16.1 NetQoS Support for WAN and Application Optimization

NetQoS performance monitoring supports all phases of WAN and application optimization: predeployment, deployment, and postdeployment. The following sections describe each phase.

9.16.1.1 Predeployment Support

When developing a WAN and application optimization strategy, a first step is to profile traffic patterns and resource bottlenecks and establish a baseline of the performance of links, servers, and applications. Profiling involves identifying the applications running on the network, and understanding the consumption of WAN resources by different types of business and non-business traffic.

NetQoS NetVoyant leverages NBAR to produce reports such as the one shown in Figure 9-42. Using NBAR statistics, IT staff can identify the applications running on the network.

Figure 9-42. NBAR Statistics by Protocol

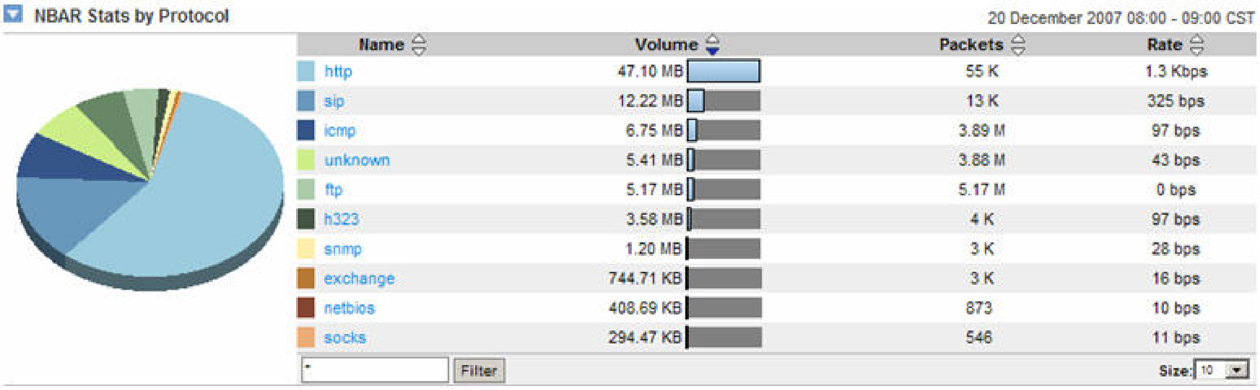

NBAR Stats by Protocol

20 December 2007 08:00 - 09:00 CST

	Name	Volume	Packets	Rate
■	http	47.10 MB	55 K	1.3 Kbps
■	sip	12.22 MB	13 K	325 bps
■	icmp	6.75 MB	3.89 M	97 bps
■	unknown	5.41 MB	3.88 M	43 bps
■	ftp	5.17 MB	5.17 M	0 bps
■	h323	3.58 MB	4 K	97 bps
■	snmp	1.20 MB	3 K	28 bps
■	exchange	744.71 KB	3 K	16 bps
■	netbios	408.69 KB	873	10 bps
■	socks	294.47 KB	546	11 bps

NetQoS ReporterAnalyzer also uses NetFlow information to show traffic profiles of WAN resources. For example, Figure 9-43 is a report from NetQoS ReporterAnalyzer that shows the traffic profile for a hypothetical New York branch site. The report identifies applications on the WAN link and shows that most of the traffic is voice traffic (83.47%).

Figure 9-43. Protocol Summary Report for a Branch WAN Link

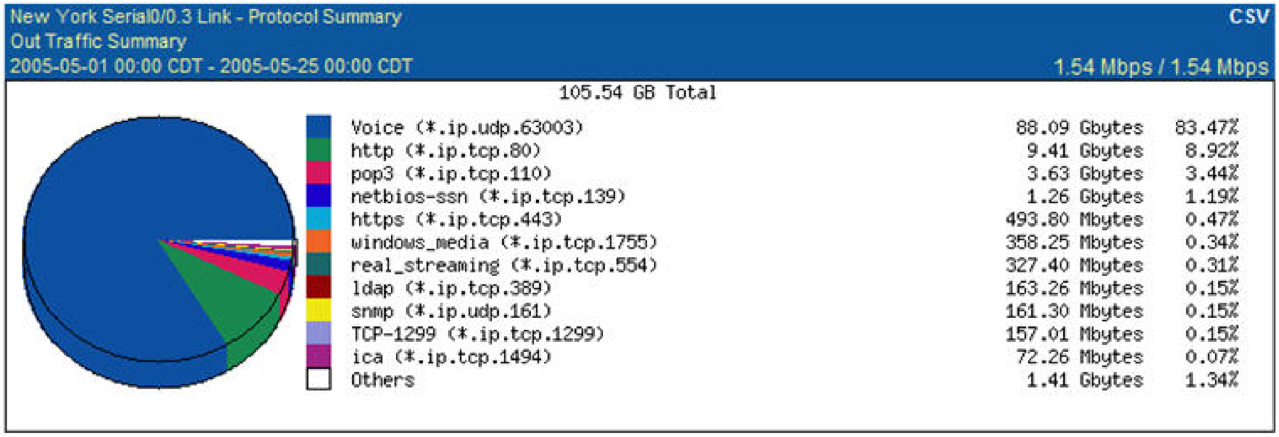

New York Serial0/0.3 Link - Protocol Summary
Out Traffic Summary
2005-05-01 00:00 CDT - 2005-05-25 00:00 CDT

1.54 Mbps / 1.54 Mbps

105.54 GB Total

■	Voice (*.ip.udp.63003)	88.09 Gbytes	83.47%
■	http (*.ip.tcp.80)	9.41 Gbytes	8.92%
■	pop3 (*.ip.tcp.110)	3.63 Gbytes	3.44%
■	netbios-ssn (*.ip.tcp.139)	1.26 Gbytes	1.19%
■	https (*.ip.tcp.443)	493.80 Mbytes	0.47%
■	windows_media (*.ip.tcp.1755)	358.25 Mbytes	0.34%
■	real_streaming (*.ip.tcp.554)	327.40 Mbytes	0.31%
■	ldap (*.ip.tcp.389)	163.26 Mbytes	0.15%
■	snmp (*.ip.udp.161)	161.30 Mbytes	0.15%
■	TCP-1299 (*.ip.tcp.1299)	157.01 Mbytes	0.15%
■	ica (*.ip.tcp.1494)	72.26 Mbytes	0.07%
□	Others	1.41 Gbytes	1.34%

Figure 9-43 and Figure 9-44 show more ReporterAnalyzer custom reports generated from the NetFlow instrumentation provided on Cisco routers. These reports list the WAN links with the most time over a user-selected threshold, and traffic composition for two of the links.

A time-over-threshold report (shown in Figure 9-43) can indicate which links on the network might be good candidates for optimization. The protocol reports (Figure 9-43 and Figure 9-44) help to predict the likely outcome of different optimization strategies. For example, the report for a New York link (shown previously in Figure 9-43) shows high volumes of UDP (VoIP) traffic, suggesting that attention to routing and QoS policies might be a useful first step in optimizing this link. Conversely, the heavily-utilized Houston link (Figure 9-44) shows a mix of TCP/IP application traffic that would probably benefit most from WAN optimization appliances, such as Cisco WAAS.

Figure 9-44. ReporterAnalyzer Custom Report Showing Networks Having the Most Time over a Selected Threshold

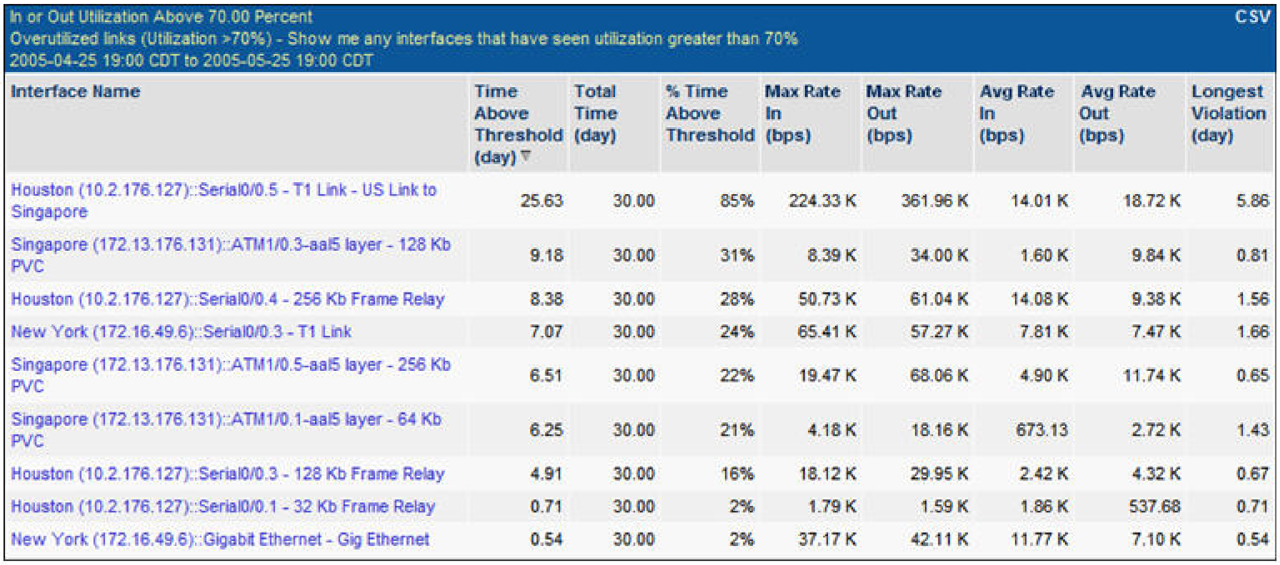

In or Out Utilization Above 70.00 Percent
Overutilized links (Utilization >70%) - Show me any interfaces that have seen utilization greater than 70%
2005-04-25 19:00 CDT to 2005-05-25 19:00 CDT

CSV

Interface Name	Time Above Threshold (day) ▿	Total Time (day)	% Time Above Threshold	Max Rate In (bps)	Max Rate Out (bps)	Avg Rate In (bps)	Avg Rate Out (bps)	Longest Violation (day)
Houston (10.2.176.127)::Serial0/0.5 - T1 Link - US Link to Singapore	25.63	30.00	85%	224.33 K	361.96 K	14.01 K	18.72 K	5.86
Singapore (172.13.176.131)::ATM1/0.3-aal5 layer - 128 Kb PVC	9.18	30.00	31%	8.39 K	34.00 K	1.60 K	9.84 K	0.81
Houston (10.2.176.127)::Serial0/0.4 - 256 Kb Frame Relay	8.38	30.00	28%	50.73 K	61.04 K	14.08 K	9.38 K	1.56
New York (172.16.49.6)::Serial0/0.3 - T1 Link	7.07	30.00	24%	65.41 K	57.27 K	7.81 K	7.47 K	1.66
Singapore (172.13.176.131)::ATM1/0.5-aal5 layer - 256 Kb PVC	6.51	30.00	22%	19.47 K	68.06 K	4.90 K	11.74 K	0.65
Singapore (172.13.176.131)::ATM1/0.1-aal5 layer - 64 Kb PVC	6.25	30.00	21%	4.18 K	18.16 K	673.13	2.72 K	1.43
Houston (10.2.176.127)::Serial0/0.3 - 128 Kb Frame Relay	4.91	30.00	16%	18.12 K	29.95 K	2.42 K	4.32 K	0.67
Houston (10.2.176.127)::Serial0/0.1 - 32 Kb Frame Relay	0.71	30.00	2%	1.79 K	1.59 K	1.86 K	537.68	0.71
New York (172.16.49.6)::Gigabit Ethernet - Gig Ethernet	0.54	30.00	2%	37.17 K	42.11 K	11.77 K	7.10 K	0.54

Figure 9-45. Protocol Summary Report for another Branch WAN Link

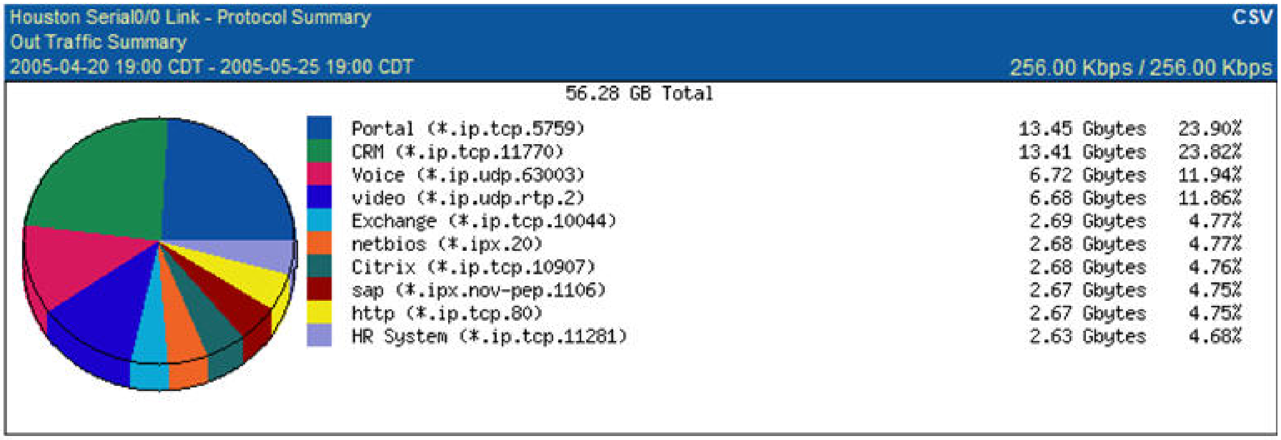

Houston Serial0/0 Link - Protocol Summary
Out Traffic Summary
2005-04-20 19:00 CDT - 2005-05-25 19:00 CDT

CSV

256.00 Kbps / 256.00 Kbps

56.28 GB Total

	Protocol	Volume	Share
🟦	Portal (*.ip.tcp.5759)	13.45 Gbytes	23.90%
🟩	CRM (*.ip.tcp.11770)	13.41 Gbytes	23.82%
🟥	Voice (*.ip.udp.63003)	6.72 Gbytes	11.94%
🟦	video (*.ip.udp.rtp.2)	6.68 Gbytes	11.86%
🟦	Exchange (*.ip.tcp.10044)	2.69 Gbytes	4.77%
🟧	netbios (*.ipx.20)	2.68 Gbytes	4.77%
🟩	Citrix (*.ip.tcp.10907)	2.68 Gbytes	4.76%
🟥	sap (*.ipx.nov-pep.1106)	2.67 Gbytes	4.75%
🟨	http (*.ip.tcp.80)	2.67 Gbytes	4.75%
🟪	HR System (*.ip.tcp.11281)	2.63 Gbytes	4.68%

Baselining involves determining typical application and network behavior based on traffic loads on the WAN resources and service levels experienced by end users. Understanding the user experience, and the contribution of the network to the delays experienced at each site, enables IT staff to prioritize the applications and links to be optimized. Knowing the composition of traffic consuming the links can help to predict the impact of WAN and application optimization on network utilization and future capacity.

IT staff must establish baselines with respect to the traffic loads on WAN resources and the composition of traffic. NetQoS ReporterAnalyzer leverages NetFlow data to help compute baseline and trends. IT staff also needs to establish baselines for network, server, and application performance. That can be achieved in two ways: using active or passive measurements. Active measurements are based on synthetic traffic, i.e. network traffic generated strictly for the purpose of measuring a network/server/application characteristic, while passive measurements are based on actual end-user traffic.

Each measurement approach has advantages and disadvantages. The best practice is to use both active and passive measurements. Passive measurement is the most accurate approach for the end-user application traffic. However, passive measurement is limited to measuring the performance of existing traffic types (which may not be present on the network at all times) and existing traffic patterns (which may not reflect patterns for new and future applications).

Active measurements have the advantage of being more controllable; performance can be measured between any two points in the network, the type, frequency and traffic class can be specified for the generated traffic. On the other hand, active measurements are only an approximation for the performance of actual traffic. An additional disadvantage is that in order to perform active measurements, some traffic is injected in the network.

To establish baselines using active measurements, NetVoyant uses Cisco IP SLA to report performance metrics, and uses IP SLA synthetic traffic to compute SLA baselines. Figure 9-45 shows an example of a VoIP performance report where the current performance (98.7ms) is compared to the baseline (108.8ms).

Figure 9-46. VoIP Performance Report Example

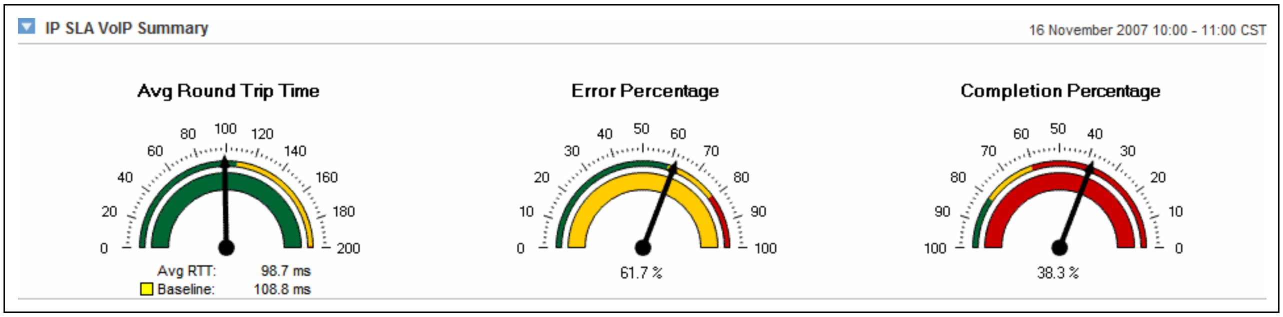

IP SLA VoIP Summary

16 November 2007 10:00 - 11:00 CST

Avg Round Trip Time

Avg RTT: 98.7 ms

🟨 Baseline: 108.8 ms

Error Percentage

61.7 %

Completion Percentage

38.3 %

To establish baselines using passive measurements, NetQoS SuperAgent provides reports based on actual network traffic. SuperAgent product calculates baseline information for applications, networks, and servers each hour to indicate when performance conditions are normal for that hour of the day, factoring in the previous week, day of the week, and day of the month.

Figure 9-47 shows application response time performance maps from NetQoS SuperAgent. These reports list remote sites sorted by worst overall performance (as experienced by end users) for a selected application, and the contribution of network latency to the delay at each site. SuperAgent performance maps can serve as guides for prioritizing sites and applications for WAN and application optimization. Users in the Durham office experience the longest transaction time and the longest network latency, so the Durham site might be considered a top candidate for WAN and application optimization.

Figure 9-47. SuperAgent Performance Maps for a Selected Application

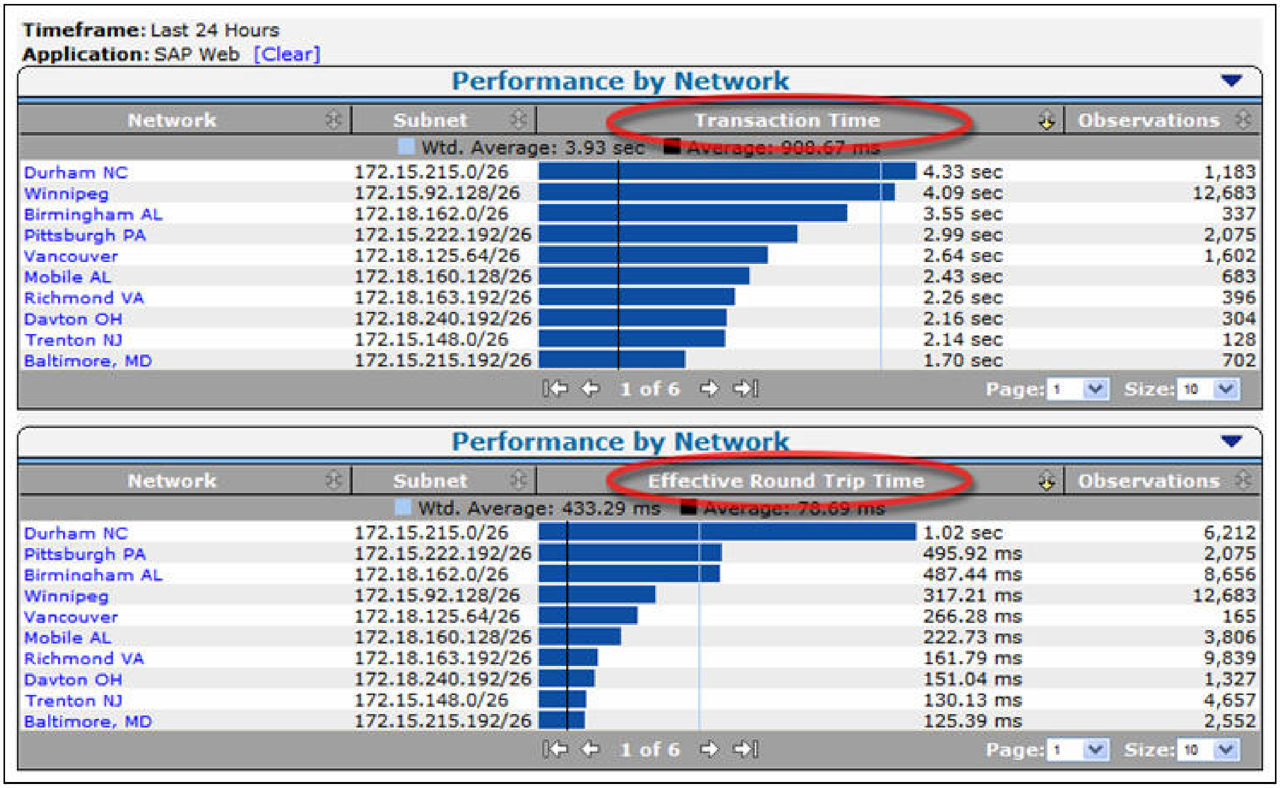

Timeframe: Last 24 Hours
Application: SAP Web [Clear]

Performance by Network

Wtd. Average: 3.93 sec ■ Average: 908.67 ms

Network	Subnet	Transaction Time	Observations
Durham NC	172.15.215.0/26	4.33 sec	1,183
Winnipeg	172.15.92.128/26	4.09 sec	12,683
Birmingham AL	172.18.162.0/26	3.55 sec	337
Pittsburgh PA	172.15.222.192/26	2.99 sec	2,075
Vancouver	172.18.125.64/26	2.64 sec	1,602
Mobile AL	172.18.160.128/26	2.43 sec	683
Richmond VA	172.18.163.192/26	2.26 sec	396
Davton OH	172.18.240.192/26	2.16 sec	304
Trenton NJ	172.15.148.0/26	2.14 sec	128
Baltimore, MD	172.15.215.192/26	1.70 sec	702

1 of 6 Page: 1 Size: 10

Performance by Network

Wtd. Average: 433.29 ms ■ Average: 78.69 ms

Network	Subnet	Effective Round Trip Time	Observations
Durham NC	172.15.215.0/26	1.02 sec	6,212
Pittsburgh PA	172.15.222.192/26	495.92 ms	2,075
Birmingham AL	172.18.162.0/26	487.44 ms	8,656
Winnipeg	172.15.92.128/26	317.21 ms	12,683
Vancouver	172.18.125.64/26	266.28 ms	165
Mobile AL	172.18.160.128/26	222.73 ms	3,806
Richmond VA	172.18.163.192/26	161.79 ms	9,839
Davton OH	172.18.240.192/26	151.04 ms	1,327
Trenton NJ	172.15.148.0/26	130.13 ms	4,657
Baltimore, MD	172.15.215.192/26	125.39 ms	2,552

1 of 6 Page: 1 Size: 10

Passive response time measurements help pinpoint remote sites with the biggest performance issues, but the measurements must be used with measurements of utilization and traffic composition for individual links, and for the entire infrastructure. (Examples were shown previously in Figure 9-43, Figure 9-44, and Figure 9-45.

Protocols and transactions respond differently to optimization technologies. Therefore, the traffic mix must be well understood to determine an effective optimization strategy. As mentioned previously, TCP traffic is more likely to benefit from caching and compression technologies, while voice traffic is more likely to benefit from changes to routing and QoS policies.

Equipped with measurements of end-user experience and network latency for business applications across the network, and with measurements of utilization and protocol distribution on key links, IT personnel can undertake wider WAN and application optimization deployments with greater confidence. An example of a predeployment support use case is presented in 9.17, "Use Case 1: Predeployment Baselining."

9.16.1.2 Support during Deployment

Network performance monitoring products can help IT personnel measure the effectiveness of WAN and application optimization deployments. Users can see:

- Detailed views of how the bandwidth consumption of target applications will change after optimization.
- Changes in application response times as experienced by end users at remote locations.
- For any optimized application, detailed maps showing how WAN and application optimization affects the volume of data transmitted by datacenter components, on the WAN, and at remote sites.

- How server offload made possible by WAN and application optimization can change server response times and data volumes.

This section focuses mainly on validating Cisco WAAS, but some graphs and reports apply to validating other WAN and application optimization initiatives, such as changing QoS policies or deploying PfR.

Figure 9-48 shows how Cisco WAAS reduced the bandwidth usage of NetBIOS from a custom application on a WAN link. NetQoS ReporterAnalyzer shows the data rate before and after WAAS deployment. It also shows how the data rate relates to a baseline of data for the same period collected before deployment– the baseline is shown in light gray. Similar reports can be generated to show in detail how WAN and application optimization lowers individual applications' bandwidth consumption and overall link utilization.

Figure 9-48. ReporterAnalyzer Displaying a Predeployment Baseline

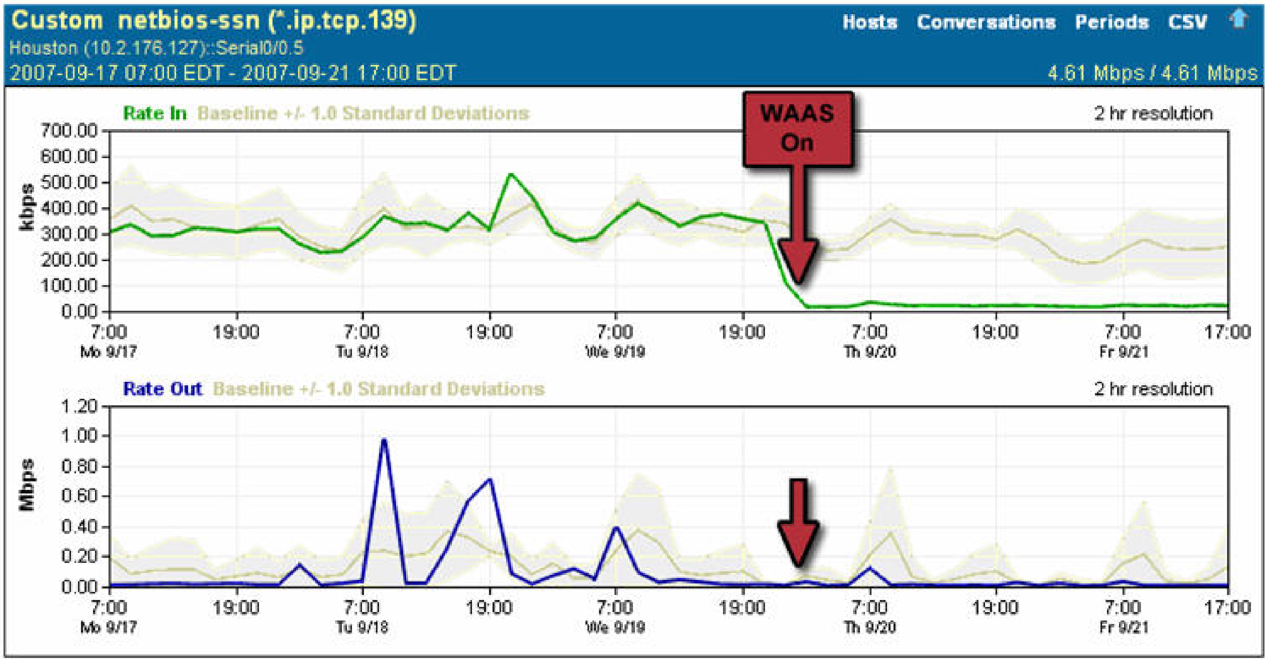

Figure 9-49 shows the impact of Cisco WAAS on application response times experienced by users of an interactive web application at a particular site. Note the pattern of daily transaction delays before WAAS deployment. Users are likely to be dissatisfied and unproductive, since transaction delays are highly variable and last up to 3 seconds.

While the observed number of transactions (represented by the light gray "Observations" line) shows a consistent daily pattern, the average response time for users of the application at this location drops from 562 milliseconds before WAAS deployment to 205 milliseconds after, as indicated by SuperAgent response-time reports (not shown here). Furthermore, the WAAS deployment resulted in far more consistent application delivery to this remote site – with no transactions approaching the three-second delays seen before deployment – and is therefore likely to make a significant, positive impact on user satisfaction and productivity.

Figure 9-49. SuperAgent Reporting that WAAS Improves Application Performance

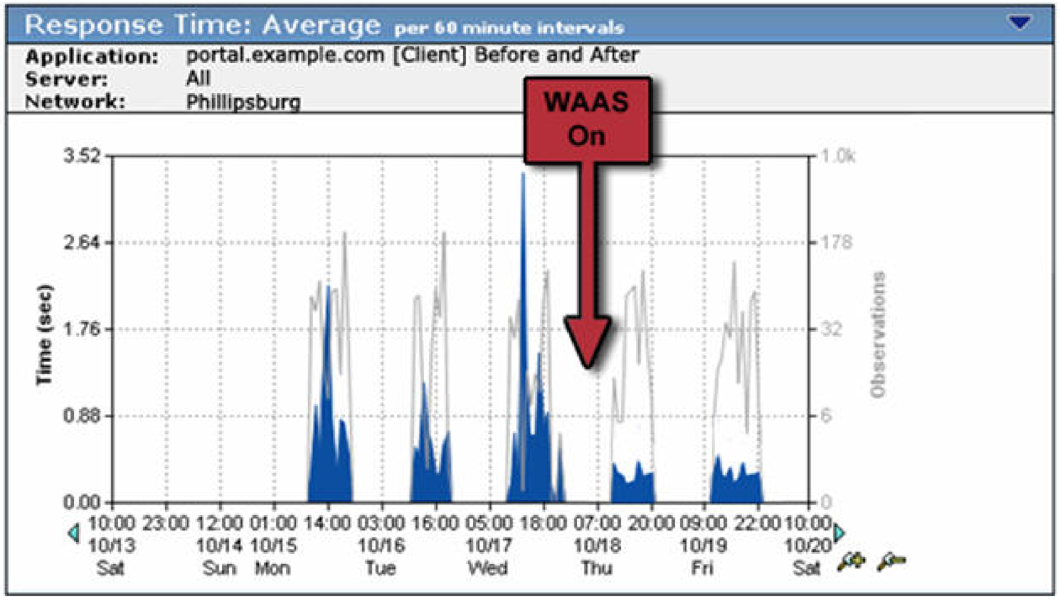

Figure 9-50 shows how network latency for the application dropped significantly after WAAS deployment because of reduced congestion on the optimized link.

Figure 9-50. SuperAgent Reporting Reduced WAN Segment Latency after WAAS Optimization

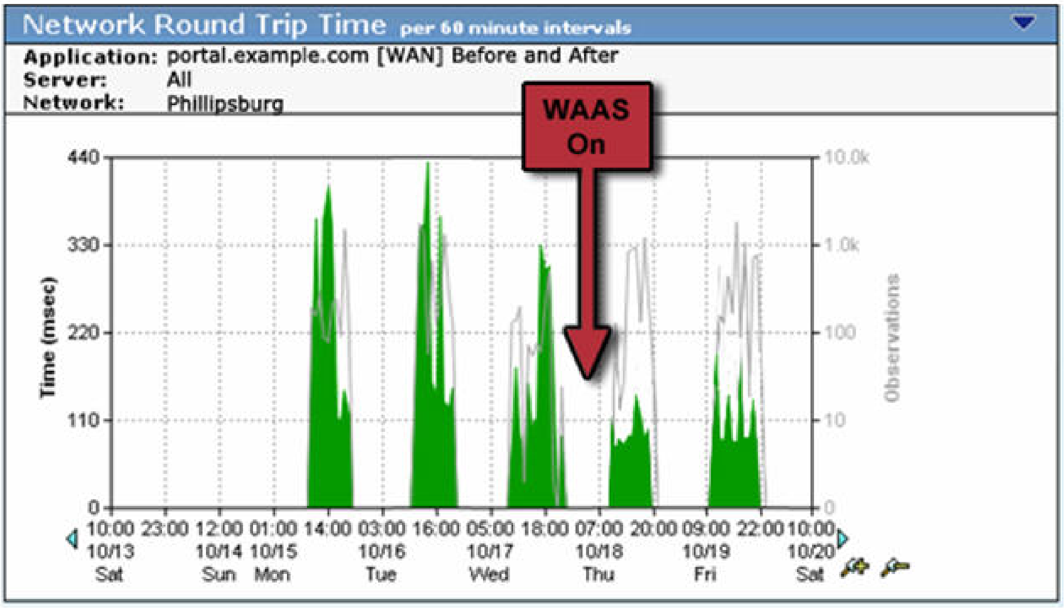

Application delays caused by network retransmission, shown in Figure 9-50, were also significantly reduced by the WAAS deployment.

Figure 9-51. SuperAgent Reporting Decreased Network Retransmission Delay after WAAS Optimization

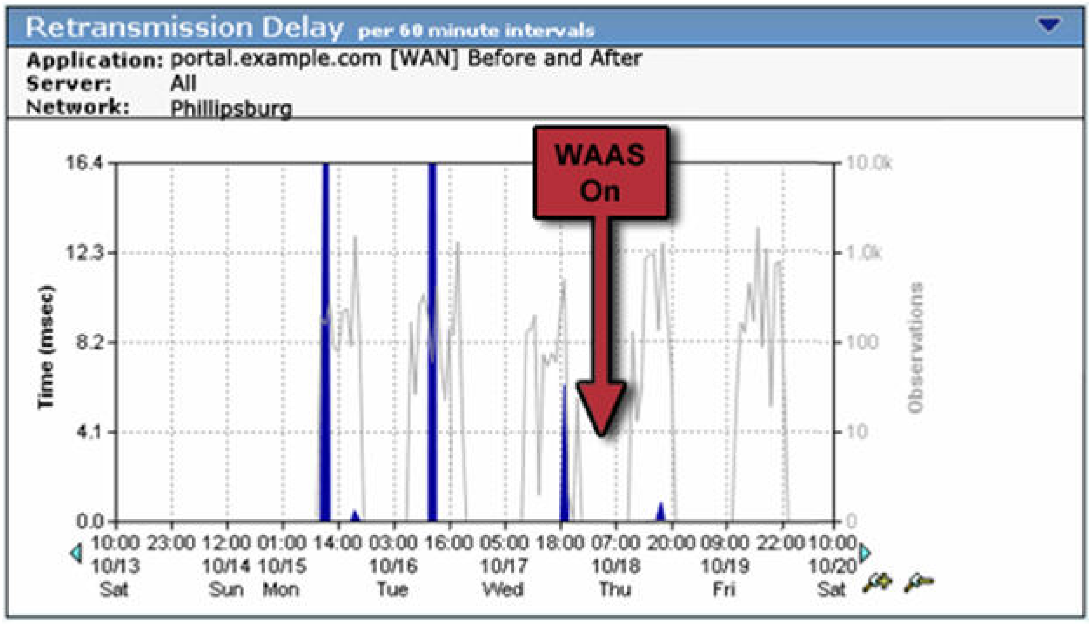

Retransmission Delay per 60 minute intervals

Application: portal.example.com [WAN] Before and After
Server: All
Network: Phillipsburg

The effect of server offload made possible by Cisco WAAS can be observed in several ways. Figure 9-52 shows improved response time for a server hosting this CIFS (NetBIOS) intensive application after WAAS deployment.

Figure 9-52. SuperAgent Reporting Faster, More Consistent Server Response Times after Server Offload

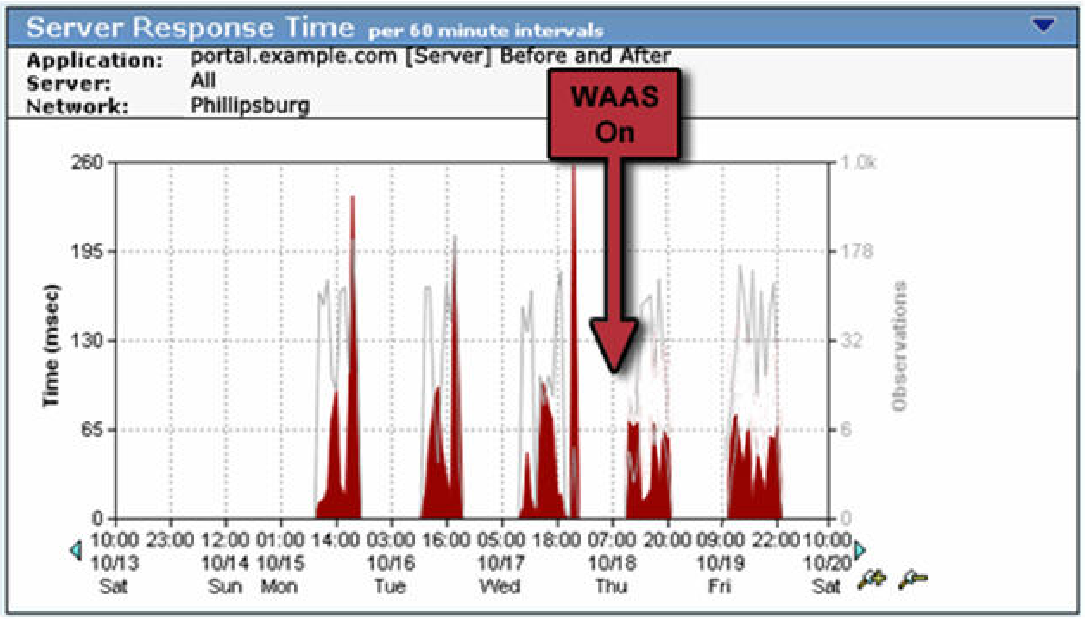

Server Response Time per 60 minute intervals

Application: portal.example.com [Server] Before and After
Server: All
Network: Phillipsburg

Figure 9-53 shows data volume reduction for a bandwidth-hungry email application. This SuperAgent performance map, configured to report email traffic over a single optimized link, shows that advanced data reduction and compression features of Cisco WAAS have reduced 180 MB of server data volume to only 99 MB of volume over the WAN.

Figure 9-53. SuperAgent Performance Map Showing Reduced WAN Data Volumes after WAAS Optimization

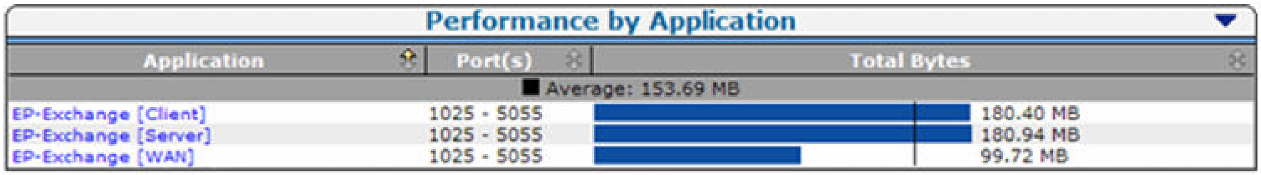

Performance by Application

Application	Port(s)	Total Bytes
	■ Average: 153.69 MB	
EP-Exchange [Client]	1025 - 5055	180.40 MB
EP-Exchange [Server]	1025 - 5055	180.94 MB
EP-Exchange [WAN]	1025 - 5055	99.72 MB

Detailed documentation of the performance improvements experienced by end-users after WAN and application optimization deployment – in combination with performance metrics that show how data volume and network latency improvements can lead to better overall performance – help to prove the effectiveness of each WAN and application optimization initiative and can help to justify and fine-tune future efforts. An example of a during deployment support use case is presented in section 9.18.

9.16.1.3 Postdeployment Support

NetQoS performance monitoring products, in combination with Cisco WAN and application optimization technologies, help IT professionals maintain the visibility necessary to ensure efficient application delivery regardless of how network, datacenter, and usage conditions change. Effective network performance monitoring also makes possible more efficient troubleshooting and capacity planning to meet the organization's future needs.

The NetQoS Performance Center can alert IT staff to performance problems no matter where they may occur in an overall transaction. The following example illustrates some of the performance reports used to support troubleshooting in the example network shown in Figure 9-54. The figure shows the following appliances: NetQoS Super Agent Collector (SA-CO), Super Agent Aggregator (SA-AG), Super Agent Management Console (SA-MC), NetVoyant (NQ-NV), and Reporter Analyzer (NQ-RA). In this example, the NetQoS appliances are connected to the distribution layer switch (DS1).

Figure 9-54. Post-Deployment Support Network Example

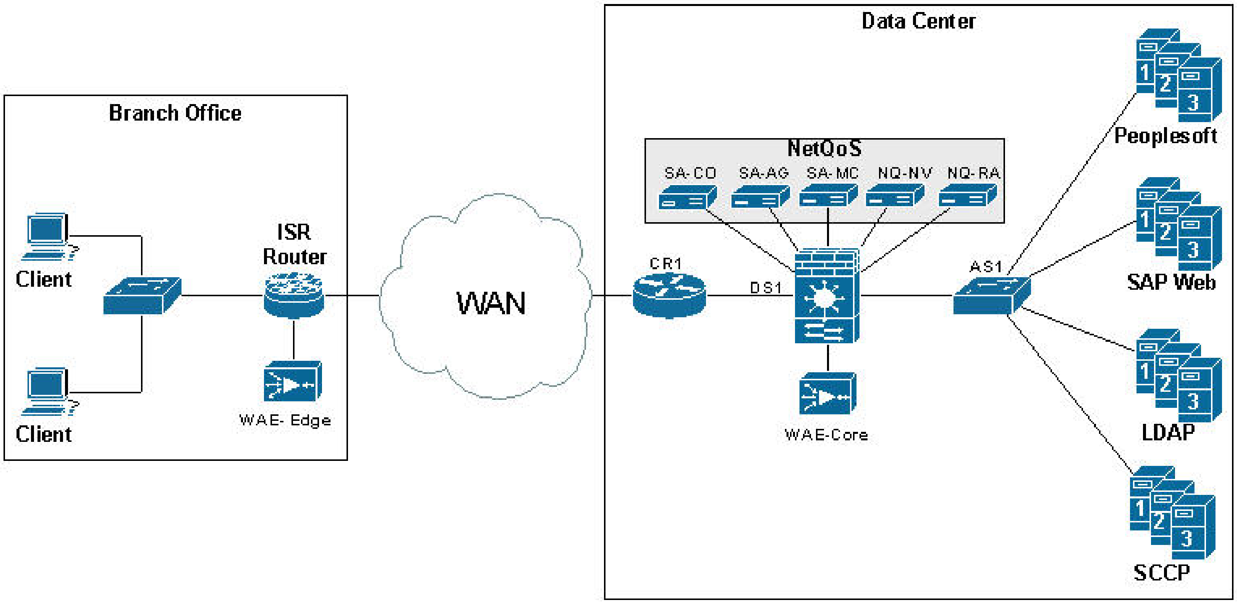

Figure 9-55 is a performance view, sorted by worst performing application, which highlights an application tier responsible for a recent slowdown.

Figure 9-55. NetQoS Performance Center Report: Performance by Application

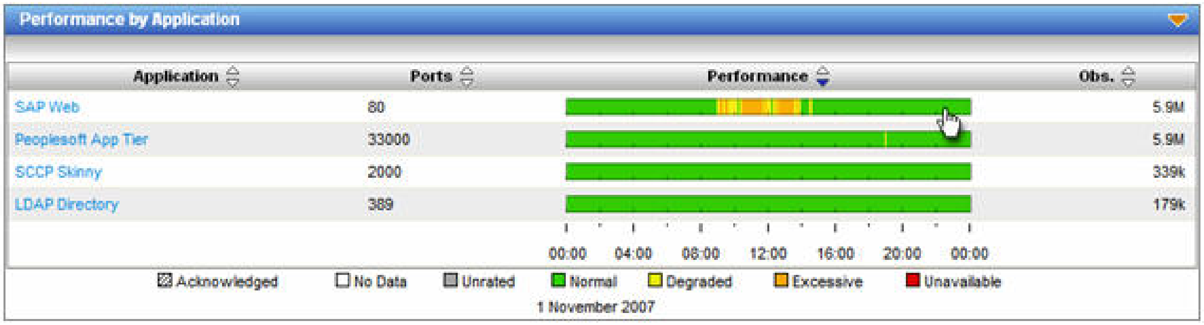

Performance by Application

Application	Ports	Performance	Obs.
SAP Web	80		5.9M
Peoplesoft App Tier	33000		5.9M
SCCP Skinny	2000		339k
LDAP Directory	389		179k

☒ Acknowledged □ No Data ▩ Unrated 🟩 Normal 🟨 Degraded 🟧 Excessive 🟥 Unavailable

1 November 2007

Users can navigate to an Engineering view, shown in Figure 9-56, to see response-time metrics for a server in the data center that appears to be responsible for the slowdown. In this example, a response time composition graph reveals an increase in SRT (shown in red) without corresponding delays in application data transfer time (shown in yellow). This condition is associated with server resource issues caused by background processes, memory leaks, faulty configuration changes, and so on.

Figure 9-56. A SuperAgent Engineering View

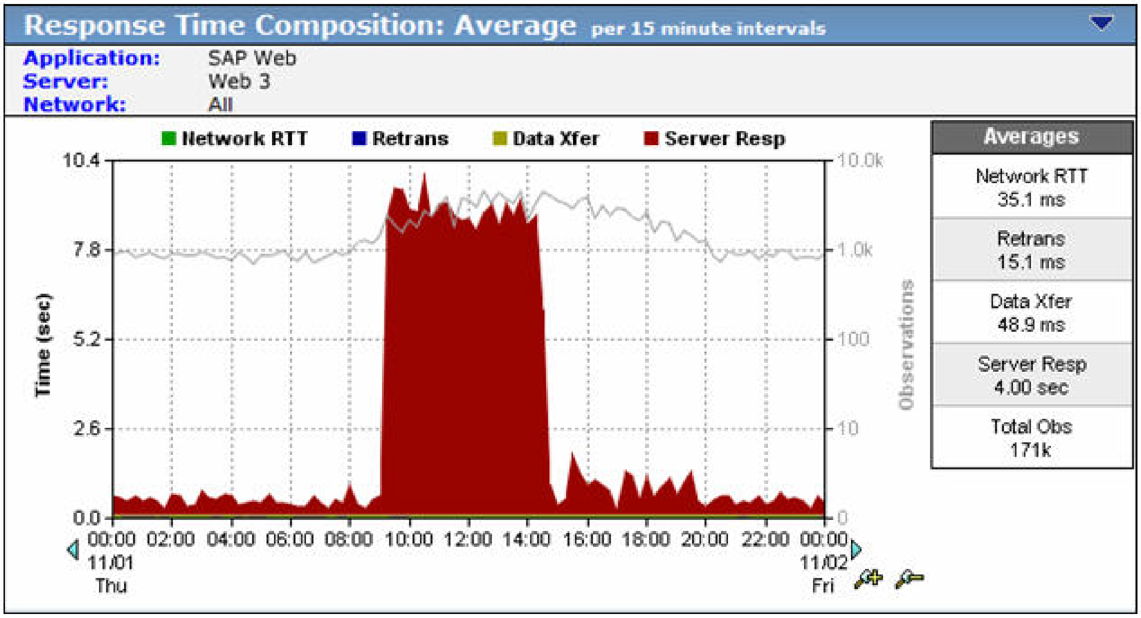

Response Time Composition: Average per 15 minute intervals

Application: SAP Web
Server: Web 3
Network: All

Averages
Network RTT 35.1 ms
Retrans 15.1 ms
Data Xfer 48.9 ms
Server Resp 4.00 sec
Total Obs 171k

SuperAgent can be configured to automatically investigate server and network performance issues. In this case, SuperAgent reports the presence of a backup process consuming significant CPU during the slowdown.

Figure 9-57. A NetVoyant Device Performance View

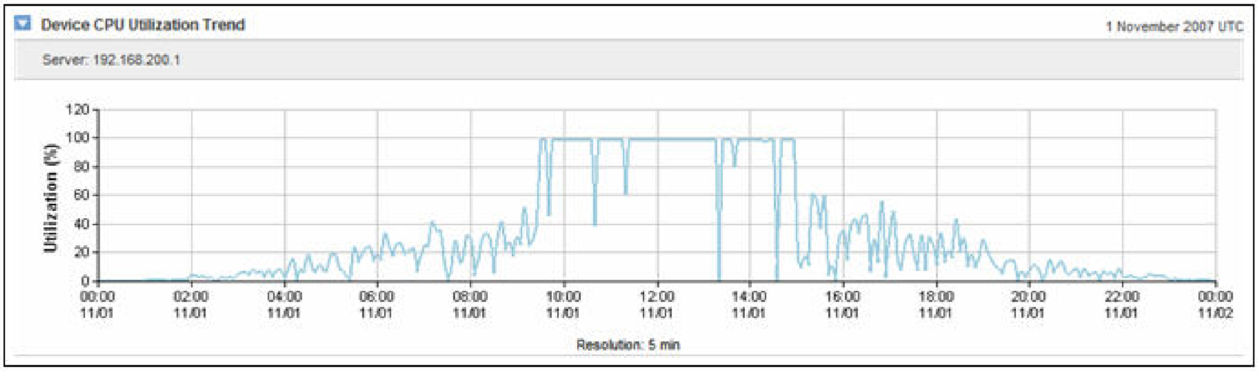

Device CPU Utilization Trend

1 November 2007 UTC

Server: 192.168.200.1

NetVoyant generates a detailed CPU utilization report, shown in Figure 9-57, and server process list, shown in Figure 9-58, and presents them in the NetQoS Performance Center. These document that a backup process, mistakenly configured to run during business hours, is responsible for the slowdown.

In the example, a combination of application response time monitoring, automated SuperAgent investigations, and SNMP device statistics help IT staff quickly and easily find the root cause of the performance problem. The data made it possible to quickly correct the backup schedule and help prevent the problem from recurring.

Figure 9-58. Process List Showing the Presence of a Backup Application

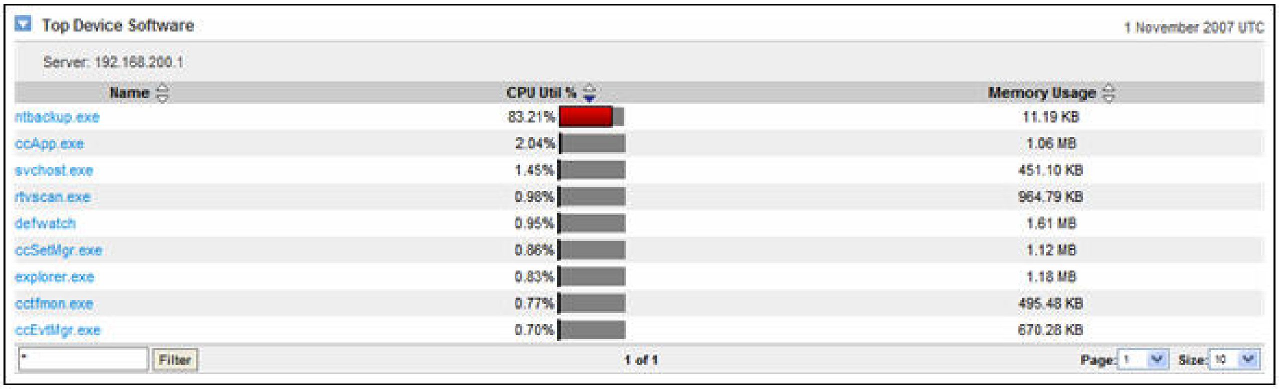

Top Device Software

1 November 2007 UTC

Server: 192.168.200.1

Name	CPU Util %	Memory Usage
ntbackup.exe	83.21%	11.19 KB
ccApp.exe	2.04%	1.06 MB
svchost.exe	1.45%	451.10 KB
rtvscan.exe	0.98%	964.79 KB
defwatch	0.95%	1.61 MB
ccSetMgr.exe	0.86%	1.12 MB
explorer.exe	0.83%	1.18 MB
cctfmon.exe	0.77%	495.48 KB
ccEvtMgr.exe	0.70%	670.28 KB

The presentation of comprehensive performance data in a single interface, including application response times, link traffic analysis, and SNMP device performance statistics, supports faster troubleshooting and better management decisions.

9.16.2 NetQoS Metrics for WAN and Application Optimization

NetQoS SuperAgent analyzes the following metrics when reporting application response times. This analysis reports total transaction time for the network, server, and application components.

Table 9-4. NetQoS Metrics

Metric relative to	Metric	Description
Networks	Network RTT	Time that a packet takes to traverse the network.
	Network connection time	Time it takes the client to confirm the server connection acknowledgment. Delay is probably caused by network latency.
	Effective RTT	Network RTT plus delays due to any retransmissions.
	Retransmission delay	Elapsed time between sending the original packet and sending of the last duplicate packet.Retransmission delay is reported as an average across all observations, not just retransmitted packets.
	Packet loss percentage	The ratio of retransmitted data to total data.
Servers	Server response time (SRT)	Time for the server to begin responding to a request.
	Server connection time	Time the server takes to acknowledge an initial client connection request.
	Refused sessions	The number of requests that the server explicitly rejected during the data collection interval.
	Unresponsive sessions	The number of essions in which a connection request was sent and the server never responded.
Applications	Data transfer time	Time to transmit a complete response, measured from initial to final packet. Excludes initial SRT, and includes network RTT only if there is the data to send does not fit in the TCP window.

Metric relative to	Metric	Description
	Transaction time	Time to complete a TCP transaction or data request within a persistent TCP connection, from the moment a client sends the request to the time that the client receives the last packet in the response.

Figure 9-59. Four Primary Metrics That Sum to Total Transaction Time

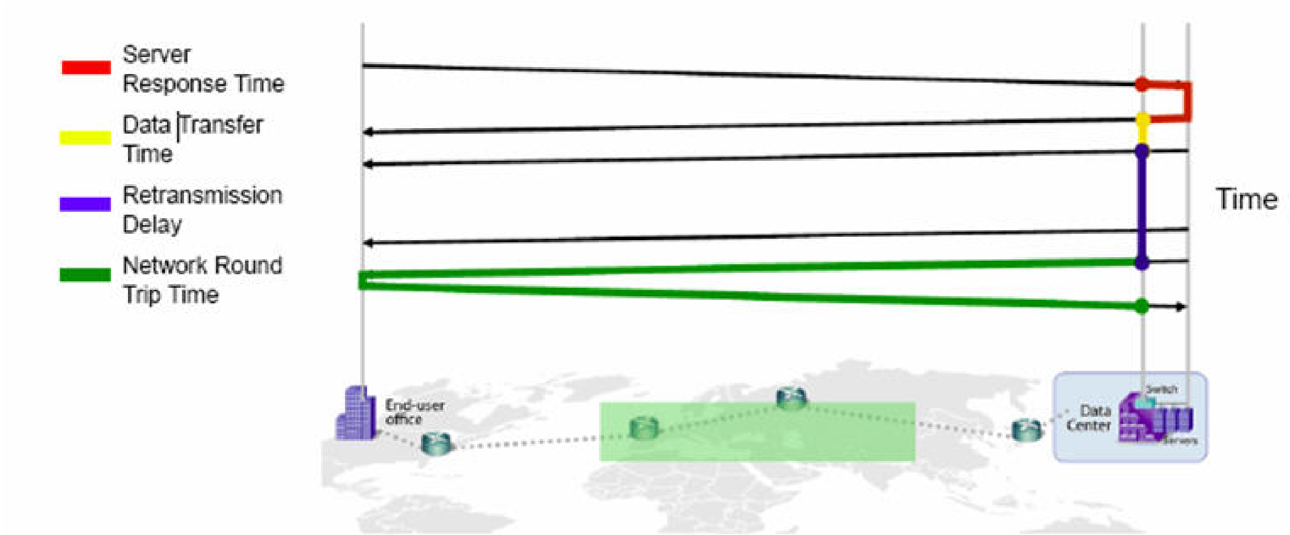

9.16.3 NetQoS Deployment Considerations

Figure 9-60 shows NetQoS deployed in the data center. The NetQoS appliances (shown in a light gray box) are deployed in the distribution layer of the Data Center. The figure shows the following appliances: NetQoS Super Agent Collector (SA-CO), Super Agent Aggregator (SA-AG), Super Agent Management Console (SA-MC), NetVoyant (NQ-NV) and Reporter Analyzer (NQ-RA). In this example, the NetQoS appliances are connected to the distribution layer switch (DS1). For redundancy, they can also be connected to the secondary distribution layer switch, DS2, which is not shown in the figure.

Figure 9-60. NetQoS Placement in the Data Center

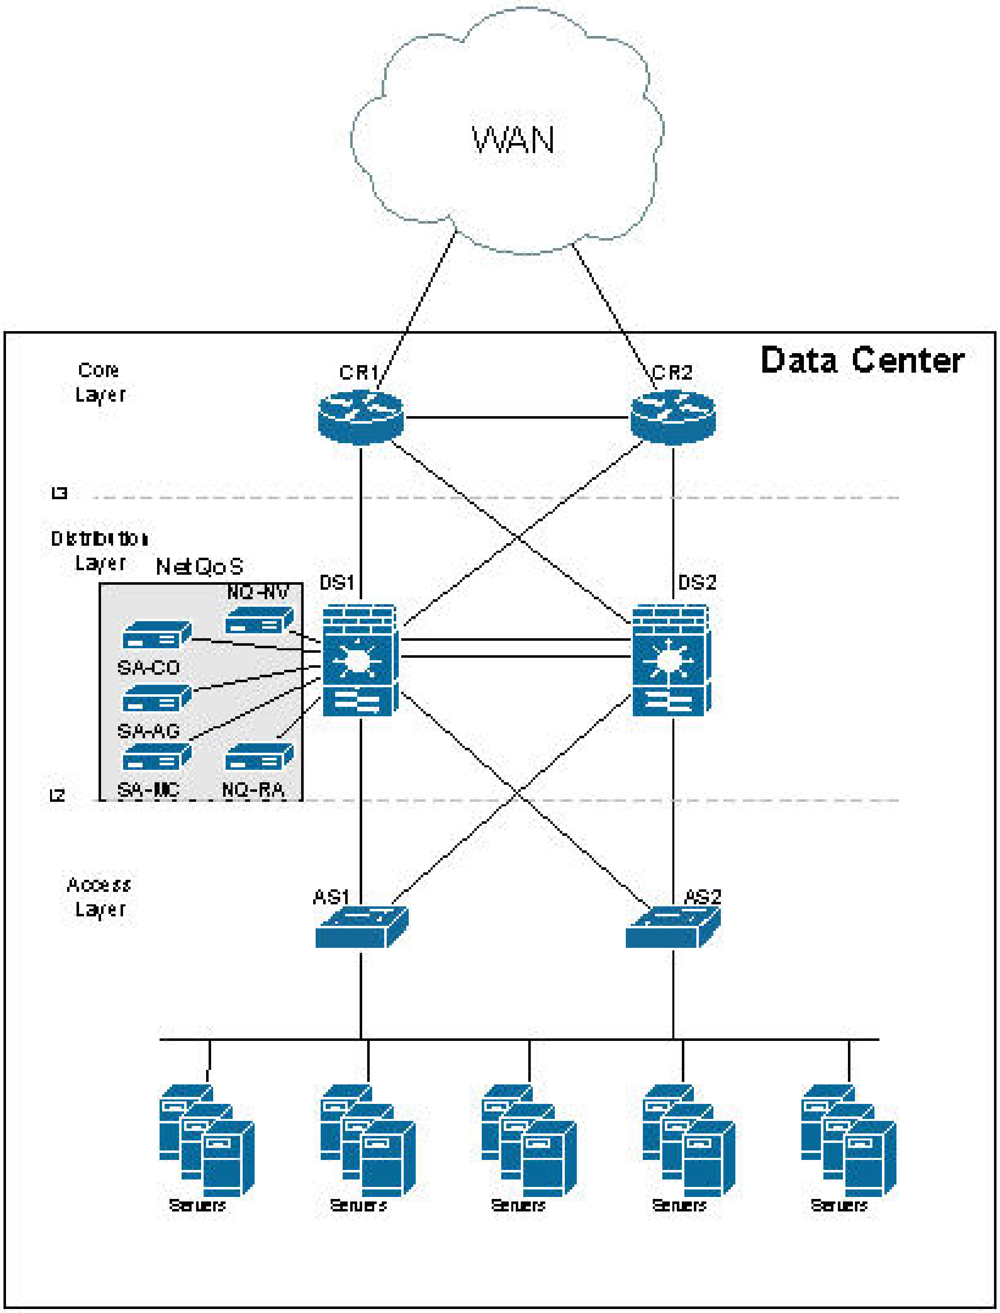

9.16.4 Application Response Time Analysis with NetQoS SuperAgent

A SuperAgent deployment comprises a SuperAgent Management Console and some combination of one or more Aggregators for collecting Flow Agent instrumentation from Cisco Wide Area Application Engine (WAE) devices, and Collectors for monitoring Switched Port Analyzer (SPAN) data from data center switches.

In a standalone configuration, a SuperAgent Management Console and one Collector or Aggregator reside on one server. In a distributed configuration, multiple Collectors and Aggregators are associated with one SuperAgent Management Console, and each component is installed on a separate server. Figure 9-61 illustrates a distributed configuration using a collector (SA-CO) and an aggregator (SA-AG) with a separate server Management Console (SA-MC).

Figure 9-61. SuperAgent Distributed Configuration Example

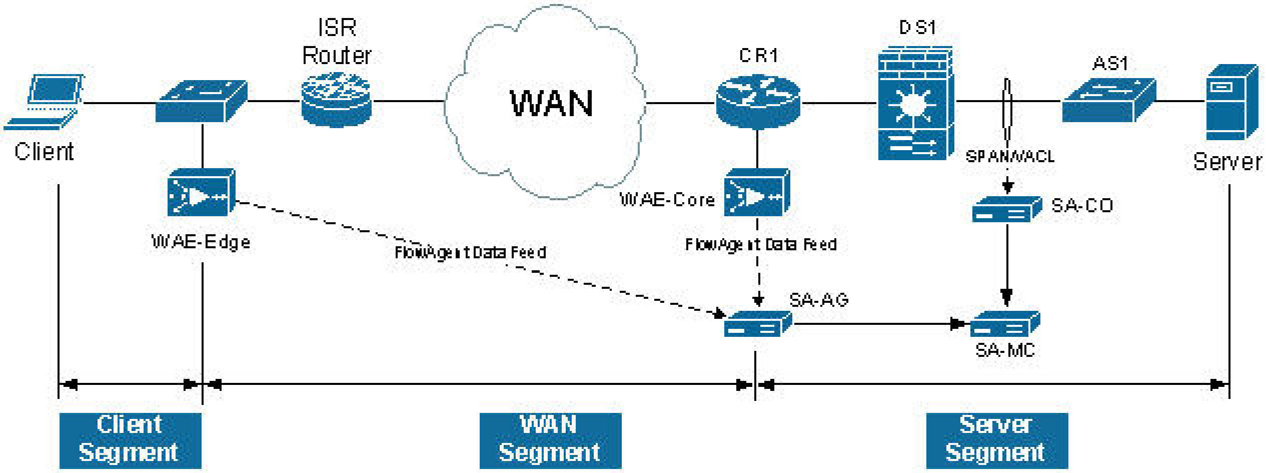

9.16.4.1 SuperAgent Data Feeds

When installing a SuperAgent Collector, IOS commands are executed on a distribution-layer switch (DS1) to configure SPAN settings. Advanced spanning techniques such as Remote SPAN (RSPAN)/ Encapsulated Remote SPAN (ERSPAN) and controls such as VLAN ACLs (VACLs) can be configured to direct only traffic of interest to the SuperAgent Collector. The SuperAgent Collector capacity is limited to 1Gb/s, so only business critical traffic should be monitored.

When installing a SuperAgent Aggregator, FlowAgent export is enabled on all deployed WAE devices. Configuration is through a CLI on individual WAE devices or GUI-enabled commands for an entire device group of WAE devices on the Devices tab of the WAAS Central Manager. Once enabled, WAE devices periodically poll a configuration file on the SuperAgent Management Console that specifies the set of data to be exported to the SuperAgent Aggregator.

Figure 9-61 illustrates data feeds for the different segments. Core WAE devices, such as WAE-Core, measure the WAN segment. Branch WAE devices, such as WAE-Edge, measure Client segments. There are two ways to get the data for the Server segment: use SPAN/VACL data from the SuperAgent Collector (SA-CO), or use FlowAgent data from the core WAE (WAE-Core).

The core WAE device (WAE-Core) sends server segment data to the aggregator. However, the SuperAgent Management Console uses the SPAN/VACL data for the server segment if SPAN/VACL data exists. Server segment FlowAgent data is used only if there is no SPAN/VACL data. This is because SPAN/VACL data is more accurate and more complete. SPAN/VACL data is measured closer to the server, so server response time (SRT) is slightly more accurate. Depending on how SPAN/VACL is configured, SPAN/VACL data can also include the server-to-server traffic, which makes the data more complete.

Therefore, the recommended practice is to use SPAN/VACL data for the Server Segment using the SuperAgent Collector. The alternative, using FlowAgent data from the core WAE, is possible for limited WAAS deployments but is not recommended.

Figure 9-62 shows an example network to illustrate how local SPAN can capture traffic on the L2 VLANs interconnecting the distribution switch and the access switch. The following snippet shows an example configuration at the distribution layer switch.

Figure 9-62. Monitoring the Server Segment Example Deployment

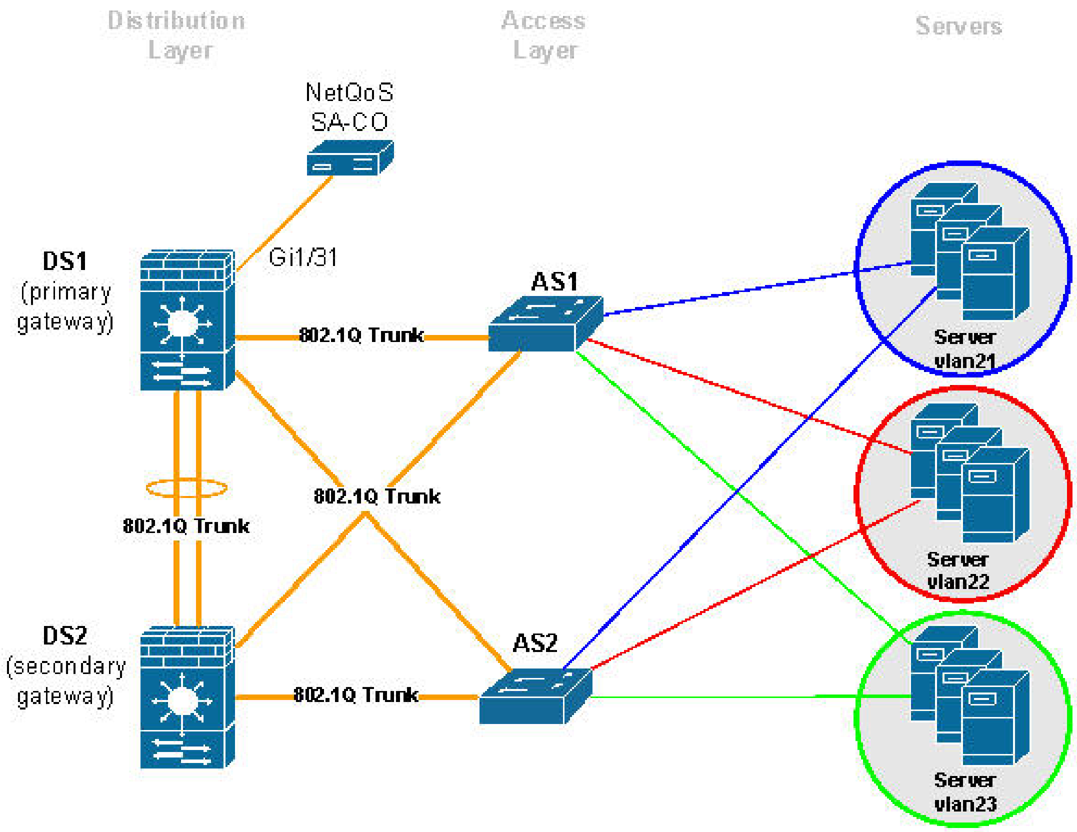

The example assumes that NetQoS SuperAgent Collector (SA-CO) connects to the Gi1/31 port of DS1. NetQoS SuperAgent Collector must monitor VLANs 21 and 22; VLAN 23 traffic does not need to be monitored.

```
1. CONFIGURE SOURCE OF SPAN TRAFFIC
DS1# monitor session 1 source vlan 21-22
2. CONFIGURE NETQOS SUPER AGENT COLLECTOR AS A DESTINATION TO SPAN TRAFFIC
DS1# monitor session 1 destination interface Gi1/31
```

The following commands can be used to verify the configuration.

```
3. VERIFY SPAN CONFIGURATION
DS1# show monitor session 1
Session 1
---------
Type                   : Local Session
Source VLANs           :
    Both               : 21-22
Destination Ports      : Gi1/31
```

9.16.4.2 SuperAgent Data Collection

To configure data collection on a Management Console, the administrator imports or enters a list of network subnets to be monitored. The administrator also configures network types so that groups of networks with similar performance characteristics can be managed collectively. SuperAgent can detect applications and servers, and application names are assigned based on well-know ports. SuperAgent can also import configuration information from third-party managers.

The administrator next assigns servers to Collectors, assigns WAE devices to Aggregators, and ensures that the graphs populate and active sessions are reported to verify the configuration. The administrator can also create aggregations for applications, networks, or servers so that they can be reported and compared as groups.

9.16.4.3 SuperAgent Baselines and Thresholds

SuperAgent automatically computes baselines and thresholds after 48 hours of data collection, and then periodically adapts them to changing conditions. Threshold sensitivity can be adjusted by application, metric, or network type, or can be set to a fixed value. These performance thresholds determine the Normal/Degraded/Excessive ratings, which are visible on the SuperAgent Operations page, for applications, networks, and servers.

When any threshold is violated, SuperAgent launches incidents and can respond to those incidents with email, SNMP trap notifications, and automated investigations. Automated investigations can be configured by application and network type, and can include trace routes, packet captures, application port connects, SNMP polling of affected resources, and ping tests.

The administrator enters thresholds to be referenced in the reports to configure performance SLA reporting. SuperAgent provides application and availability service-level reports and enables drill-down to display time-and resource-based patterns of compliance. This can reveal, for example, that violations tend to occur at certain times or only on certain servers.

The administrator also configures password, role, and permission settings to tailor access and functionality by user. All configurations offers default settings, with full customization available.

9.16.5 Link Traffic Analysis using NetQoS ReporterAnalyzer

In a distributed ReporterAnalyzer configuration, a Harvester collects and condenses raw NetFlow data from routers and switches while a Flow Manager gathers data from the Harvesters and sends it to one or more Data Storage Appliances. A Console processes the stored data and displays it in a web interface. In a standalone configuration, one server containing all of these software components is installed.

When installing ReporterAnalyzer, it is necessary to specify basic information about each router to be monitored, including the source address, SNMP read community string[1], and NetFlow version. It may also be necessary to configure firewalls and other access control lists (ACLs) to allow NetFlow, SNMP, and ReporterAnalyzer process communication.

[1] SNMP community strings are needed in order to obtain router and interface information. ReporterAnalyzer polls devices to obtain SNMP information such as: sysName, ifName, ifDescr, ifSpeed, and so on. ReporterAnalyzer uses these names and descriptions when displaying routers and interfaces. If SNMP community strings are unavailable or incorrect, ReporterAnalyzer displays only device IP addresses and interface indexes.

To begin installation, the administrator executes IOS commands on selected routers to enable NetFlow export and index persistence and confirms that the ReporterAnalyzer Harvester is receiving the data. A management utility displays the interfaces associated with each router and enables the administrator to select or clear the interfaces to be monitored.

The administrator then modifies or confirms report settings that control the frequency of DNS host name resolution, reported interface speeds, router name resolution rules, and the inclusion of router-generated traffic. The administrator also defines roles and user accounts for any users who do not have full administrative privileges for ReporterAnalyzer.

After initial setup, communication settings for third-party applications are configured as required to automatically emailing reports and send SNMP traps. ReporterAnalyzer sends notifications if routers cease to send NetFlow data, or if product components are unavailable or exceed performance thresholds.

The administrator then configures additional features to make reports easy to understand. For example, interfaces can be grouped based on geography, speed, or other attributes for easier selection and comparative reporting. Interfaces can be aggregated so that they can appear together in reports. User-defined protocol groups can be configured so that protocol ranges for each custom application can be reported together. Type of Service (ToS) values are also labeled and grouped for reporting application classes.

9.16.6 Device Performance Analysis using NetQoS NetVoyant

NetVoyant uses SNMP management information bases (MIBs) to define the types of data to be collected. Users can view the MIBs to create event notifications or assist in other administrative tasks, and can compile new MIBs to add functionality to NetVoyant.

NetVoyant creates events and alarms to alert users to issues with NetVoyant services, missed SNMP polls, exceeded utilization, and other configurable thresholds. Administrators can set thresholds, create notifications that NetVoyant triggers in response to events, and view event and alarm logs.

NetVoyant provides a wizard for configuring Cisco IP SLA operations for supported devices. These operations provide data for IP SLA reports in the NetVoyant Reporting Interface.

In a standalone configuration, a NetVoyant Master Console performs administrative, reporting, and polling tasks. In a distributed configuration, the Master Console performs administrative and reporting tasks. Remote servers, called Polling Stations, discover and poll devices on the network.

During initial configuration, a wizard guides the process of discovering network devices. The configuration wizard prompts the user for read community strings of the devices to be polled; both read and write community strings are required for devices that NetVoyant configures to run IP SLA tests.

The configuration wizard prompts the administrator to enter or import the network address ranges of the devices to be discovered. If gateway or backbone routers, or other connectivity devices, are not discoverable (for example, because they are separated from NetVoyant by a public network), it might be necessary to provide the names or addresses of these devices. During setup, a user can also select device classes (for example, printers or workstations) and models to be excluded from automatic polling after discovery.

It might be necessary to configure access lists for routers and other network devices to accept SNMP polls from NetVoyant. Access to a naming service, such as DNS, is also required to present device names in NetVoyant reports.

After the configuration wizard performs initial discovery, NetVoyant performs rediscovery each Midnight by default to update device information. Rediscovery settings can be changed after setup, and rediscovery can also be initiated manually.

9.17 Use Case 1: Predeployment Baselining

Use case 1 illustrates the process of network baselining during the predeployment phase.

The following use case describes an implementation that was deployed in a test lab. The use case provides an example of how the WAN and application optimization solution might work in an real-world scenario. However, the use case is not intended to reflect actual performance or behavior in any environment.

9.17.1 Objectives

Identify candidate sites for WAN and application optimization.

9.17.2 Assumptions

1. WAAS and QoS Policies are not deployed.
2. NetQoS SuperAgent is deployed at the data center Catalyst 6000 Series distribution switch.
3. NetQoS SuperAgent Collector uses SPAN to monitor the DC links and VLANs.
4. NetFlow Data Export is configured on WAN routers to send data to NetQoS ReporterAnalyzer.
5. NetQoS SuperAgent has been deployed for more than 48 hours and has established a baseline.

The following use case example illustrates NetQoS capabilities and features that are useful for identifying candidate sites for WAN and application optimization.

9.17.3 Use Case Example

A multinational enterprise is considering WAN and application optimization and wants to identify the sites that would benefit most from WAN and application optimization. NetQoS has been deployed, and a baseline has been established.

9.17.4 Use Case Workflow

Step 1. Identify worst performing sites and use NetQoS Performance Center to create a customized report.

1. Create a new page containing the following information (see Figure 9-63).
2. Add Router/Switch | Interfaces Over Threshold.
3. Add Network | Bytes by Network
4. Add Network | Incident Count by Network.

5. Add Network | Network Round Trip Time by Network.

6. Add Network | Packet Loss by Network

7. Add Network | Performance by Network

8. Save and view report (see Figure 9-63).

Note: The Interfaces Over Threshold view shows that the Singapore and New York interfaces have average daily utilizations above 95%. Performance reports from SuperAgent indicate that these subnets have high data volumes, high latency, and high packet loss. The Performance by Network view shows excessively poor performance with several network incidents.

We must also ensure that the protocols on these links will benefit from WAN and application optimization.

Step 2. Use ReporterAnalyzer to identify the protocols on the interfaces with performance issues:

1. In the On Interfaces Over Threshold view, click on most utilized interface to get the protocol distribution and volumes optimization, similar to Figure 9-43, Figure 9-44, and Figure 9-45.

2. Repeat for each interface over threshold.

Note: Protocol distribution on interface shows whether WAN and application optimization will be effective on an interface. Protocol distribution also shows what type of WAN and application optimization is most suitable. The Singapore and New York links both appear to be good candidates for WAAS optimization.

The customer chooses New York for the initial WAAS evaluation, without optimizing the Singapore link.

Figure 9-63. NetQoS Performance Center Identifying Candidate Sites for Optimization

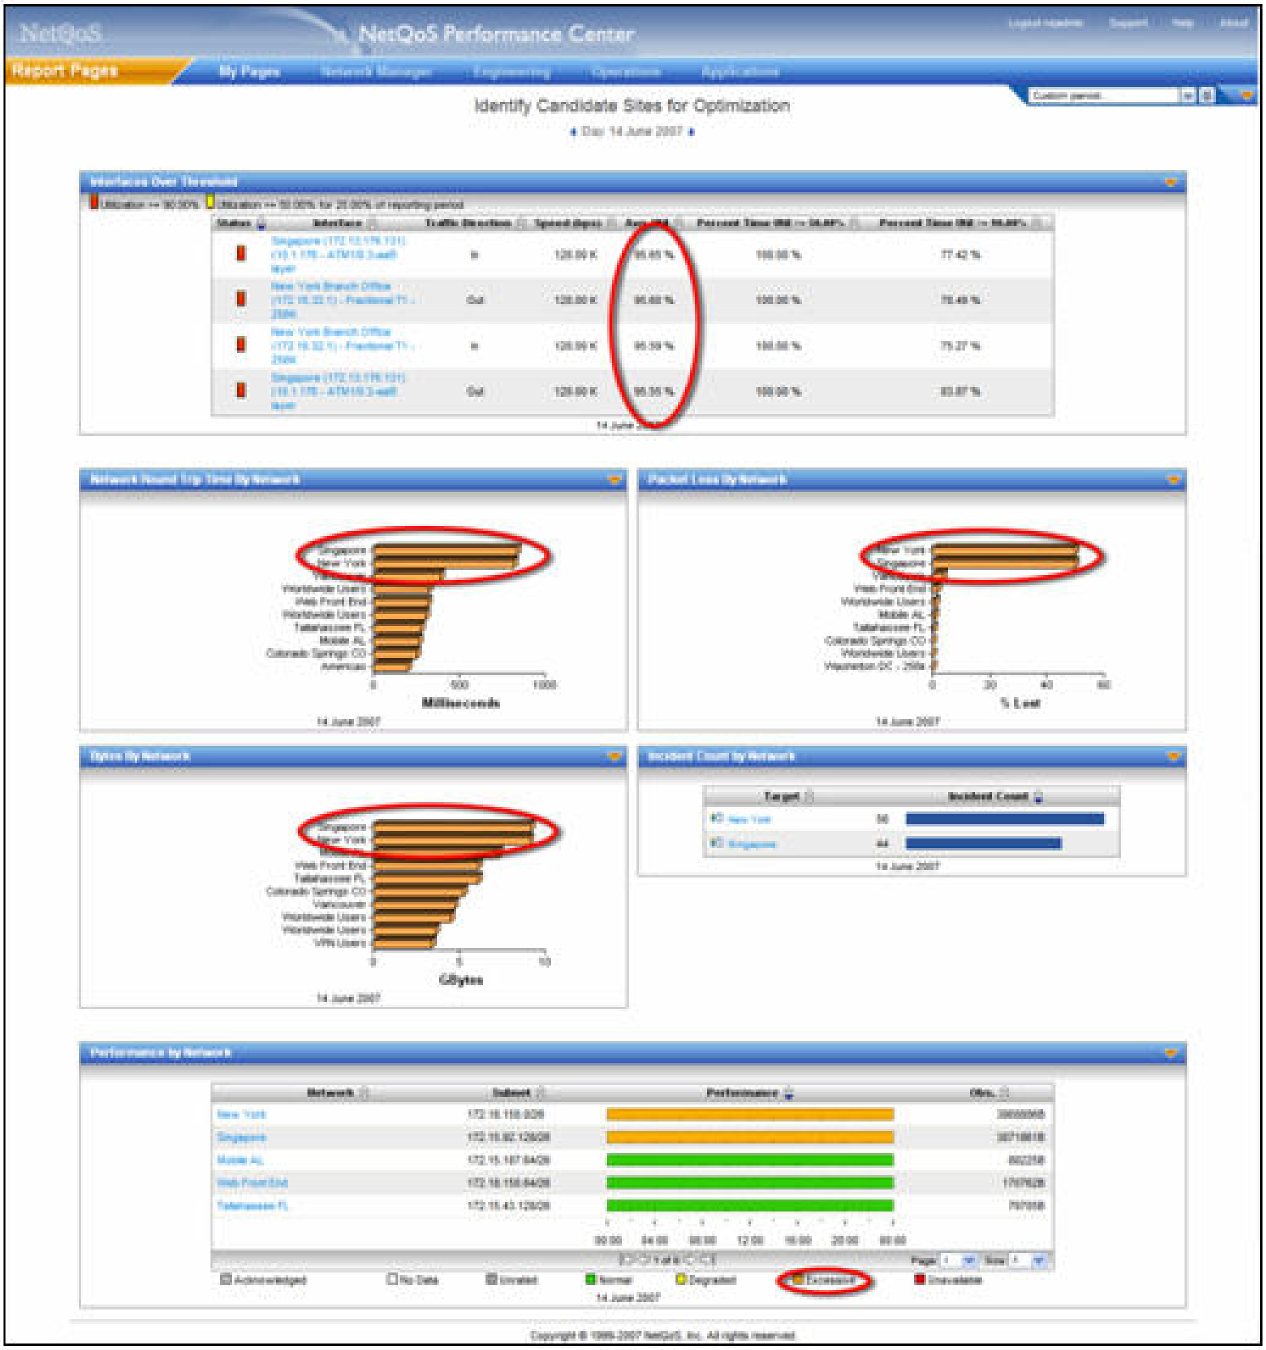

Identify Candidate Sites for Optimization

Day: 14 June 2007

Interfaces Over Threshold

🟥 Utilization >= 90.00% 🟨 Utilization >= 50.00% for 25.00% of reporting period

Status	Interface	Traffic Direction	Speed (bps)	Avg. Util.	Percent Time Util >= 50.00%	Percent Time Util >= 90.00%
🟥	Singapore (172.13.176.131) (10.1.175 - ATM1/0.3-aal5 layer)	In	128.00 K	95.65 %	100.00 %	77.42 %
🟥	New York Branch Office (172.16.32.1) - Fractional T1 - 256K	Out	128.00 K	95.60 %	100.00 %	75.49 %
🟥	New York Branch Office (172.16.32.1) - Fractional T1 - 256K	In	128.00 K	95.59 %	100.00 %	75.27 %
🟥	Singapore (172.13.176.131) (10.1.175 - ATM1/0.3-aal5 layer)	Out	128.00 K	95.55 %	100.00 %	83.87 %

14 June 2007

Network Round Trip Time By Network

14 June 2007

Packet Loss By Network

14 June 2007

Bytes By Network

14 June 2007

Incident Count by Network

Target	Incident Count
New York	56
Singapore	44

14 June 2007

Performance by Network

Network	Subnet	Performance	Obs.
New York	172.16.116.0/26		39555565
Singapore	172.15.92.128/26		30718818
Mobile AL	172.15.187.64/26		802258
Web Front End	172.18.158.64/26		1707628
Tallahassee FL	172.15.43.128/26		797058

▨ Acknowledged □ No Data ▨ Unrated 🟩 Normal 🟨 Degraded 🟧 Excessive 🟥 Unavailable

14 June 2007

Copyright © 1999-2007 NetQoS, Inc. All rights reserved.

9.18 Use Case 2: Validating WAAS Effectiveness

Use case 2 illustrates the process of demonstrating that deploying WAAS helps to optimize WAN and application performance.

The following use case describes an implementation that was deployed in a test lab. The use case provides an example of how the WAN and application optimization solution might work in an real-world scenario. However, the use case is not intended to reflect actual performance or behavior in any environment.

9.18.1 Objectives

Validate performance improvements made possible by WAN and application optimization initiatives.

9.18.2 Assumptions

1. WAAS and QoS policies are deployed.

2. NetQoS SuperAgent is deployed at the data center Catalyst 6000 Series distribution switch.

3. NetQoS SuperAgent Collector monitors the DC links/VLANs using SPAN.

4. NetFlow Data Export is configured to send data to NetQoS ReporterAnalyzer.

5. NetQoS has been deployed for more than 48 hours and has established a baseline.

The following example use case shows how the NetQoS Performance Center was used to validate performance improvements made possible by Cisco WAAS.

9.18.3 Use Case Example

The multinational enterprise in the previous use case has started a pilot WAN and application optimization deployment. During the pilot, WAAS was deployed at 12:00 PM on June 15 to optimize a remote site (New York) that has experienced poor performance. The customer uses NetQoS Performance Center to demonstrate to senior management that Cisco WAAS improves application service delivery, thereby justifying further deployments at other remote sites.

9.18.4 Use Case Workflow

Step 1. Check performance by network; use the customized report from the previous use case in NetQoS Performance Center.

Note: As shown in Figure 9-64,users in New York experience improved network performance as soon as WAAS and QoS policies are implemented. Latency and packet loss for New York show immediate improvement, while the nonoptimized Singapore branch network continues to have serious performance issues.

Figure 9-64. NetQoS Performance Center Showing Improved Behavior

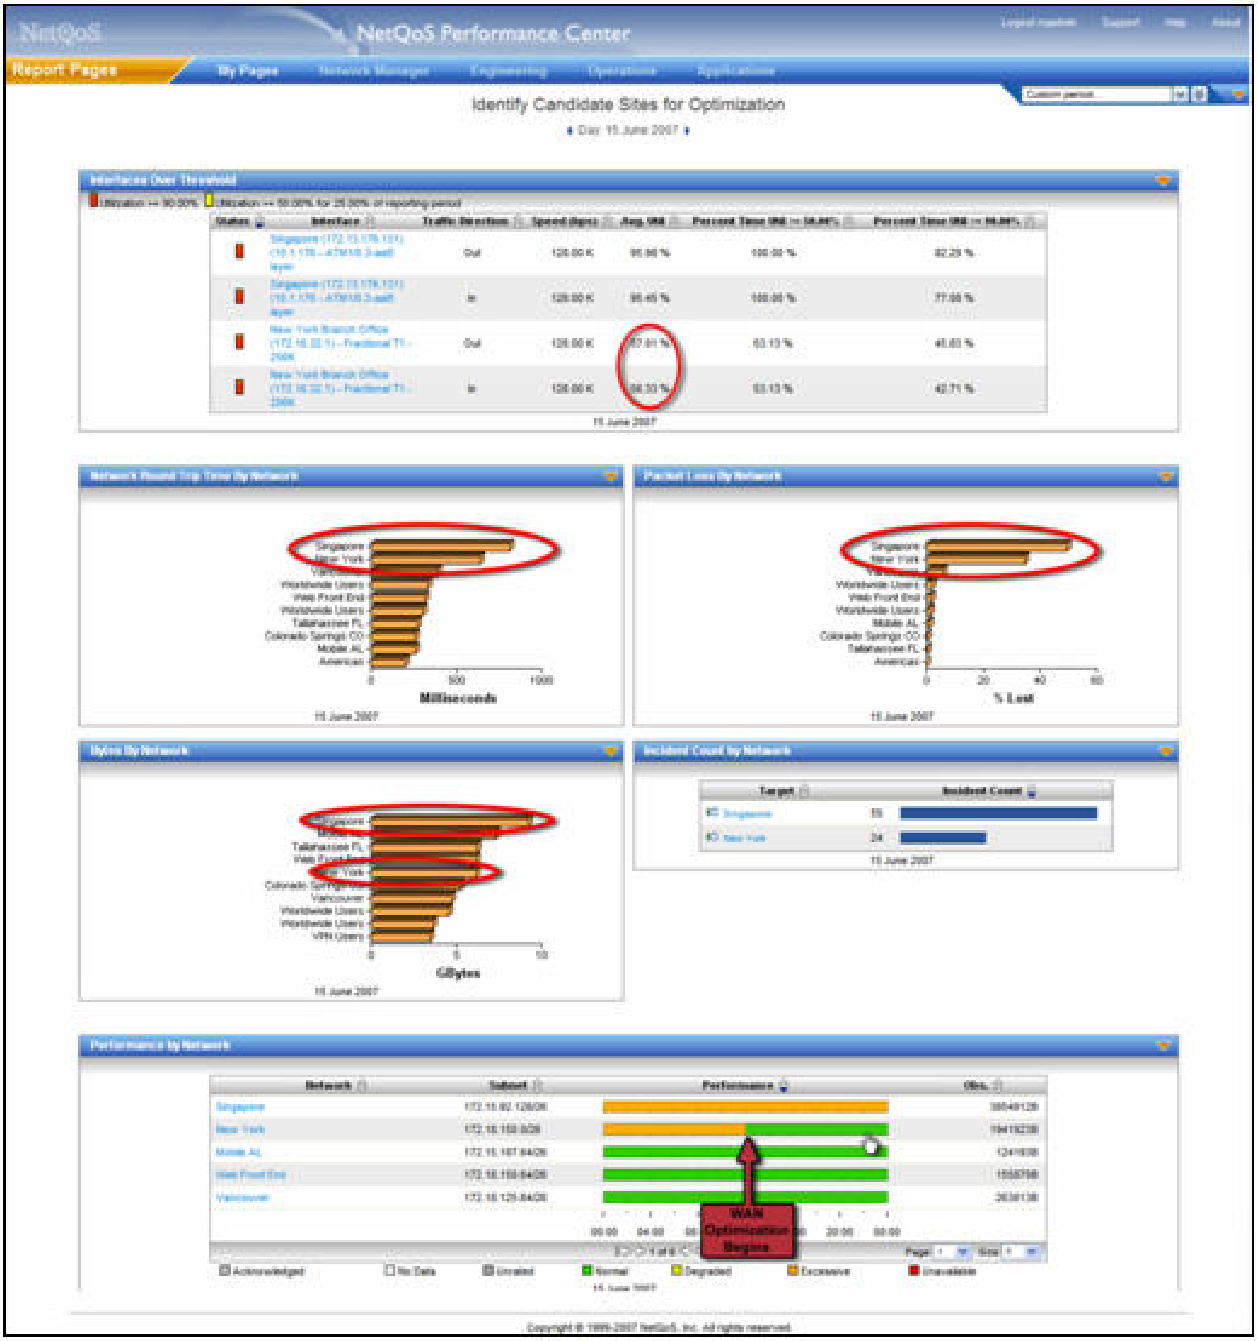

Identify Candidate Sites for Optimization

Day 15 June 2007

Interfaces Over Threshold

🟥 Utilization >= 90.00% 🟨 Utilization >= 50.00% for 25.00% of reporting period

Status	Interface	Traffic Direction	Speed (bps)	Avg. Util	Percent Time Util >= [illegible]	Percent Time Util >= [illegible]
🟥	Singapore (172.13.176.131) (10.1.176 - ATM1/0.3-aal5 layer	Out	128.00 K	95.90 %	100.00 %	82.29 %
🟥	Singapore (172.13.176.131) (10.1.176 - ATM1/0.3-aal5 layer	In	128.00 K	95.45 %	100.00 %	77.08 %
🟥	New York Branch Office (172.16.32.1) - Fractional T1 - 256K	Out	128.00 K	87.01 %	53.13 %	45.83 %
🟥	New York Branch Office (172.16.32.1) - Fractional T1 - 256K	In	128.00 K	88.33 %	53.13 %	42.71 %

15 June 2007

Incident Count by Network

Target	Incident Count
Singapore	55
New York	24

15 June 2007

Performance by Network

Network	Subnet	Performance	Obs.
Singapore	172.15.92.128/26		[illegible]
New York	172.16.150.0/26		[illegible]
Mobile AL	172.15.187.64/26		[illegible]
Web Front End	172.16.150.64/26		[illegible]
Vancouver	172.16.125.64/26		[illegible]

☒ Acknowledged □ No Data ☒ Unrated 🟩 Normal 🟨 Degraded 🟧 Excessive 🟥 Unavailable

15 June 2007

Copyright © 1999-2007 NetQoS, Inc. All rights reserved.

Step 2. Examine the impact of optimization on New York traffic.

1. Click on the New York link in the Performance by Network view (see the pointer in Figure 9-64). This opens the SuperAgent Operations page (Figure 9-65), which shows four network metrics with Degraded or Excessive behavior before the WAAS rollout. The autogenerated performance thresholds, which exist for all Application, Network, Server, and Metric combinations, reveal the operational status (Normal, Excessive, or Degraded).

Note: In this example, (New York Network, Network Round Trip Time, CRM Server) shows Excessive status until WAAS is activated at 12:00PM.

Figure 9-65. Operations Page Showing Dramatic Improvement

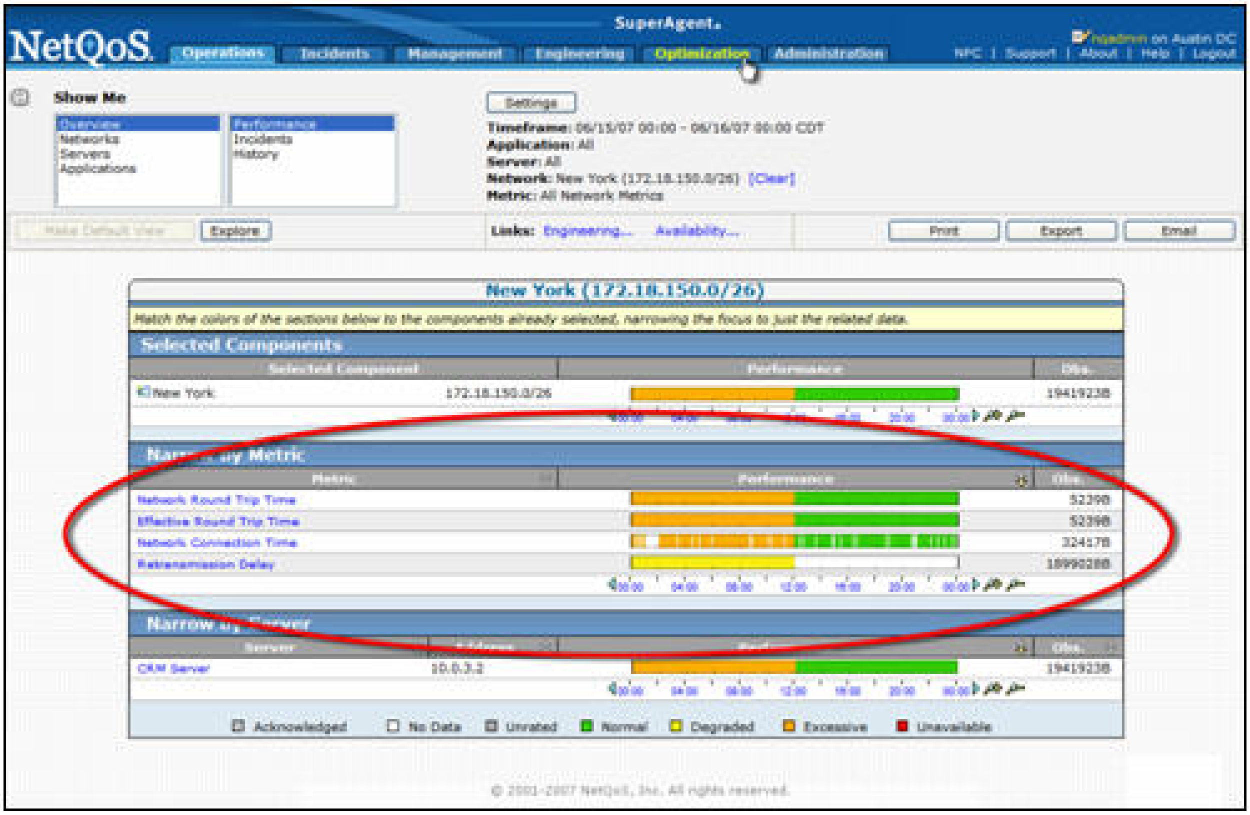

Timeframe: 06/15/07 00:00 - 06/16/07 00:00 CDT
Application: All
Server: All
Network: New York (172.18.150.0/26) [Clear]
Metric: All Network Metrics

Links: Engineering... Availability...

New York (172.18.150.0/26)

Match the colors of the sections below to the components already selected, narrowing the focus to just the related data.

Selected Components

Selected Component		Performance	Obs.
New York	172.18.150.0/26	🟧🟩	1941923B

Narrow by Metric

Metric	Performance	Obs.
Network Round Trip Time	🟧🟩	5239B
Effective Round Trip Time	🟧🟩	5239B
Network Connection Time	🟧🟩	32417B
Retransmission Delay	🟨	1899028B

Narrow by Server

Server	Address	Performance	Obs.
CRM Server	10.0.3.2	🟧🟩	1941923B

© 2001-2007 NetQoS, Inc. All rights reserved.

Step 3. View reports of optimized transaction performance as experienced by users in New York.

1. Click on the Optimization tab (see Figure 9-65) to open the Client Experience for Optimized Transactions view, which shows transaction times for each application (not shown).

2. Click the CRM application to open a Components Report, which includes several views for the application: Response Time (Figure 9-66), SRT (Figure 9-67), Network RTT (Figure 9-68), Retransmission Delay (Figure 9-69), Data Rate (Figure 9-70) and Data Volume (Figure 9-71).

3. Repeat for other key applications.

Figure 9-66 shows a view of response time for users of a CRM application. Although the number of CRM transactions (as represented by the light gray Observations line) remains relatively constant, response time for these transactions drops significantly. These measurements, showing true user experience, are from data collected at the New York branch office WAE device.

Figure 9-66. Response Time View Showing a Five-Fold Performance Improvement

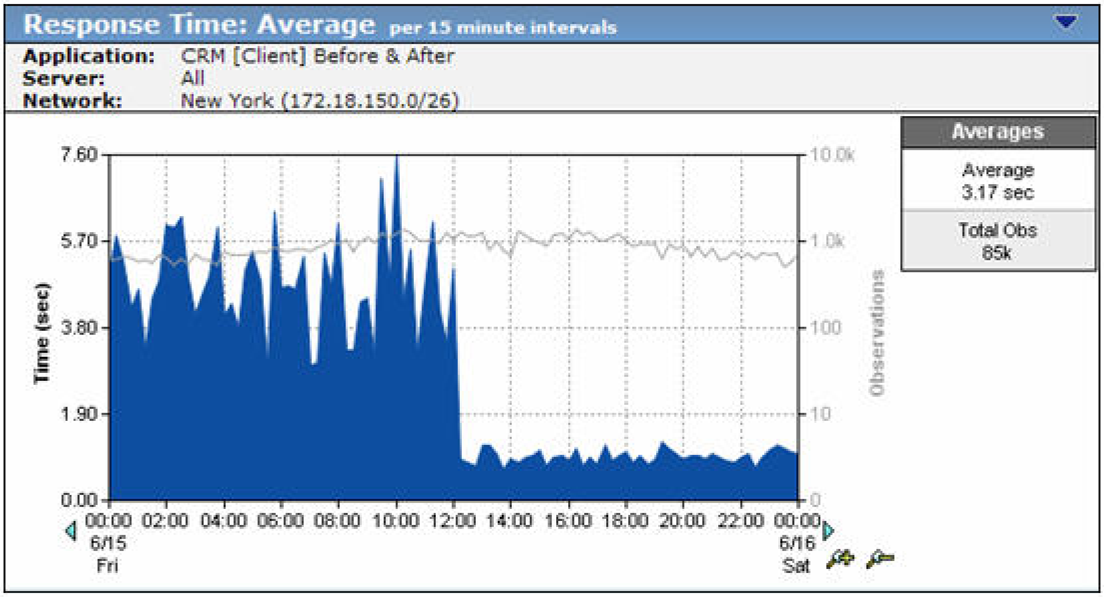

Response Time: Average per 15 minute intervals

Application: CRM [Client] Before & After
Server: All
Network: New York (172.18.150.0/26)

Averages
Average 3.17 sec
Total Obs 85k

On the same page, an SRT graph, measured from the data center WAE (Figure 9-67) shows how much WAAS improves server performance. This specific CRM application uses Windows file servers running CIFS in the data center. While Figure 9-66 shows that the observation count remains constant on the client, the observation count measured at the server declines dramatically because of the server offload made possible by CIFS optimization. With less load on the server after WAAS activation, SRT decreases and performance improves.

Figure 9-67. SRT Showing Server Offload Provided by WAAS

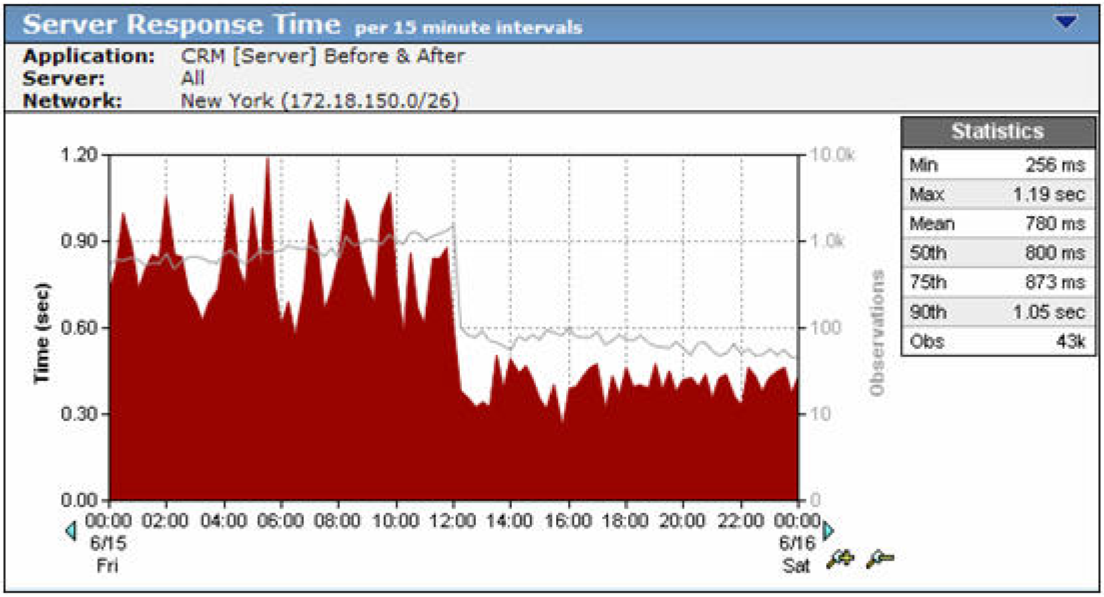

Server Response Time per 15 minute intervals

Application: CRM [Server] Before & After
Server: All
Network: New York (172.18.150.0/26)

Statistics	
Min	256 ms
Max	1.19 sec
Mean	780 ms
50th	800 ms
75th	873 ms
90th	1.05 sec
Obs	43k

A Network RTT graph (Figure 9-68) shows how WAN latency and the number of transactions observed on the link decrease after WAAS deployment. WAN latency is reduced because of optimizing a previously congested link. The use of WAAS optimization techniques such as DRE, LZ, Transport Flow Optimization

(TFO), and caching not only reduces WAN bandwidth, but also decreases the number of TCP transactions on the WAN.

Figure 9-68. Network RTT Showing the Effect of TFO on Network Latency

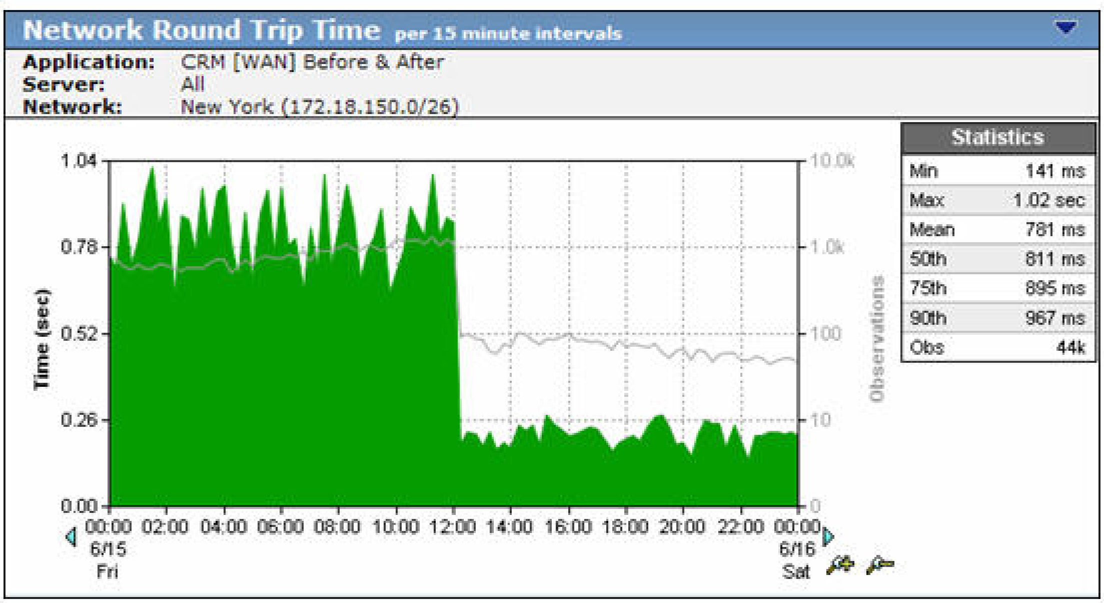

Application: CRM [WAN] Before & After
Server: All
Network: New York (172.18.150.0/26)

Statistics	
Min	141 ms
Max	1.02 sec
Mean	781 ms
50th	811 ms
75th	895 ms
90th	967 ms
Obs	44k

Retransmission delay, shown in Figure 9-69) disappears when WAAS is enabled. This is a typical result. After WAAS is deployed, fewer packets are be retransmitted (the bandwidth reduction causes less congestion), and less time is spent retransmitting packets.

Figure 9-69. Retransmission Delay Virtually Disappears after WAAS Deployment

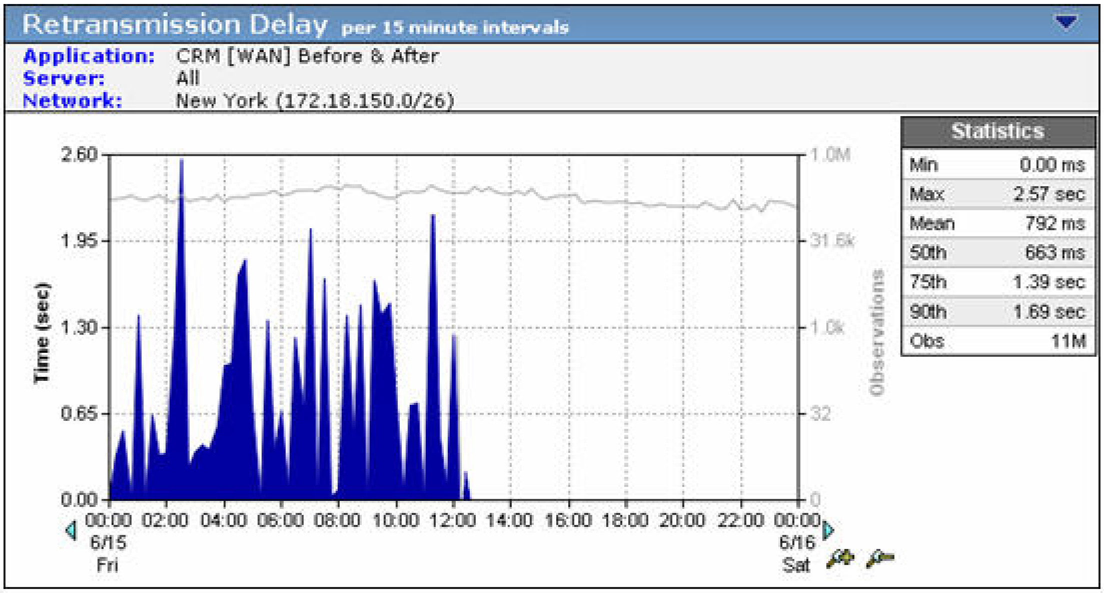

Application: CRM [WAN] Before & After
Server: All
Network: New York (172.18.150.0/26)

Statistics	
Min	0.00 ms
Max	2.57 sec
Mean	792 ms
50th	663 ms
75th	1.39 sec
90th	1.69 sec
Obs	11M

On the same page, the Data Rate graph (Figure 9-70) and Data Volume graph (Figure 9-71) show the impact of WAAS data compression and suppression on WAN traffic. The reduced WAN bandwidth consumption shown in the graphs is the result of Cisco WAAS data redundancy elimination (DRE) and LZ compression.

To document the impact of WAAS the user can access similar views for every critical TCP/IP application and every network.

Figure 9-70. Data Rate over the WAN Showing a Decrease after WAAS Deployment

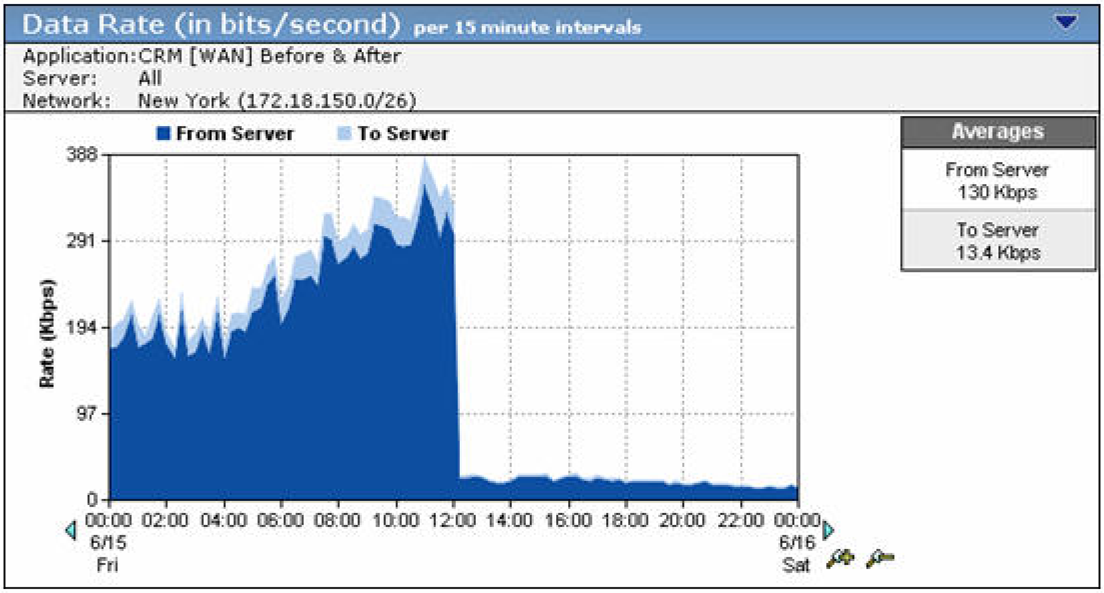

Data Rate (in bits/second) per 15 minute intervals

Application: CRM [WAN] Before & After
Server: All
Network: New York (172.18.150.0/26)

Averages
From Server 130 Kbps
To Server 13.4 Kbps

Figure 9-71. Data Volume over the WAN Decreasing Because of WAAS DRE and LZ Compression

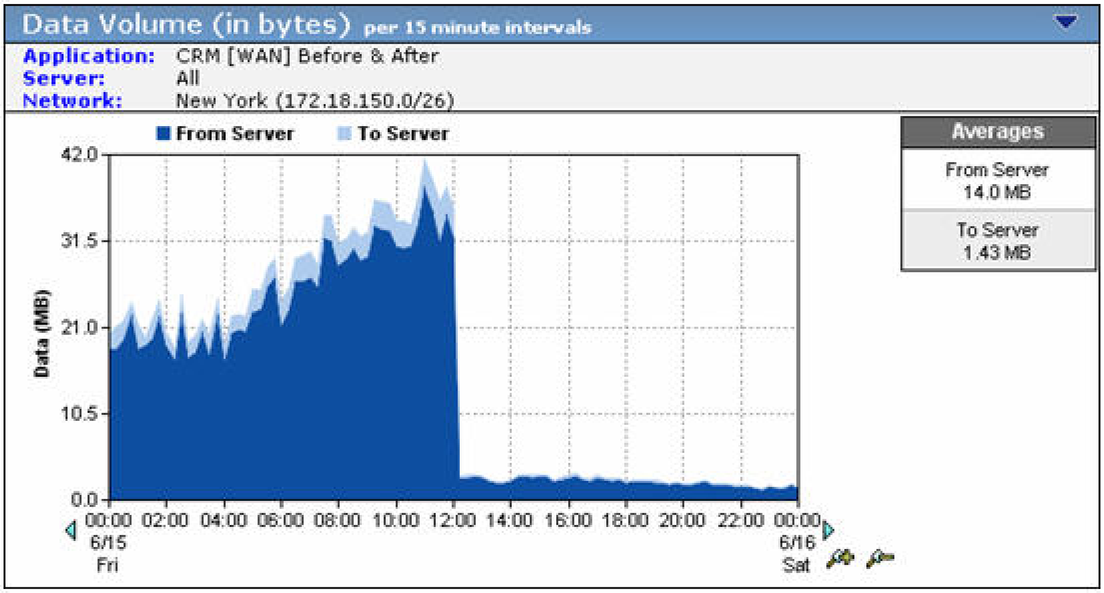

Data Volume (in bytes) per 15 minute intervals

Application: CRM [WAN] Before & After
Server: All
Network: New York (172.18.150.0/26)

Averages
From Server 14.0 MB
To Server 1.43 MB

Step 4. View the overall effect of Cisco WAAS on the New York interface.

1. Click a link to return to the NetQoS Performance Center, and then click the New York interface in the Interfaces Over Threshold view to see stacked protocol trend plots (Figure 9-72).

Note: All but one protocol is reduced significantly by the WAAS deployment. The KM protocol is a pass-through application that is not optimized by WAAS. The volume of KM traffic increases because the lower overall link utilization made possible by WAAS frees capacity for this application.

This result suggests that valid QoS policies should be implemented with WAN and application optimization devices to avoid having newly-freed bandwidth consumed by non-business traffic, such as Internet radio and recreational streaming video. Effective QoS policies can ensure that freed bandwidth is first allocated to business-critical applications.

Figure 9-72. A Stacked Protocol Trend Report Showing Reduced Bandwidth Consumption

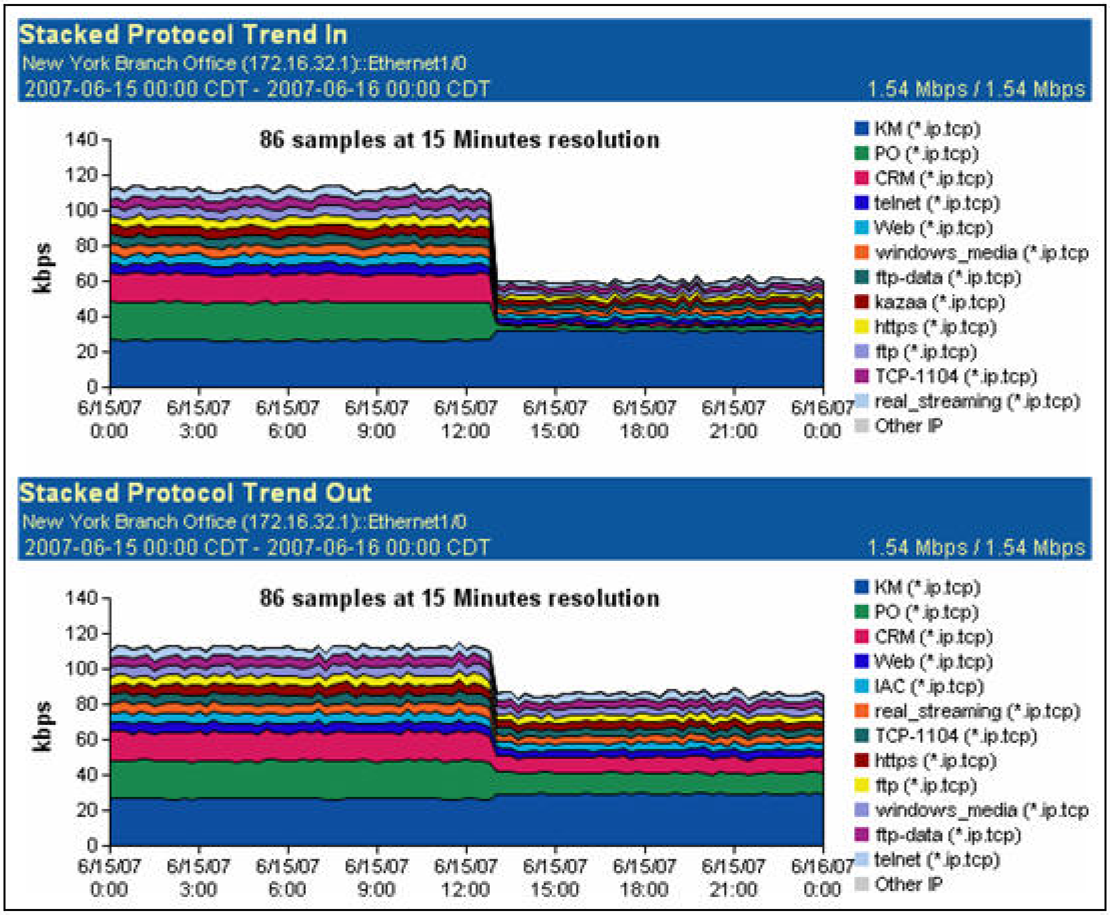

Step 5. View the effect of Cisco WAAS on the New York interface on the second day of the WAAS pilot deployment.

1. Open the customized troubleshooting page in the NetQoS Performance Center (Figure 9-73).

Note: The page shows no further application incidents and no packet loss for New York. Because it is not among the worst performing networks, New York no longer appears in the Performance by Network view.

After WAAS deployment, the NetQoS Performance Center continues to baseline performance so that any deviation from new, faster norms triggers incidents and automatic investigations. The NetQoS Performance Center continues to isolate problems, regardless of whether they occur on optimized links, on nonoptimized links, or in the datacenter.

Figure 9-73. The New York Network No Longer Appears in the Performance by Network View

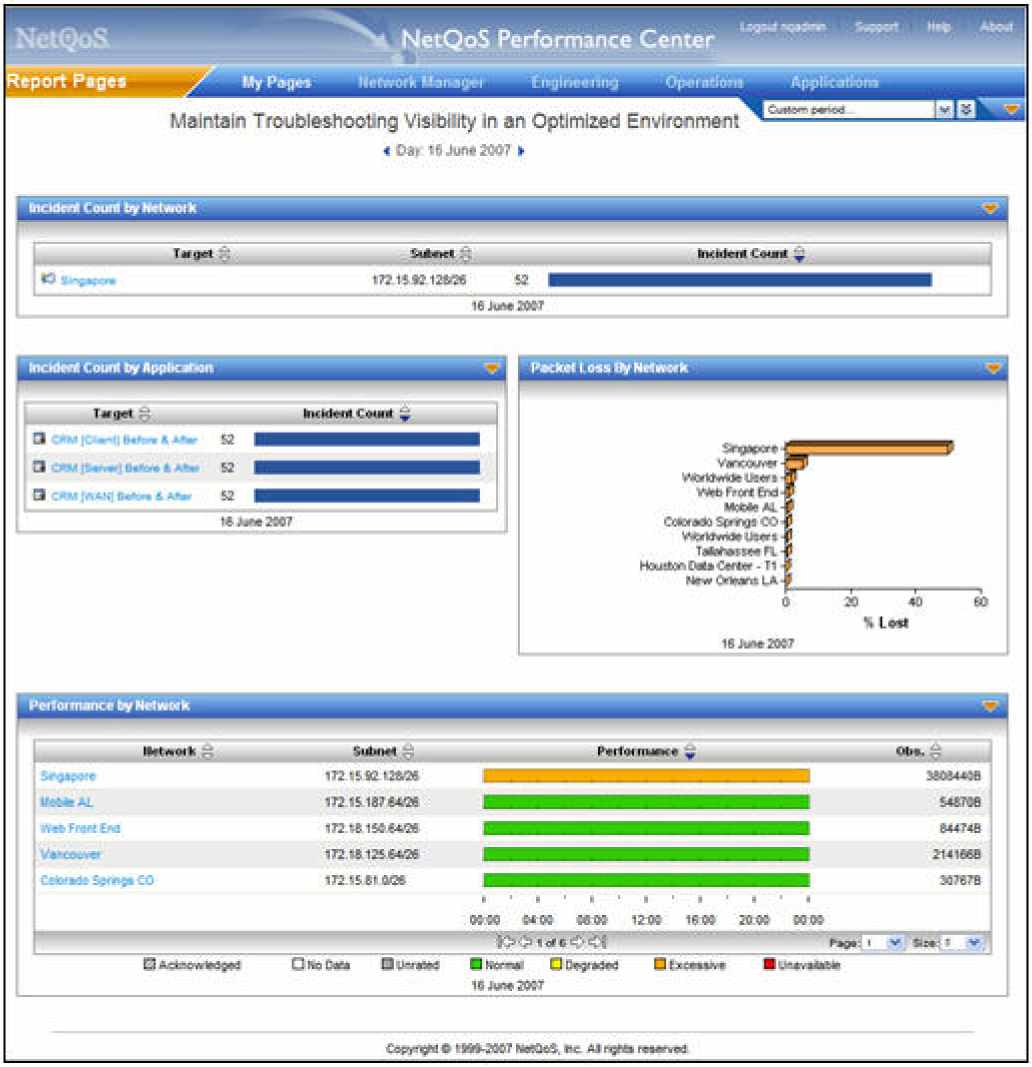

Maintain Troubleshooting Visibility in an Optimized Environment

Day: 16 June 2007

Incident Count by Network

Target	Subnet	Incident Count
Singapore	172.15.92.128/26	52

16 June 2007

Incident Count by Application

Target	Incident Count
CRM [Client] Before & After	52
CRM [Server] Before & After	52
CRM [WAN] Before & After	52

16 June 2007

Packet Loss By Network

16 June 2007

Performance by Network

Network	Subnet	Performance	Obs.
Singapore	172.15.92.128/26		3808440B
Mobile AL	172.15.187.64/26		54870B
Web Front End	172.18.150.64/26		84474B
Vancouver	172.18.125.64/26		214166B
Colorado Springs CO	172.15.81.0/26		30767B

☒ Acknowledged ☐ No Data ▨ Unrated 🟩 Normal 🟨 Degraded 🟧 Excessive 🟥 Unavailable

16 June 2007

Copyright © 1999-2007 NetQoS, Inc. All rights reserved.

NetQoS reports helped the user to make the case for further Cisco WAN and application optimization deployments because they proved that:

- The organization's remote sites and business-critical applications were good candidates for optimization by WAN and application optimization.
- WAN and application optimization improved the performance of critical applications, as experienced by end users.
- After WAN and application optimization deployment, it was possible to maintain network visibility necessary for effective troubleshooting.

9.19 Cisco NAM Use Cases for WAN and Application Optimization

This section focuses on use cases that are based on the deployment of NAM-2 running NAM 3.6 software in the data center. Using NME-NAMs in branch sites to support this solution is not covered in this guide.

The following use cases describes an implementation that was deployed in a test lab. The use case provides an example of how the WAN and application optimization solution might work in an real-world scenario. However, the use case is not intended to reflect actual performance or behavior in any environment.

9.19.1 NAM-2 Support for WAN and Application Optimization

Like NetQoS, NAM can be used to support the different phases of WAN and application optimization deployment. However, there are some differences in how the tools should be used. NAM is not a comprehensive monitoring solution like NetQoS. Instead, NAM is a versatile network tool that provides detailed packet data from the network for conversation-level traffic and response-time analysis. NAM is easy to deploy and configure through a Web graphical user interface.

9.19.1.1 Predeployment support

After NAM-2 is deployed in the data center, users can quickly generate TopN reports and real time charts that identify the application host pairs having the most traffic or the worst performance. This information helps to identify applications and branch sites for WAN and application optimization initiatives (see Figure 9-74).

Figure 9-74. NAM-2 Top Conversations

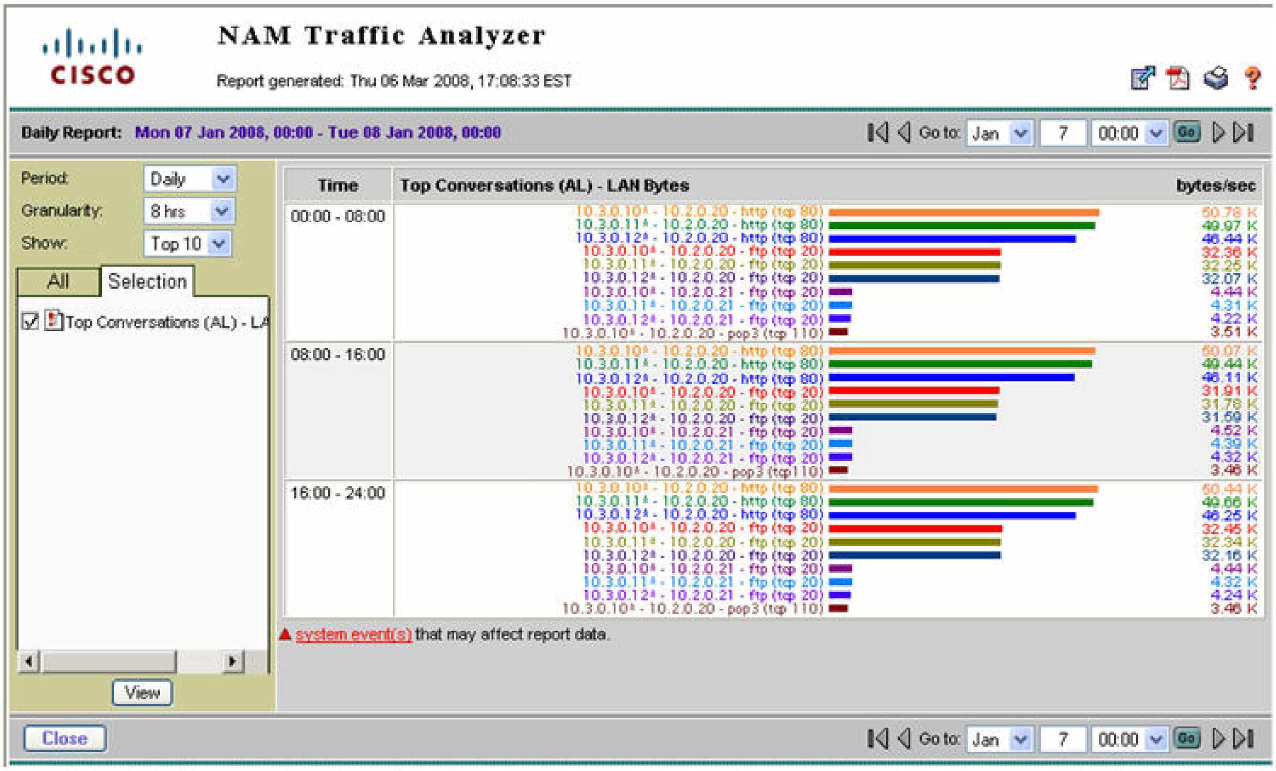

NAM Traffic Analyzer

Report generated: Thu 06 Mar 2008, 17:08:33 EST

Daily Report: Mon 07 Jan 2008, 00:00 - Tue 08 Jan 2008, 00:00

Go to: Jan 7 00:00

Period: Daily

Granularity: 8 hrs

Show: Top 10

☑ Top Conversations (AL) - LA

Time	Top Conversations (AL) - LAN Bytes	bytes/sec
00:00 - 08:00	10.3.0.10 - 10.2.0.20 - http (tcp 80)	50.78 K
	10.3.0.11 - 10.2.0.20 - http (tcp 80)	49.97 K
	10.3.0.12 - 10.2.0.20 - http (tcp 80)	46.44 K
	10.3.0.10 - 10.2.0.20 - ftp (tcp 20)	32.36 K
	10.3.0.11 - 10.2.0.20 - ftp (tcp 20)	32.25 K
	10.3.0.12 - 10.2.0.20 - ftp (tcp 20)	32.07 K
	10.3.0.10 - 10.2.0.21 - ftp (tcp 20)	4.44 K
	10.3.0.11 - 10.2.0.21 - ftp (tcp 20)	4.31 K
	10.3.0.12 - 10.2.0.21 - ftp (tcp 20)	4.22 K
	10.3.0.10 - 10.2.0.20 - pop3 (tcp 110)	3.51 K
08:00 - 16:00	10.3.0.10 - 10.2.0.20 - http (tcp 80)	50.07 K
	10.3.0.11 - 10.2.0.20 - http (tcp 80)	49.44 K
	10.3.0.12 - 10.2.0.20 - http (tcp 80)	46.11 K
	10.3.0.10 - 10.2.0.20 - ftp (tcp 20)	31.91 K
	10.3.0.11 - 10.2.0.20 - ftp (tcp 20)	31.78 K
	10.3.0.12 - 10.2.0.20 - ftp (tcp 20)	31.59 K
	10.3.0.10 - 10.2.0.21 - ftp (tcp 20)	4.52 K
	10.3.0.11 - 10.2.0.21 - ftp (tcp 20)	4.39 K
	10.3.0.12 - 10.2.0.21 - ftp (tcp 20)	4.32 K
	10.3.0.10 - 10.2.0.20 - pop3 (tcp110)	3.46 K
16:00 - 24:00	10.3.0.10 - 10.2.0.20 - http (tcp 80)	50.44 K
	10.3.0.11 - 10.2.0.20 - http (tcp 80)	49.66 K
	10.3.0.12 - 10.2.0.20 - http (tcp 80)	46.25 K
	10.3.0.10 - 10.2.0.20 - ftp (tcp 20)	32.45 K
	10.3.0.11 - 10.2.0.20 - ftp (tcp 20)	32.34 K
	10.3.0.12 - 10.2.0.20 - ftp (tcp 20)	32.16 K
	10.3.0.10 - 10.2.0.21 - ftp (tcp 20)	4.44 K
	10.3.0.11 - 10.2.0.21 - ftp (tcp 20)	4.32 K
	10.3.0.12 - 10.2.0.21 - ftp (tcp 20)	4.24 K
	10.3.0.10 - 10.2.0.20 - pop3 (tcp 110)	3.46 K

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

Unlike NetQoS, NAM-2 reporting shows specific conversation metrics and does not aggregate the information by site pair. The user must manually map host IP addresses to their site subnets to see the global site-to-site traffic pattern and its composition.

9.19.1.2 Support during Deployment

NAM-2 enables the user to examine the traffic or performance metrics for a specific conversation, either in real-time display or in predefined reports. To support WAAS optimization validation, NAM-2 can be configured to monitor both the WAN segment and the Server segment simultaneously.

NAM enables the user to examine the following WAAS benefits for a specific conversation:

- The effectiveness of compression in real-time and history reports of traffic volume on the server and WAN segments (see Figure 9-75).
- Bandwidth reduction on the WAN segment immediately after WAAS is turned on (Figure 9-77).
- The response time metrics reduction on the Server segment immediately after WAAS is turned on.

Figure 9-75. Real-Time NAM-2 Reports Comparing Traffic Volume on the WAN and Server Segments

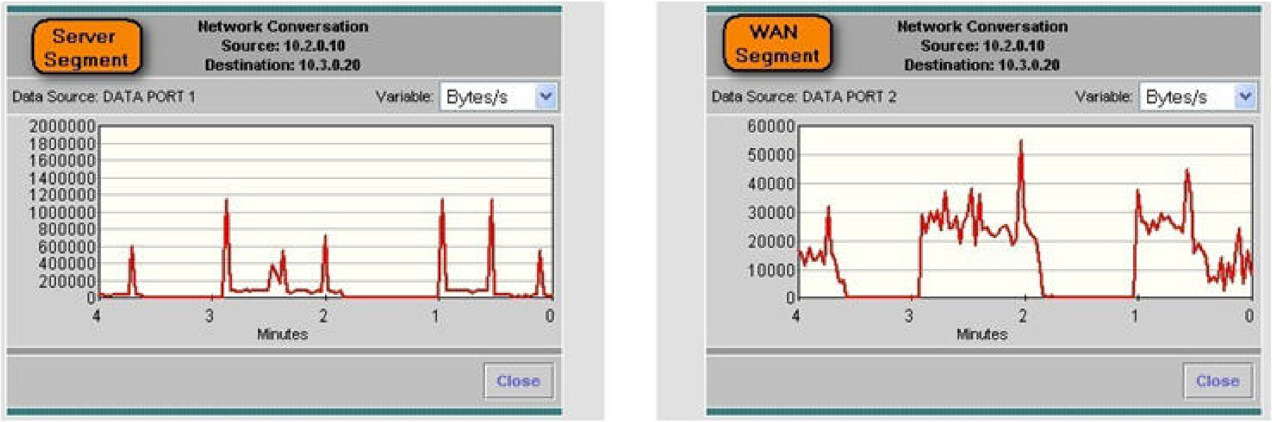

Network Conversation
Source: 10.2.0.10
Destination: 10.3.0.20

Data Source: DATA PORT 1   Variable: Bytes/s

Network Conversation
Source: 10.2.0.10
Destination: 10.3.0.20

Data Source: DATA PORT 2   Variable: Bytes/s

Figure 9-76. NAM-2 History Reports Showing Traffic Reduction on the WAN Segment

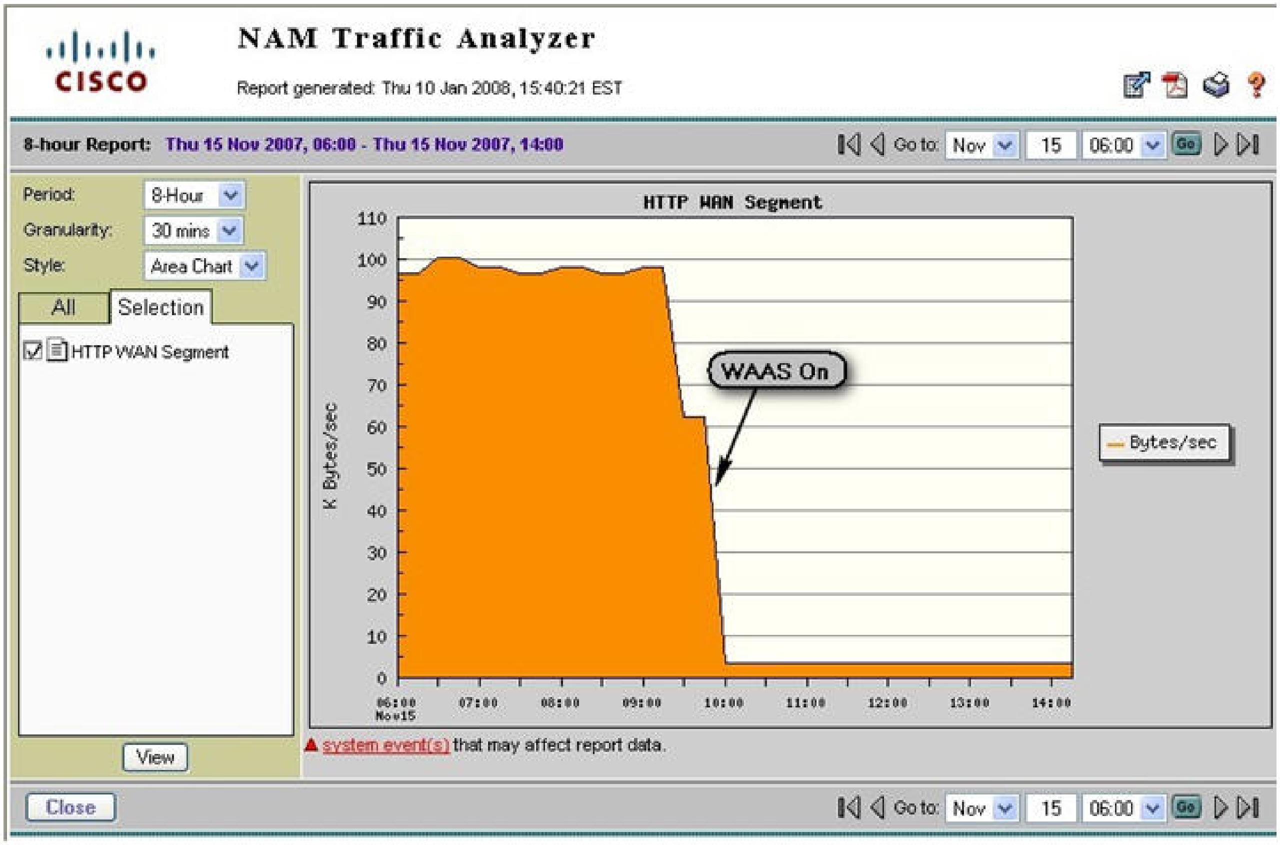

NAM Traffic Analyzer

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

Go to: Nov 15 06:00

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

☑ HTTP WAN Segment

system event(s) that may affect report data.

9.19.1.3 Postdeployment Support

NAM is particularly useful for troubleshooting performance problems, whether or not WAAS is deployed. Because NAM can obtain traffic flow information from different points of the network using SPAN, RSPAN/ERSPAN, VACL, or NetFlow, NAM can be used to isolate problems and help determine whether a problem is with an application server, the DC server network, the WAN access links, or the service provider (SP) WAN.

If WAAS is deployed in the data center, the troubleshooting procedure also includes a step to check whether WAAS is optimizing the target applications. When WAAS is deployed, the user must consider that the original client/server TCP session will now be split into three separate, interrelated segments: client, WAN, and server. NAM-2 can monitor and analyze traffic and performance on the server and WAN segments.

Figure 9-77 shows some questions that are useful to ask when troubleshooting performance problems.

Figure 9-77. Troubleshooting Performance Problems Using NAM-2

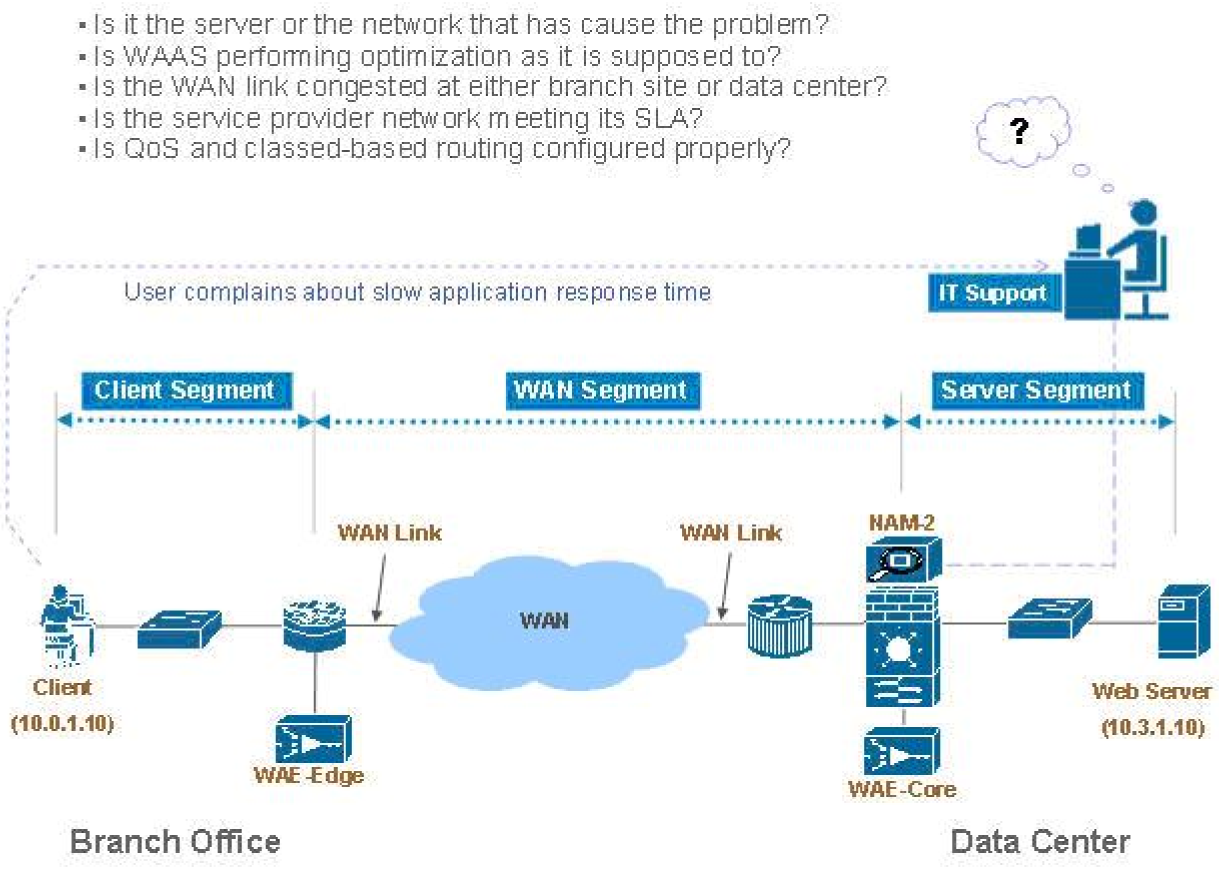
▪ Is it the server or the network that has cause the problem?
▪ Is WAAS performing optimization as it is supposed to?
▪ Is the WAN link congested at either branch site or data center?
▪ Is the service provider network meeting its SLA?
▪ Is QoS and classed-based routing configured properly?
?
User complains about slow application response time
IT Support
Client Segment
WAN Segment
Server Segment
WAN Link
WAN Link
NAM-2
WAN
Client
(10.0.1.10)
WAE-Edge
WAE-Core
Web Server
(10.3.1.10)
Branch Office
Data Center

9.20 NAM 3.6 Metrics for WAN and Application Optimization

NAM-2 provides several ways to capture, measure, and report on application traffic and performance that are relevant to the WAN and application optimization solution. Two categories of reported metrics are useful for this solution: traffic metrics and response-time metrics.

Traffic metrics, such as volume and rate, can be monitored or reported along different aggregation dimensions (for example, application, host, conversation). Top N Charts based on traffic metrics can also be generated, as shown in Figure 9-74).

Response time metrics can be monitored and reported along different aggregation dimensions. Table 9-5 shows some key NAM-2 response time metrics relevant to WAN and application optimization.

Table 9-5. Key NAM-2 Response Time Metrics

Metric relative to	Metric	Description
Networks	Network Delay (ND)	The client-server round-trip delay during TCP connection setup.
	Client Network Delay (CND)	RTT between NAM and client during TCP connection setup.
Servers	Application Delay (AD)	Time f the application server to start responding to a request from the client.
	Server Network Delay (SND)	RTT between NAM and server during TCP connection setup.
	# of clients	Number of branch office clients connecting to a server in the data center.
	# of connections	Number of TCP connections connecting to a server in the data center.
Applications	Transaction Time (TT)	Time it takes to complete a client-server application transaction.
	Number of transactions	Number of client-server application transactions seen by the NAM.

Figure 9-77 illustrates segments monitored by NAM-2 after WAAS is deployed. WAAS creates three network segments: client, WAN, and server. NAM-2 can monitor WAN and server segments.

In the server segment, NAM-2 monitors traffic between the WAE-Core and the server. Transaction time in this segment measures the time for a server to complete an application transaction. This time reflects the server experience, not the client experience.

In the WAN segment, NAM-2 monitors traffic between the WAE-Core and the WAE-Edge. The structure of the transactions is different from in the server segment because the WAAS devices exchange additional control traffic. Monitoring the WAN segment is useful to show WAAS benefits, such as bandwidth reduction on the WAN.

Figure 9-78. NAM-2 Monitoring Segments in the Presence of WAAS

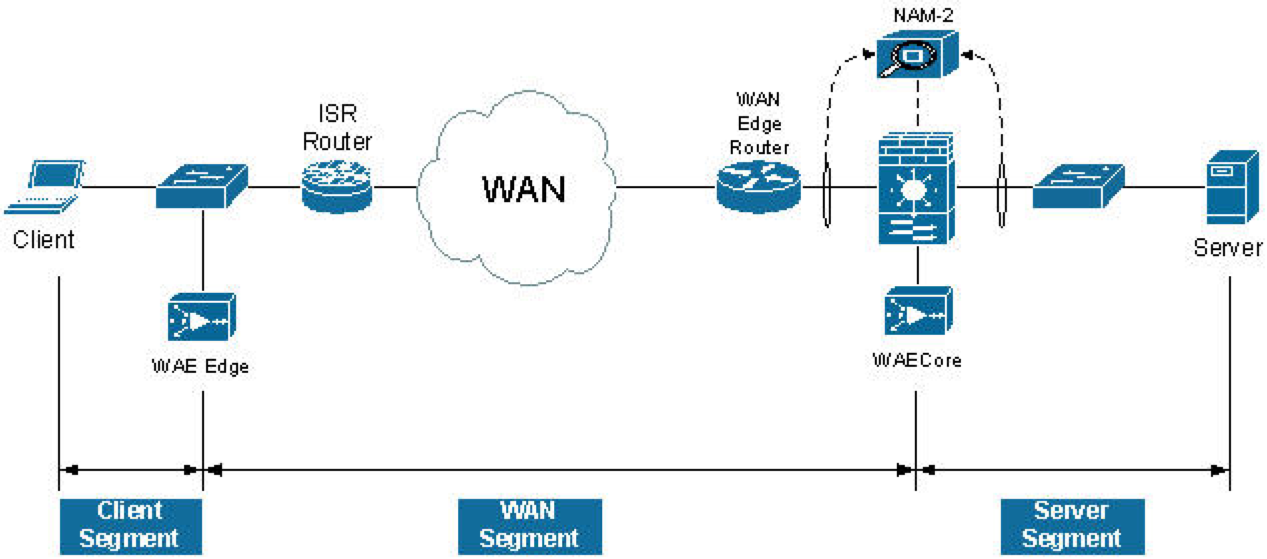

9.21 NAM-2 Deployment Considerations

In a typical deployment scenario for WAN and application optimization, NAM-2 is deployed on a Catalyst 6000 Series distribution switch in the data center. This section describes two data center WAAS deployment scenarios.

In Data Center Deployment Scenario 1, the WAN optimization controllers (WAE1 and WAE2) are deployed at the distribution layer (Figure 9-79). They are connected to the Catalyst 6000 Series distribution switches (DS1 and DS2). Each NAM-2 (NAM1 and NAM2) is installed on a slot of the DS1 and DS2. NAM-2 is a network blade on the Catalyst 6000 Series switch.

Figure 9-79. Data Center WAAS Deployment Scenario 1

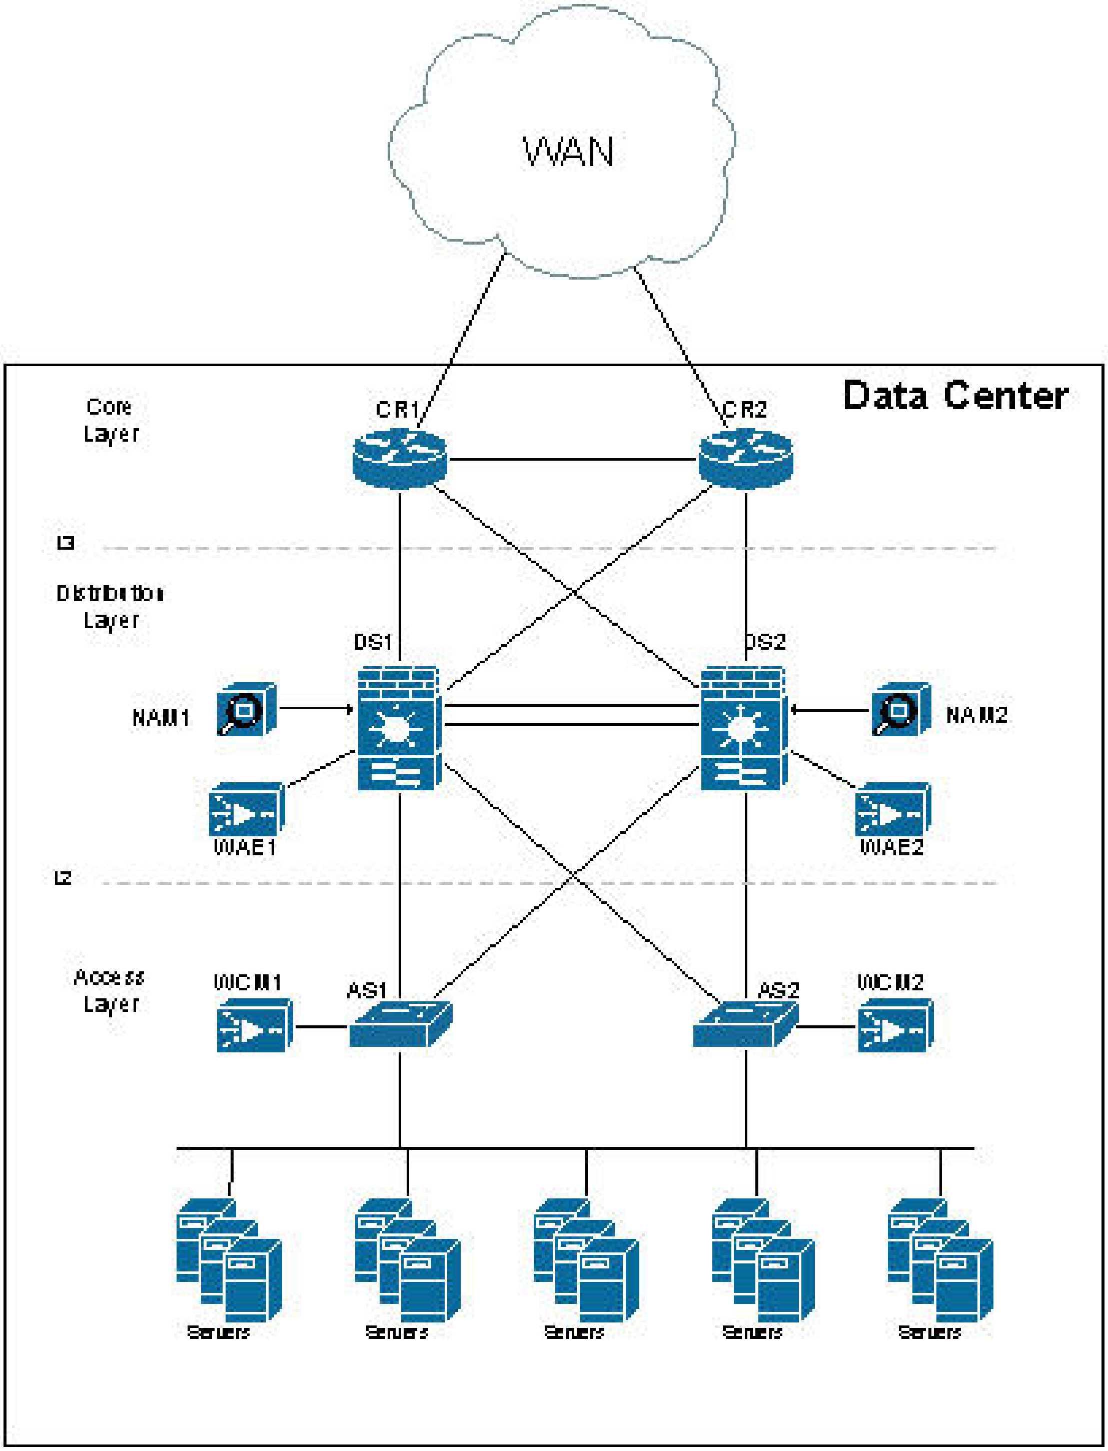

In Data Center Deployment Scenario 2, the WAN optimization controllers (WAE1 and WAE2) are deployed at the core layer, as shown in Figure 9-80), and are connected to the core routers (CR1 and CR2). The core routers can be Cisco 7200 or 7600 Series routers or Catalyst 6000 Series switches.

Figure 9-80. Data Center WAAS Deployment Scenario 2

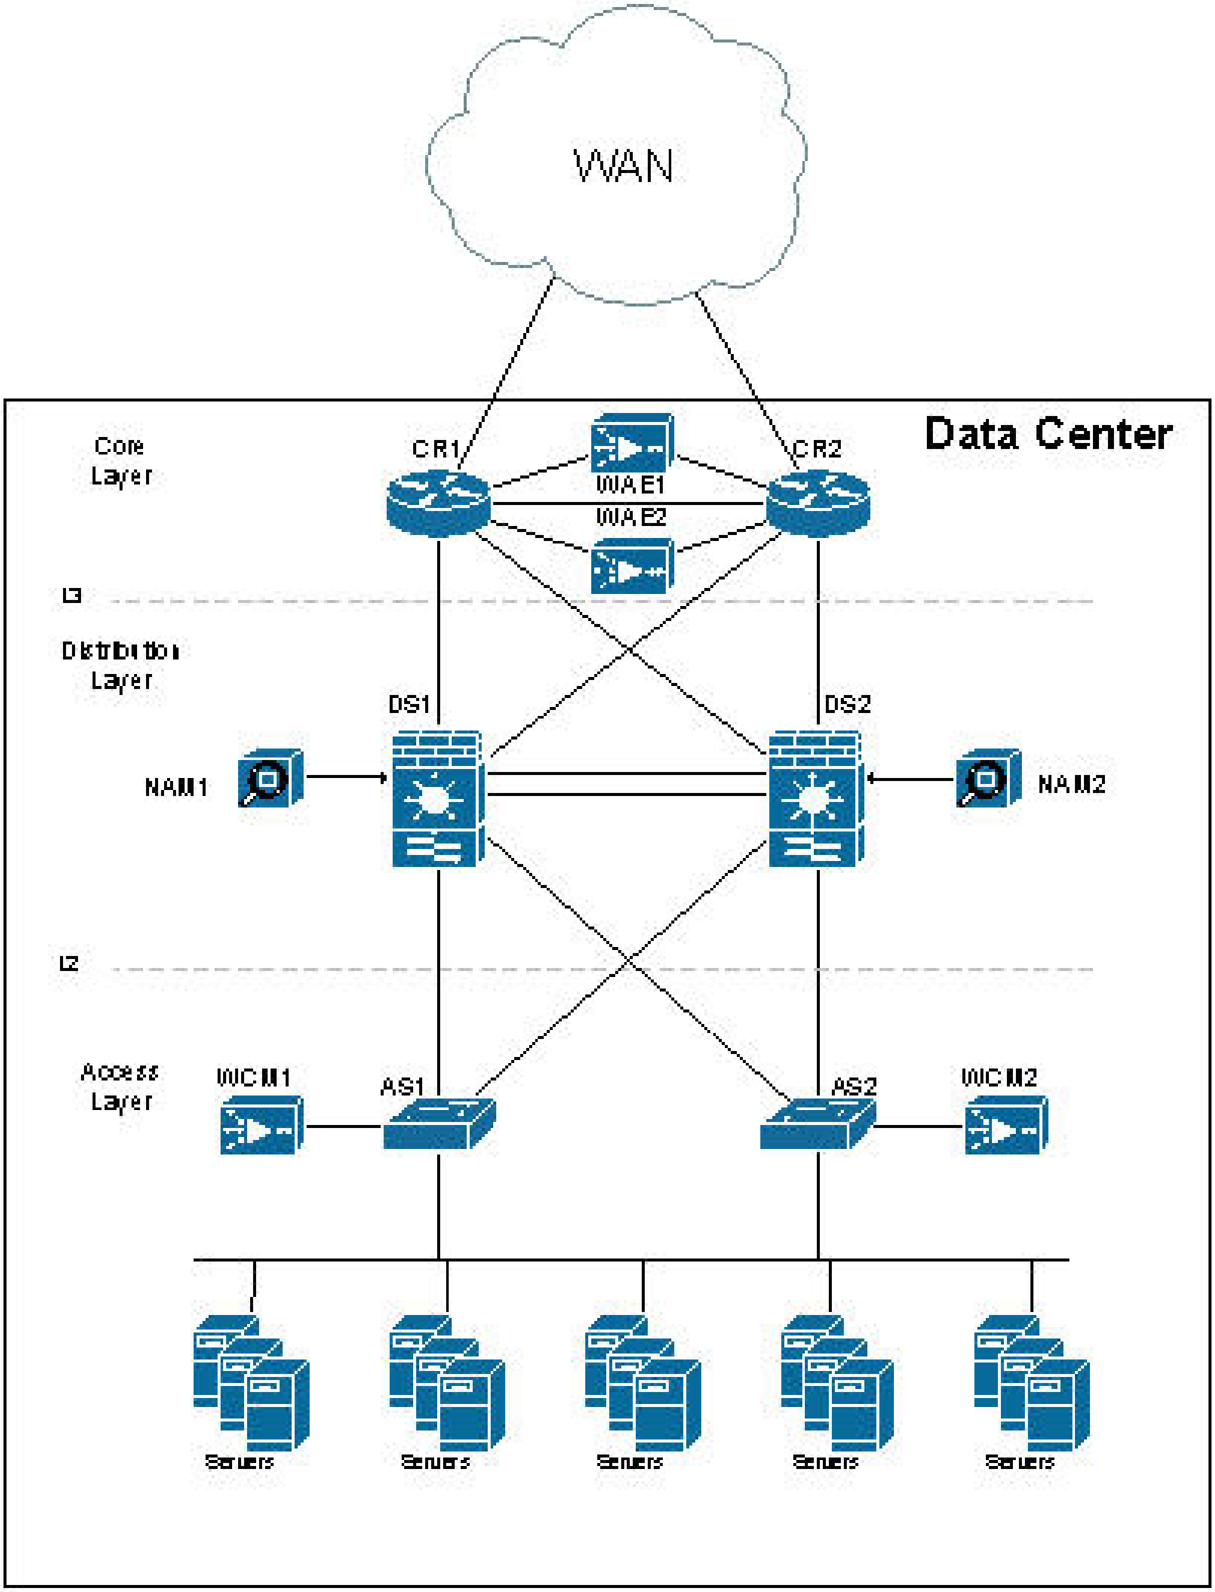

9.22 NAM-2 Data Collection for WAN and Application Optimization

NAM-2 enables the user to select data sources for computing the metrics. For NAM-2, the available data sources follow:

- Local SPAN by VLANs, with visibility for each individual VLAN
- Local SPAN by ports, aggregated by the destination NAM data port
- RSPAN/ERSPAN from another switch by VLANs, with each individual VLAN visible
- RSPAN/ERSPAN from another switch by ports, aggregated by destination data port
- NetFlow data sources either local or remote

Because NAM-2 has only a 1Gb/s capacity, only business critical applications should be monitored.

Data Center Deployment Scenario 1

Figure 9-81 illustrates how the NAM monitors data center traffic for data center deployment scenario 2. In this scenario, WCCP redirects traffic to WAAS (WAE-Core). A NAM-2 monitors both nonintercepted and intercepted traffic. For intercepted traffic, NAM-2 use can use its two data-ports to simultaneously monitor the WAN segment (WAN-WAE traffic) and the server segment (LAN-WAE traffic).

Figure 9-81. NAM-2 Monitoring Configuration for Data Center Deployment Scenario 1

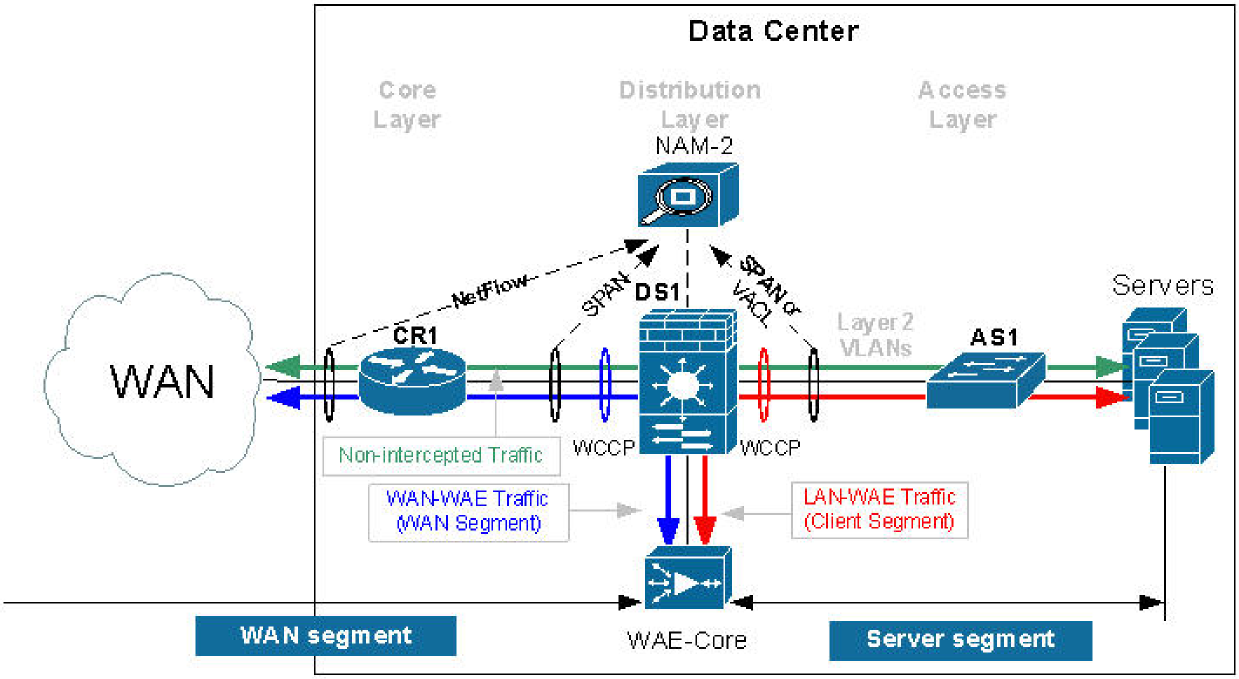

Either local SPAN or VACLs can be used to monitor the server segment. Local SPAN is easier to configure, but is limited to capturing either VLANs or ports. For more granular traffic analysis, VACLs can be used on specified VLANs to match traffic based on source IP address, destination IP address, L4 protocol type, source and destination L4 ports, and other information. VACLs are very useful for granular traffic identification and filtering.

Note that the optimized TCP traffic (WAN-WAE traffic and LAN-WAE traffic) is copied to the NAM-2, using SPAN or VACL, before the traffic is redirected to the WAAS device using Web Cache Coordination Protocol (WCCP).

To monitor the WAN segment, NetFlow Data Export on the WAN router can be configured using NAM-2 as a destination. Another option is to use SPAN on the uplinks to the core router; these are typically L3 links, and therefore it is possible to only SPAN the ports of the L3 links.

Note, however, that there is a limit of two local SPAN sessions on the Catalyst 6000 Series switch. This can become an issue if a Firewall Services Module (FWSM) is used on the distribution Catalyst 6500: FWSM uses by default a SPAN session. For more details, refer to Cisco Data Center Infrastructure Design Guide 2.1 Release Notes.

Possible workarounds to this issue are:

1. Use VACLs for the server segment and SPAN for the WAN segment.

2. Use the **no monitor session service module** command to disable the FWSM SPAN session. In this case, care must be taken so that multicast sources are not placed behind the FWSM.

3. Use RSPAN on the Access Switch for the Server Segment.

9.22.1 Monitoring the Server Segment

To monitor the server segment, the recommended method is to use local SPAN to capture traffic on the L2 VLANs interconnecting the distribution Catalyst 6000 Series switch (see Figure 9-82) and the access switch.

Figure 9-82 illustrates how the NAM should monitor data center traffic for data center deployment scenario 2.

Figure 9-82. Monitoring the Server Segment Example Deployment

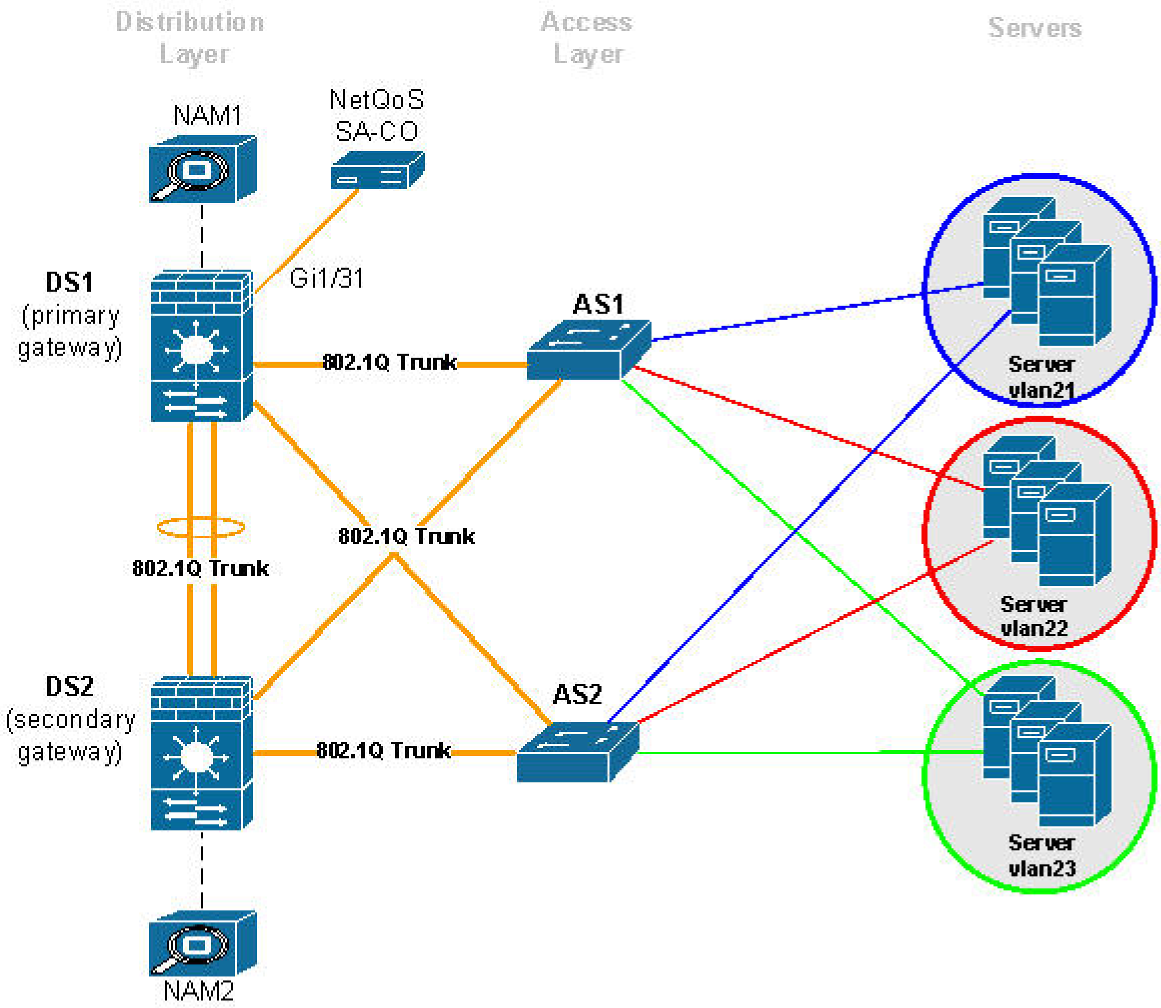

The following snippet shows an example configuration at the distribution layer switch. The example assumes that NAM-2 (NAM1) is installed on slot 9 of the Catalyst 6000 Series switch DS1. Here, VLANs 21 and 22 are monitored on NAM-2 data port 1. NAM does not monitor VLAN 23. The example also shows how to monitor the same VLANs using NetQoS SuperAgent Collector (SA-CO) connected at Gi1/31.

```
1. CONFIGURE SOURCE OF SPAN TRAFFIC
DS1# monitor session 1 source vlan 21-22

2. CONFIGURE NAM DATA PORT AS A DESTINATION TO SPAN TRAFFIC
DS1# monitor session 1 destination analysis-module 9 data-port 1

3. OPTIONAL : CONFIGURE NETQOS SUPER AGENT COLLECTOR
DS1# monitor session 1 destination interface Gi1/31
```

The following commands can be used to verify the configuration.

```
4. VERIFY SPAN CONFIGURATION
DS1# show monitor session 1
Session 1
---------
Type                    : Local Session
Source VLANs            :
    Both                : 21-22
Destination Ports       : Gi1/31analysis-module 9 data-port 1
```

An alternative option is to use VACLs. The following snippet shows an example configuration at the distribution layer switch. As before, the example assumes that NAM-2 is installed on slot 9 of the switch. Here only TCP traffic on VLANs 21, 22 and 23 is captured on data port 1 of NAM-2.

```
1. CONFIGURE NAM DATA PORT TO CAPTURE TRAFFIC
analysis module 9 data-port 1 capture
analysis module 9 data-port 1 capture allowed-vlan 1-4094

2. CONFIGURE ACCESS LISTS TO CAPTURE TRAFFIC – TCP TRAFFIC IN THIS CASE
access-list 110 permit ip any any
access-list 120 permit tcp any any

3. CONFIGURE VLAN ACCESS MAP - which packets to capture, which to forward on
vlan access-map LAN 100
 match ip address 120
 action forward capture
 exit
vlan access-map LAN 200
 match ip address 110
 action forward
 exit

4.APPLY ACCESS MAP TO VLANs
vlan filter LAN vlan-list 21-22
```

The following commands can be used to verify the configuration.

```
5. VERIFY VLAN FILTER CONFIGURATION
DS1# show vlan filter
VLAN Map LAN:
        Configured on VLANs:  21-22
            Active on VLANs:  21-22

6. VERIFY ACCESS-MAP CONFIGURATION
DS1# show vlan access-map LAN
Vlan access-map "LAN"  100
        match: ip address 120
        action: forward capture
Vlan access-map "LAN"  200
        match: ip address 110
        action: forward

7. VERIFY NAM PORT CONFIGURATION
DS1# show analysis module 9 data-port 1 state
Analysis module 9 data-port 1:

Switchport: Enabled
Administrative Mode: trunk
Operational Mode: capture
Administrative Trunking Encapsulation: dot1q
Negotiation of Trunking: Off
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Trunking VLANs Enabled: NONE
Pruning VLANs Enabled: 2-1001
Capture Mode Enabled
Capture VLANs Allowed: ALL
Vlans allowed on trunk: none
Vlans allowed and active in management domain: none
Vlans in spanning tree forwarding state and not pruned:
   none
Administrative Capture Mode: Enabled
Administrative Capture Allowed-vlans: 1-4094
```

9.22.2 Monitoring the WAN Segment

The recommended option for monitoring the WAN segment is to export NetFlow Data from the core router WAN interface link to the NAM-2 at the distribution layer switch.

The following snippet shows an example configuration on the core router CR1 based on the network shown in Figure 9-83. The example assumes that the IP address for NAM-2 is 10.10.0.2.

Figure 9-83. NetFlow Data Export to NAM Example

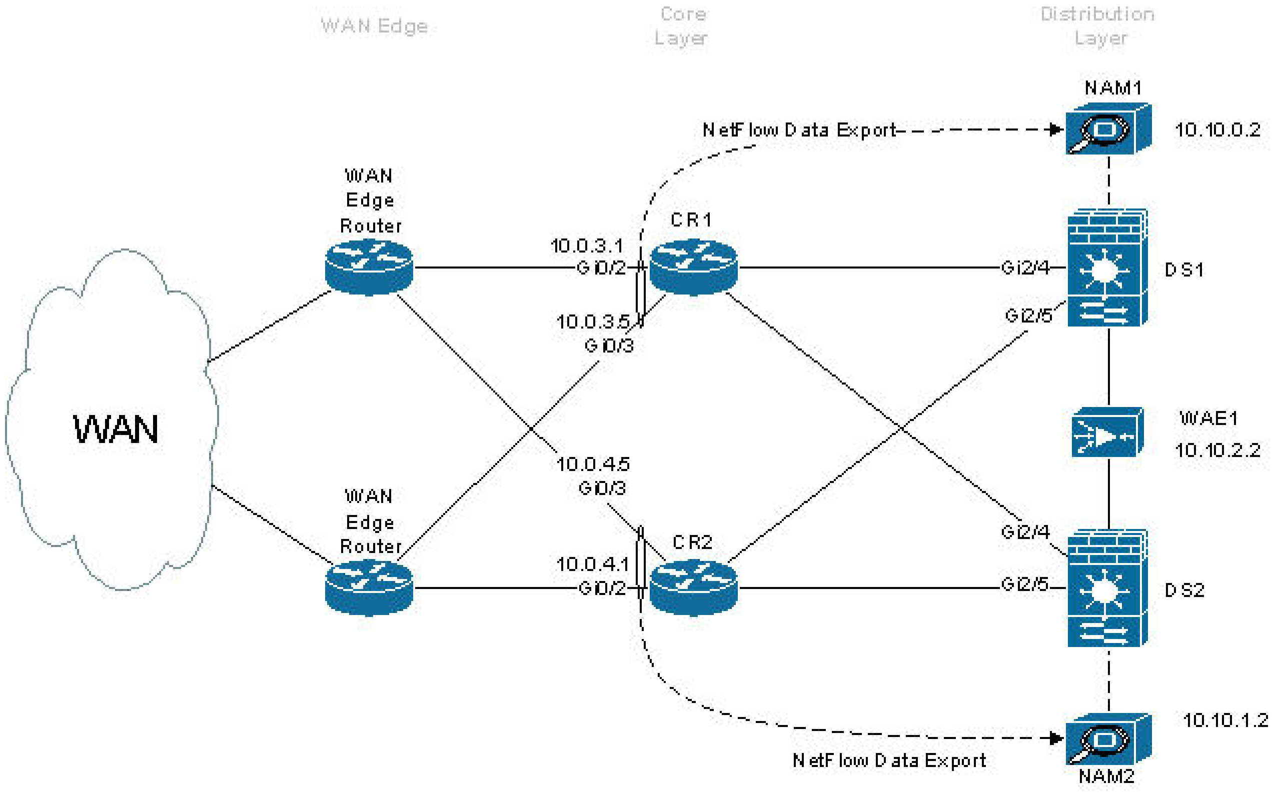

```
1. CONFIGURE CORE ROUTER CR-1 NETFLOW DATA EXPORT
ip flow-export version 9
interface GigabitEthernet0/2
 description Interface to WAN
 ip address 10.0.3.1 255.255.255.252
 ip flow ingress
 ip flow egress
 ip route-cache flow
interface GigabitEthernet0/3
 description Interface to WAN
 ip address 10.0.3.5 255.255.255.252
 ip flow ingress
 ip flow egress
 ip route-cache flow

2. CONFIGURE NAM AS A DESTINATION FOR NETFLOW DATA EXPORT
ip flow-export destination 10.10.0.2 3000
```

The user can then use the Web GUI to configure NAM-2 to monitor traffic from that NetFlow data source.

Note: NetFlow Data export supports traffic metrics in NAM-2, but does not support response time metrics.

SPAN on the uplinks to the core router can also be configured on the distribution layer Catalyst 6000 Series switch (DS1). As before, the example assumes that NAM-2 is installed on slot 9 of the switch. Here, traffic on the core uplink interfaces (Gi2/4 and Gi2/5) is captured on data port 2 of NAM-2.

```
1. CONFIGURE SOURCE OF SPAN TRAFFIC
DS1# monitor session 2 source interface Gi2/4, Gi2/5

2. CONFIGURE NAM DATA PORT AS A DESTINATION TO SPAN TRAFFIC
DS1# monitor session 2 destination analysis-module 9 data-port 2

3. VERIFY SPAN CONFIGURATION
DS1# show monitor session 2
Session 2
---------
Type                    : Local Session
Source Ports            :
    Both                : Gi2/4 Gi2/5
Destination Ports       : analysis-module 9 data-port 2
```

A similar configuration can be applied to the second NAM-2 on core router CR2 and distribution switch DS2.

9.23 Data Center Deployment Scenario 2

Figure 9-84 shows the monitoring setup options for NAM-2 in the context of WAAS Data Center Deployment Scenario 2. As before, WCCP redirects traffic to WAAS, and one NAM-2 uses its two data ports to simultaneously monitor the WAN segment and the server segment.

Figure 9-84. NAM-2 Monitoring Configuration for Data Center Deployment Scenario 2

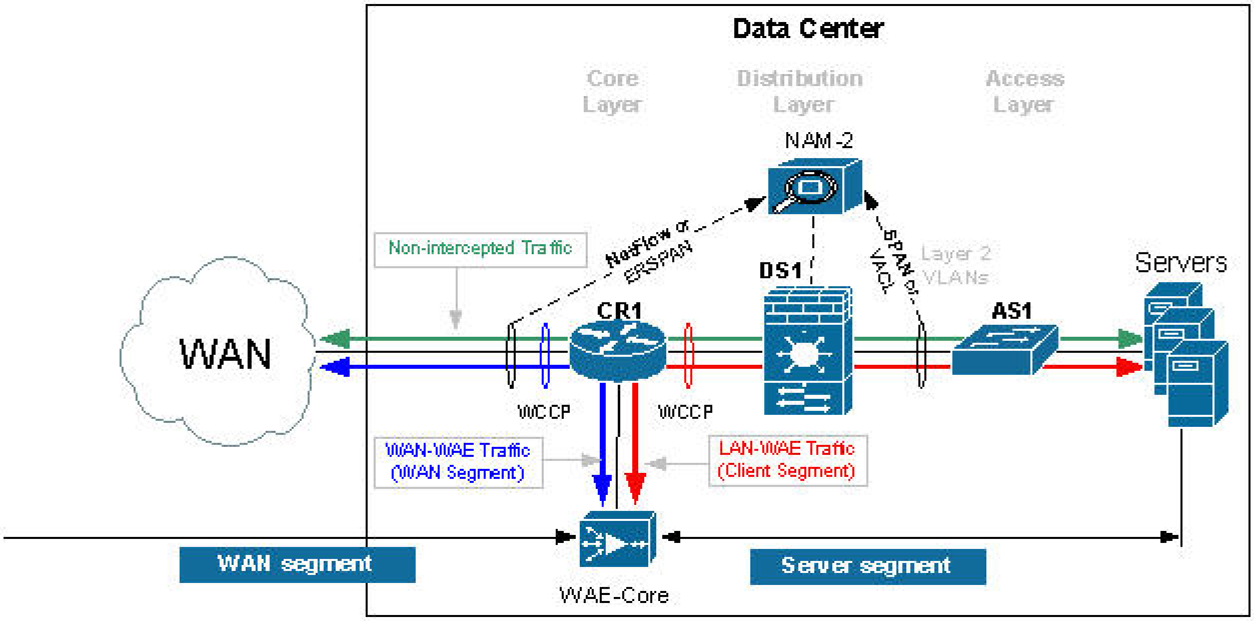

Either local SPAN or VACLs can be used to monitor the server segment on the access links as in the previous scenario. Either ERSPAN or NetFlow Data Export can be used to monitor the WAN segment. If the data center core router is a Cisco 7200 Series, only the NetFlow Data Export option is available for monitoring the WAN side.

9.23.1 Monitoring the Server Segment

As in Data Center Deployment Scenario 1, the recommended option is to use local SPAN to capture traffic on the L2 VLANs. An alternative option is to use VACLs. See example configuration of previous deployment scenario.

9.23.2 Monitoring the WAN Segment

As in Data Center Deployment Scenario 1, the recommended option for monitoring the WAN segment is to export NetFlow Data from the core router to the NAM-2 at the distribution layer Catalyst 6000 Series switch. See example configuration of previous deployment scenario.

In this particular deployment scenario, an alternative option is to use ERSPAN. ERSPAN is only available on Catalyst 6000 Series switches and Cisco 7600 series routers on IOS Release 12.2(18)SXE and later. The disadvantage of ERSPAN is that it creates additional traffic between the core layer router and the distribution layer switch.

The following snippets show example configurations on the core router and the distribution switch. Figure 9-85 shows the network used in the configurations. First, the ERSPAN traffic source must be configured on the core router 7200-1.

Figure 9-85. ERSPAN Configuration Example

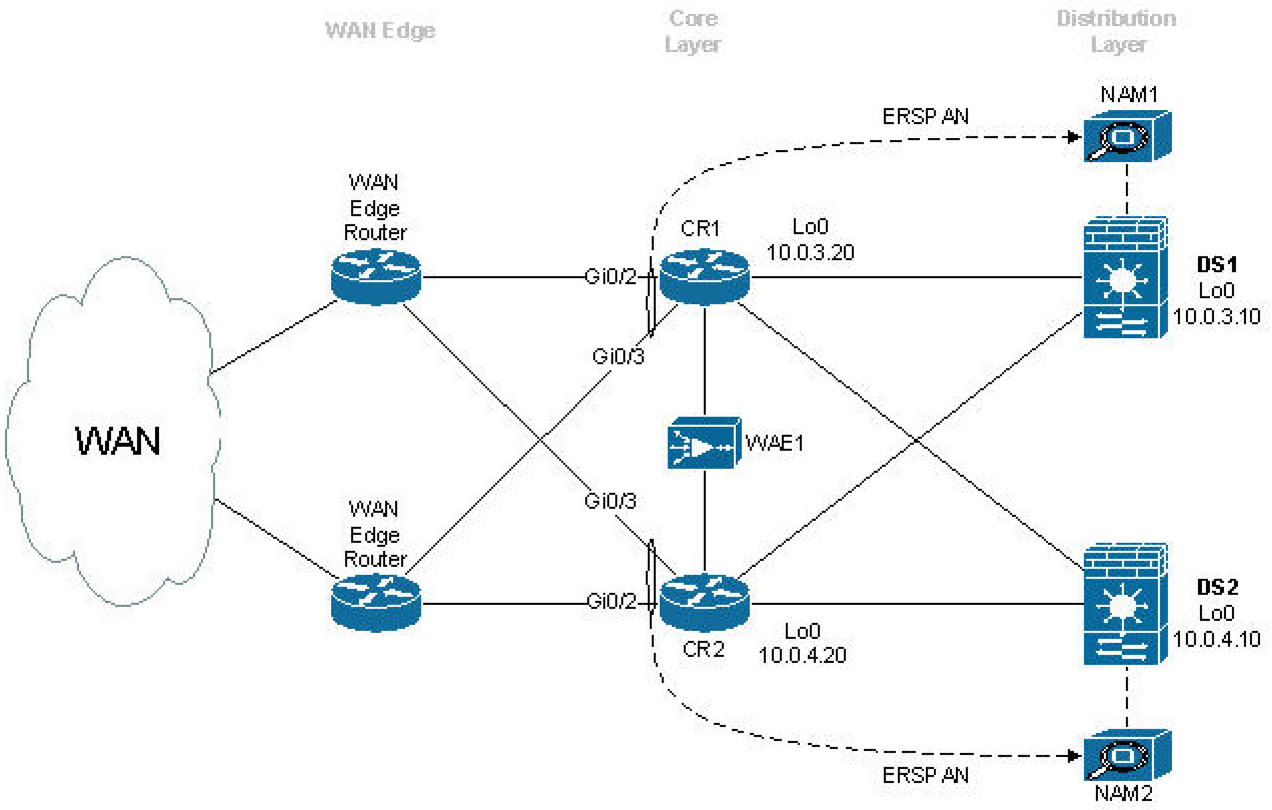

```
1. CONFIGURE SOURCE OF SPAN TRAFFIC ON CORE ROUTER CR1
monitor session 3 type erspan-source
 source interface Gi0/2
 source interface Gi0/3
 destination
```

```
   erspan-id 101
   ip address 10.0.3.10
   origin ip address 10.0.3.20
 !
```

After the ERSPAN traffic source is configured, ERSPAN traffic must be redirected to the NAM-2 on the distribution switch. Note that the same erspan-id is used at the distribution switch DS1.

```
2. CONFIGURE DESTINATION OF ERSPAN TRAFFIC ON DISTRIBUTION CAT6K SWITCH DS1
monitor session 3 type erspan-destination
 destination analysis-module 9 data-port 2
 source
  erspan-id 101
  ip address 10.0.3.10
!
```

A similar configuration can be applied to the second NAM-2 on core router CR2 and distribution switch DS2.

9.23.3 NAM-2 Deployment Caveats

- A NAM-2 data port cannot be used for two different SPAN sessions.
- A NAM-2 data port can be used as either a SPAN or VACL destination.
- NAM-2 can separate and report measurements by VLAN in the same SPAN session. However, NAM-2 cannot separate measurements by ports in the same SPAN session.
- NetFlow Data Export to NAM-2 does not support response time metrics. Only traffic metrics are available.

9.24 Use Case 1: Troubleshooting

This use case focuses on using NAM-2 in the data center to troubleshoot performance problems related to WAN and application optimization.

The following use case describes an implementation that was deployed in a test lab. The use case provides an example of how the WAN and application optimization solution might work in an real-world scenario. However, the use case is not intended to reflect actual performance or behavior in any environment.

9.24.1 Objectives

Use NAM-2 in the data center to troubleshoot performance problems related to WAN and application optimization.

9.24.2 Assumptions

1. WAAS and new QoS policies are deployed.
2. NAM-2 is deployed at the data center Catalyst 6000 Series distribution switch.
3. NAM-2 monitors the server segment using SPAN or VACL on DATA PORT 1.
4. NAM-2 monitors the WAN segment using SPAN on DATA PORT 2
5. NAM-2 monitors the WAN links of the Data Center using NetFlow Data Export.

The following example use case illustrates the NAM capabilities and features useful for troubleshooting application performance problems related to WAN and application optimization.

9.24.3 Use Case Example

A multinational enterprise recently consolidated IT services and applications into fewer data centers to reduce costs. To ensure that remote branch offices can use critical business applications with good performance, the company deployed WAAS to perform WAN and application optimization. After the WAAS deployment WAN traffic from the branch offices was reduced by more than 60%, and application performance is no longer affected by long latencies and limited bandwidth.

Two months after the deployment, IT Support received a call from a user in a remote branch office who complained about the slow response time of an important web-based business application. Because NAM was also deployed in the company data centers, IT Support has an invaluable tool for potential real-time traffic visibility and performance problem isolation and resolution.

The data center and the remote branch office are multihomed to two SP networks (SP-A and SP-B). The SP-A network is for business critical applications and has a strict SLA, while the SP-B network is for best-effort traffic with no SLA. Therefore, critical business applications run on the SP-A network, which the enterprise monitors using Cisco IP SLA.

9.24.4 Use Case Workflow

IT Support first identifies the client IP address (10.0.1.10), and then uses NAM-2 to troubleshoot the problem as follows.

Step 1. Use NAM to focus on user conversations and identify which application and server the user at the remote branch is using.

1. Go to Monitor > Conversations > Application Hosts.

2. Filter conversations using User IP Address in order to identify the application (HTTP) and the Server IP address (10.3.1.10).

The user application and server are identified.

Figure 9-86. Identifying User Conversations at the Remote Branch

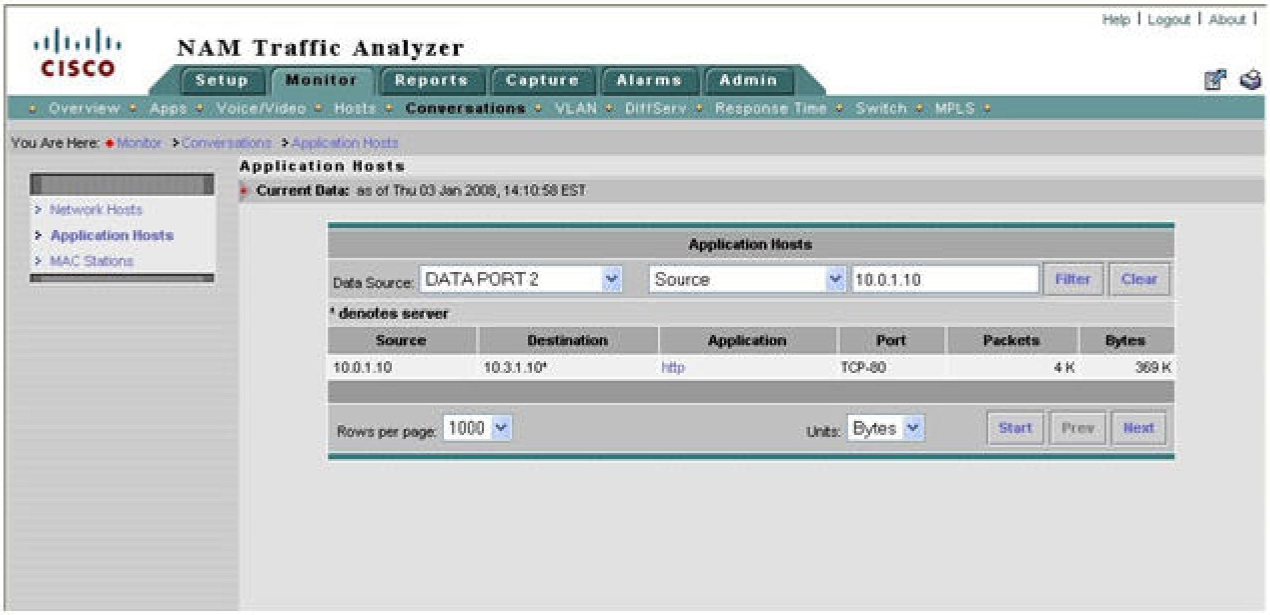

Application Hosts

Current Data: as of Thu 03 Jan 2008, 14:10:58 EST

Application Hosts

Data Source: DATA PORT 2 Source 10.0.1.10

* denotes server

Source	Destination	Application	Port	Packets	Bytes
10.0.1.10	10.3.1.10*	http	TCP-80	4 K	369 K

Rows per page: 1000 Units: Bytes

The next step is to isolate the problem. Is it in the server segment or in the WAN segment?

Step 2. [Server Segment] Check for conversation-specific application delay (that is, SRT) metrics in Real Time monitor mode to see if the server application is slow to respond.

1. Go to Monitor >Response Time > Server/Client Application.
2. Filter conversations using the user IP Address to identify the application (HTTP) and the server IP address (10.3.1.10).

Application delay as shown in Figure 9-87 appears normal, so server problems can be ruled out. The next step is to investigate network problems.

Figure 9-87. Checking Application Delay for a Specific Conversation

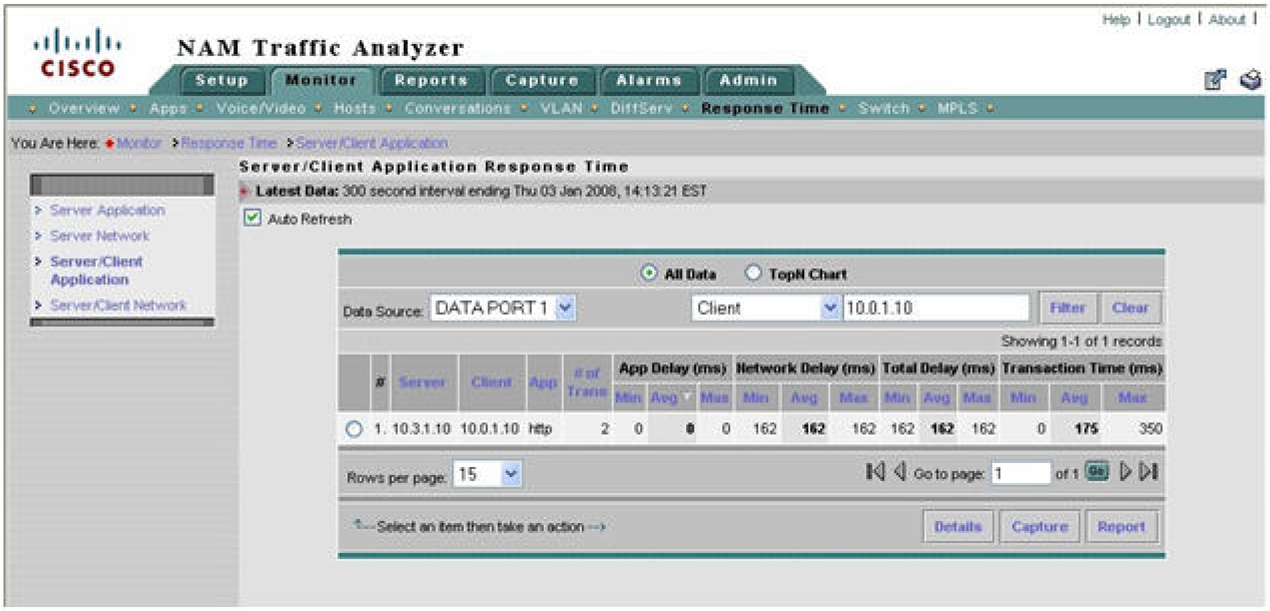

Server/Client Application Response Time

Latest Data: 300 second interval ending Thu 03 Jan 2008, 14:13:21 EST

☑ Auto Refresh

● All Data ○ TopN Chart

Data Source: DATA PORT 1 — Client 10.0.1.10 [Filter] [Clear]

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.	10.3.1.10	10.0.1.10	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 --> [Details] [Capture] [Report]

Step 3. [Server Segment] Check for conversation-specific network delays. Check server network delay. Are there long delays in the server side of the network? Check client network delay - Are there long delays in the client side of the network?

1. Go to Monitor > Response Time > Server/Client Network.
2. Focus on Server segment (DATA PORT 1).
3. Filter conversations using the user IP address (10.0.1.10).
4. Check SND and CND metrics

As shown in Figure 9-88, server network delay appears normal, but client network delay for the conversation exceeds the SLA for the SP-A network.

Figure 9-88. Check Network Delay for a Specific Conversation

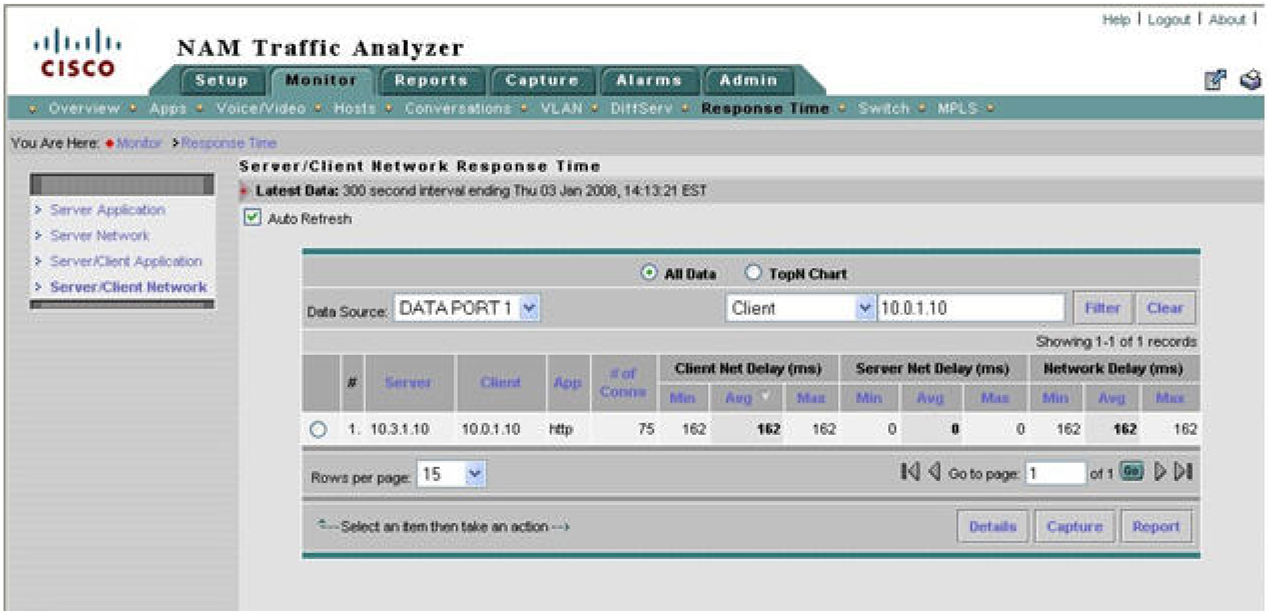

Server/Client Network Response Time

Latest Data: 300 second interval ending Thu 03 Jan 2008, 14:13:21 EST

☑ Auto Refresh

● All Data ○ TopN Chart

Data Source: DATA PORT 1 Client 10.0.1.10 Filter Clear

Showing 1-1 of 1 records

	#	Server	Client	App	# of Conns	Client Net Delay (ms)			Server Net Delay (ms)			Network Delay (ms)		
						Min	Avg	Max	Min	Avg	Max	Min	Avg	Max
○	1.	10.3.1.10	10.0.1.10	http	75	162	**162**	162	0	**0**	0	162	**162**	162

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

<--Select an item then take an action--> Details Capture Report

Now we must verify whether client network delay (CND) consistently exceeds the SLA. If so, there might be a routing problem or a problem with the SP. Otherwise, a congestion event might have delayed the packets for that conversation.

Step 4. [Server Segment] Generate reports to identify whether longer than expected CND is due to network congestion.

1. In the real-time window of Step 2, click a specific conversation and then click Report to create history reports for server network delay, CND, and application delay for the conversation. Figure 9-89 shows the result.

Figure 9-89. Create History Report for Specific Conversation

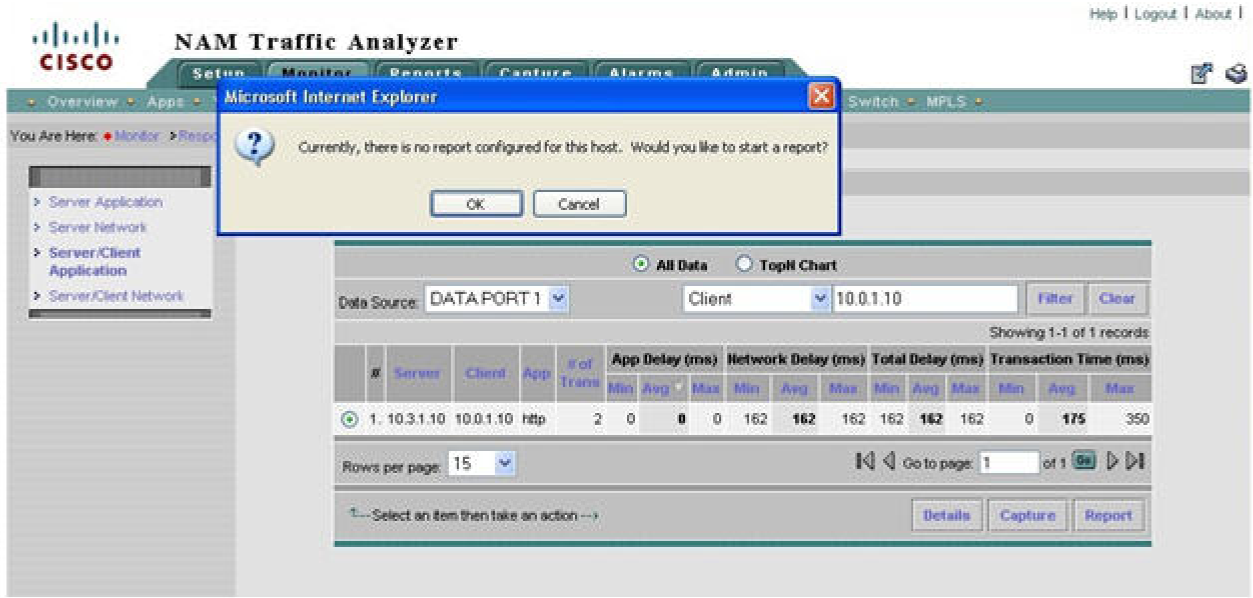

Currently, there is no report configured for this host. Would you like to start a report?

● All Data ○ TopN Chart

Data Source: DATA PORT 1 — Client 10.0.1.10

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.	10.3.1.10	10.0.1.10	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 -->

We must also check whether WAAS optimized the conversation.

Step 5. [Server and WAN Segment] Using real-time reports check whether WAAS is optimizing the conversation. Compare traffic volume in the WAN and server segments and check whether there is a reduction of traffic in the WAN segment.

1. Go to Monitor > Conversations > Network Hosts...
2. Change the data source to focus on the server segment (DATA PORT 1).
3. Use the user IP address (10.0.1.10) to filter conversations.
4. Click a specific conversation and then click Real-Time (do not close window).
5. In Monitor > Conversations > Network Hosts, change the data source to focus on the WAN segment (DATA PORT 2).
6. Use the user IP address (10.0.1.10) to filter conversations.
7. Click a specific conversation and then click Real-Time (do not close window).
8. Compare WAN segment and server segment traffic volumes in the real-time windows from the preceding steps. Check whether traffic is reduced on the WAN segment.

In this case, there is a reduction of bandwidth at the WAN segment, as shown in Figure 9-90, so it appears that WAAS is optimizing the conversation.

Figure 9-90. Checking whether WAAS Reduces WAN Traffic

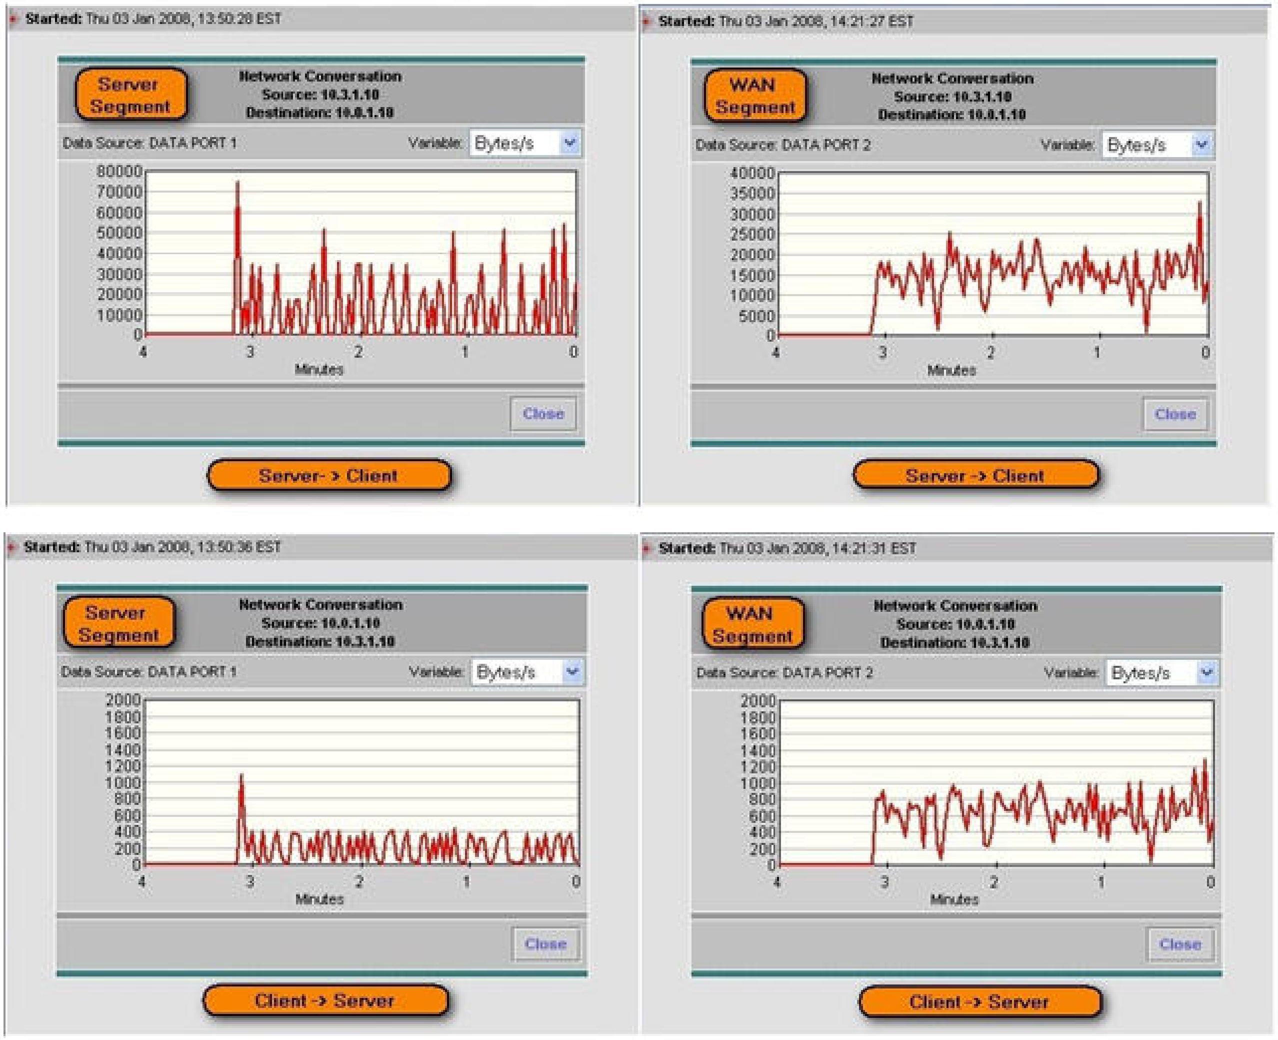

Started: Thu 03 Jan 2008, 13:50:28 EST

Network Conversation
Source: 10.3.1.10
Destination: 10.0.1.10

Data Source: DATA PORT 1 Variable: Bytes/s

Started: Thu 03 Jan 2008, 14:21:27 EST

Network Conversation
Source: 10.3.1.10
Destination: 10.0.1.10

Data Source: DATA PORT 2 Variable: Bytes/s

Started: Thu 03 Jan 2008, 13:50:36 EST

Network Conversation
Source: 10.0.1.10
Destination: 10.3.1.10

Data Source: DATA PORT 1 Variable: Bytes/s

Started: Thu 03 Jan 2008, 14:21:31 EST

Network Conversation
Source: 10.0.1.10
Destination: 10.3.1.10

Data Source: DATA PORT 2 Variable: Bytes/s

To confirm WAAS optimization, we can configure history reports for this conversation and determine whether it is being optimized by WAAS.

Step 6. [Server and WAN Segment] Configure reports to monitor WAAS traffic volume for the conversation.

1. Go to Monitor > Conversations > Network Hosts.
2. Change the data source to focus on the server segment (DATA PORT 1).
3. Use the user IP address (10.0.1.10) to filter conversations.
4. Click a specific conversation and then click Report.
5. In Monitor > Conversations > Network Hosts, change the data source to focus on the WAN Segment (DATA PORT 2).
6. Use the user IP address (10.0.1.10) to filter conversations.

7. Click a specific conversation and then click Report.

The reports generated in Step 6 are used to confirm the real-time view conclusions reached in Step 5.

Step 7. Ask the user to access the application a few more times.

The traffic generated by the user will be captured on historic reports set up in Steps 4 and 6.

Next, check for congestion on WAN links at the data center and branch site).

Step 8. [WAN Segment] While reports collect data, check whether WAN links at data center site are congested.

1. Go to Monitor > DiffServ > Traffic Stats.
2. Change the data source profile to focus on the SP-A data center WAN link Netflow Data Export (NDE-DC-WANLinkIn/NDE-DC-WANLinkOut).
3. Check traffic by aggregation class.

As shown in Figure 9-91, there is no congestion on the SP-A data center WAN link.

Figure 9-91. Checking for Congestion on the Data Center WAN Link

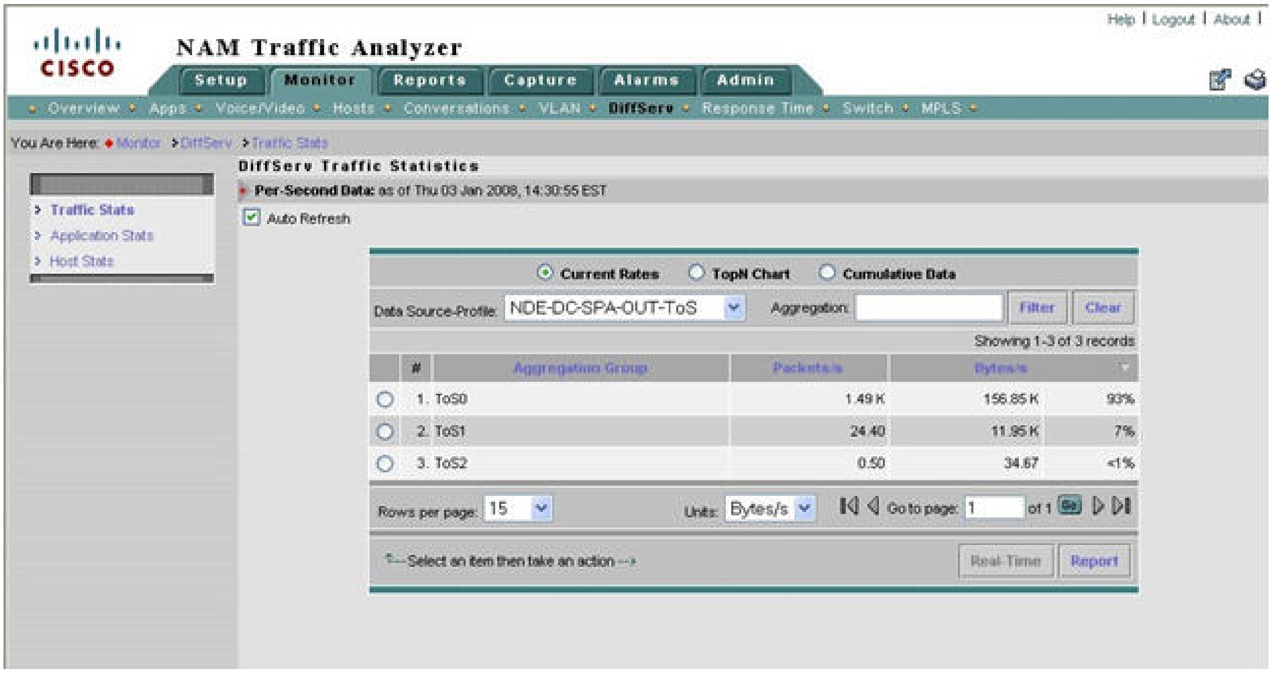

DiffServ Traffic Statistics

Per-Second Data: as of Thu 03 Jan 2008, 14:30:55 EST

☑ Auto Refresh

● Current Rates ○ TopN Chart ○ Cumulative Data

Data Source-Profile: NDE-DC-SPA-OUT-ToS

Aggregation:

Showing 1-3 of 3 records

	#	Aggregation Group	Packets/s	Bytes/s	
○	1.	ToS0	1.49 K	156.85 K	93%
○	2.	ToS1	24.40	11.95 K	7%
○	3.	ToS2	0.50	34.67	<1%

Rows per page: 15 Units: Bytes/s Go to page: 1 of 1

Select an item then take an action -->

Step 9. [WAN Segment] Check whether the SP-A WAN link at the branch office site is congested.

Either check the router using CLI, use a third-party tool, or configure NetFlow Export to NAM-2 at DC and follow these steps:

1. Go to Monitor > DiffServ > Traffic Stats.

2. Change the data source profile to focus on the branch office WAN link NetFlow Data Export (NDE-DC-WANLinkIn/NDE-DC-WANLinkOut).

3. Check traffic by aggregation class.

As shown in Figure 9-92, there is no congestion on the SP-A branch site WAN link.

Figure 9-92. Checking for Congestion at the Remote Site WAN Link

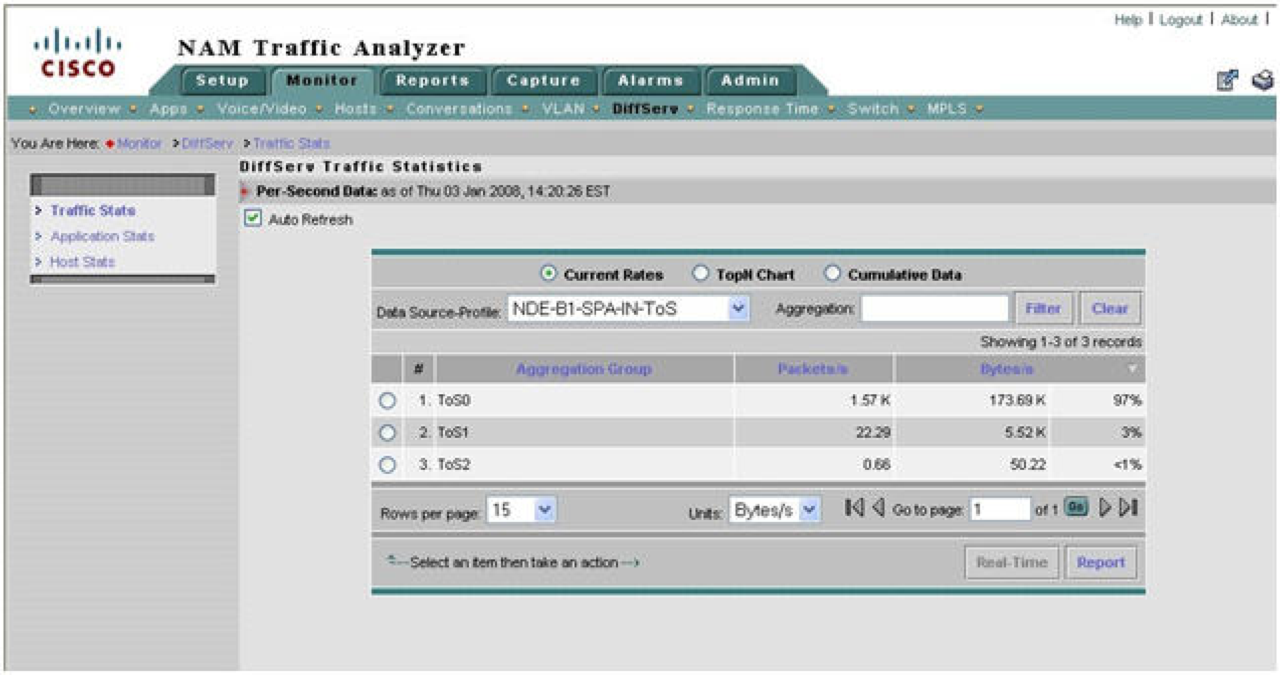

DiffServ Traffic Statistics

Per-Second Data: as of Thu 03 Jan 2008, 14:20:26 EST

☑ Auto Refresh

● Current Rates ○ TopN Chart ○ Cumulative Data

Data Source-Profile: NDE-B1-SPA-IN-ToS

Aggregation:

Showing 1-3 of 3 records

	#	Aggregation Group	Packets/s	Bytes/s	
○	1.	ToS0	1.57 K	173.69 K	97%
○	2.	ToS1	22.29	5.52 K	3%
○	3.	ToS2	0.66	50.22	<1%

Rows per page: 15 Units: Bytes/s Go to page: 1 of 1

Select an item then take an action

Review the history reports to confirm real-time view metrics.

Step 10. [Server Segment] Review the conversation-specific history reports that were set up in Steps 4 and 6. Compare Network Delay (SND and CND) with IP SLA reports for the SP-A network.

1. Go to Reports > Basic Reports.

2. Select Type: Response Type and Source (DATA PORT 1).

3. Use the user IP address (10.0.1.10) to filter conversations.

4. Click Application Delay, Server Network Delay, and Client Network Delay Reports for the user conversation.

Figure 9-93 shows that CND and SND (160ms) is much longer than the IP SLA reported value (80ms) for the SP-A network. This suggests that the conversation is probably routed onto the wrong WAN link.

Figure 9-93. Network Delay History Report for a Specific Conversation

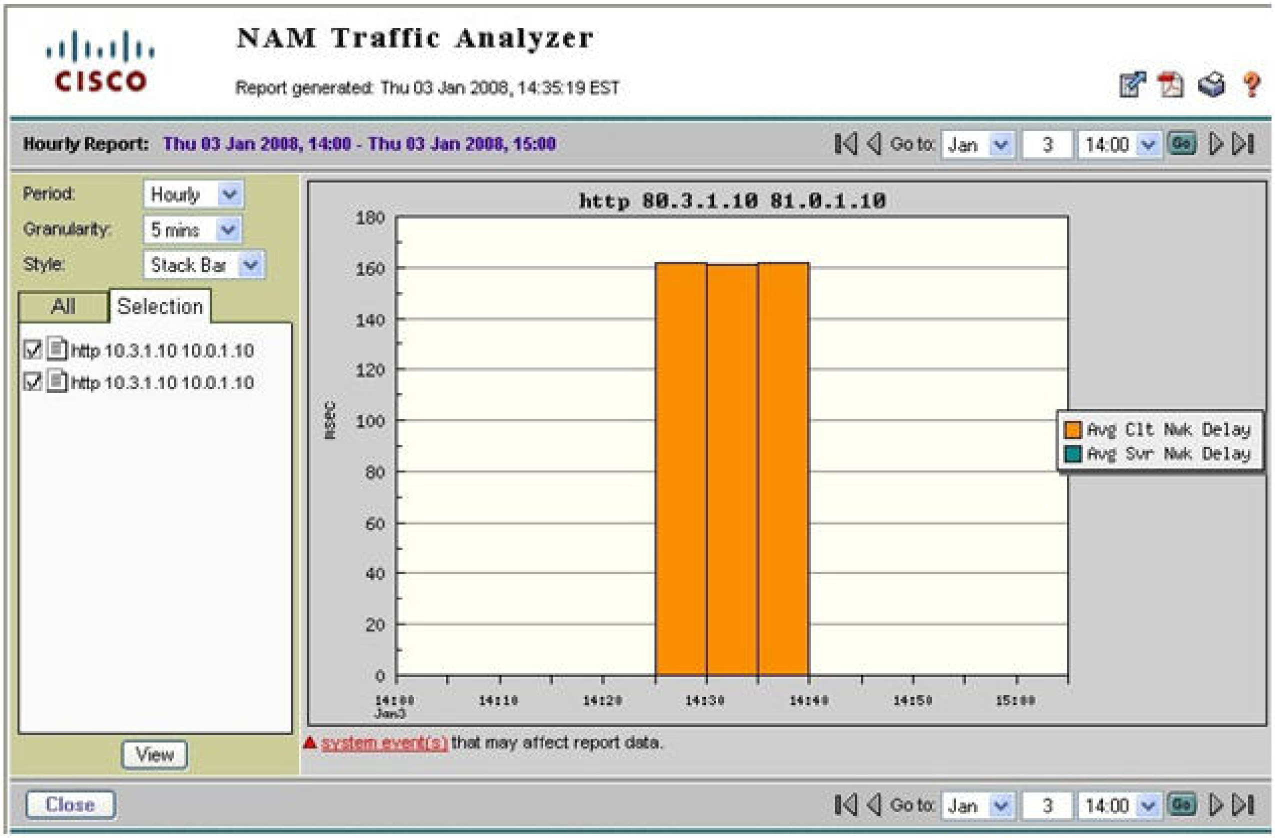

NAM Traffic Analyzer

Report generated: Thu 03 Jan 2008, 14:35:19 EST

Hourly Report: Thu 03 Jan 2008, 14:00 - Thu 03 Jan 2008, 15:00

Go to: Jan 3 14:00

Period: Hourly

Granularity: 5 mins

Style: Stack Bar

☑ http 10.3.1.10 10.0.1.10

☑ http 10.3.1.10 10.0.1.10

system event(s) that may affect report data.

Next, compare bandwidth for the specific conversation on the server and WAN segments.

1. Go to Reports > Basic Reports.
2. Select Type: Conversation and filter by Source: Server Segment (DATA PORT 1).
3. Click on reports for the user conversation.
4. Select Type: Conversation and filter by Source: WAN Segment (DATA PORT 2).
5. Click on reports for the user conversation.
6. Compare conversation reports for server and WAN segments.

The reports in Figure 9-94 and Figure 9-95 show that conversation traffic on the server segment (18KB/s) is significantly greater than traffic on the WAN segment (5KB/s). This confirms that WAAS is optimizing conversation traffic.

Figure 9-94. History Report for Server Segment Traffic

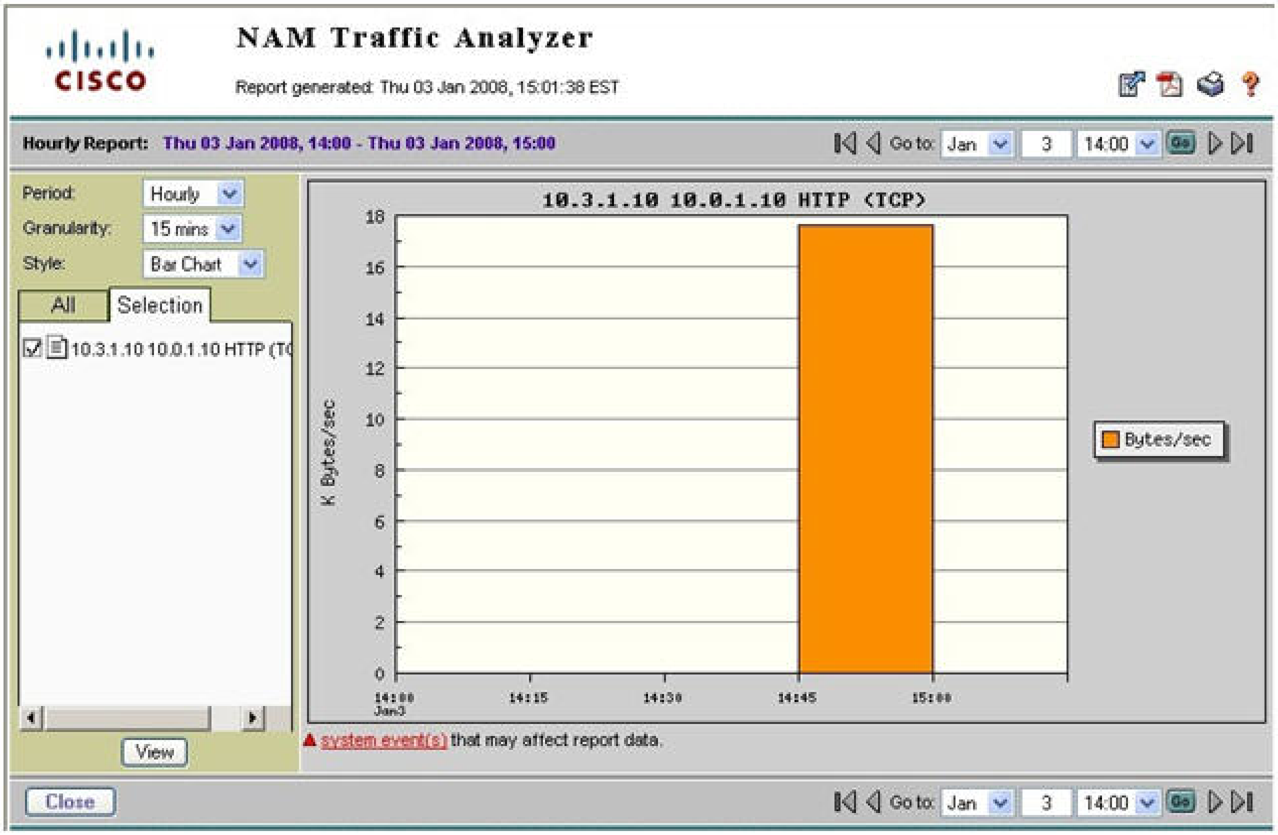

NAM Traffic Analyzer

Report generated: Thu 03 Jan 2008, 15:01:38 EST

Hourly Report: Thu 03 Jan 2008, 14:00 - Thu 03 Jan 2008, 15:00

Period: Hourly

Granularity: 15 mins

Style: Bar Chart

☑ 10.3.1.10 10.0.1.10 HTTP (TC

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

Figure 9-95. History Report for WAN Segment Traffic

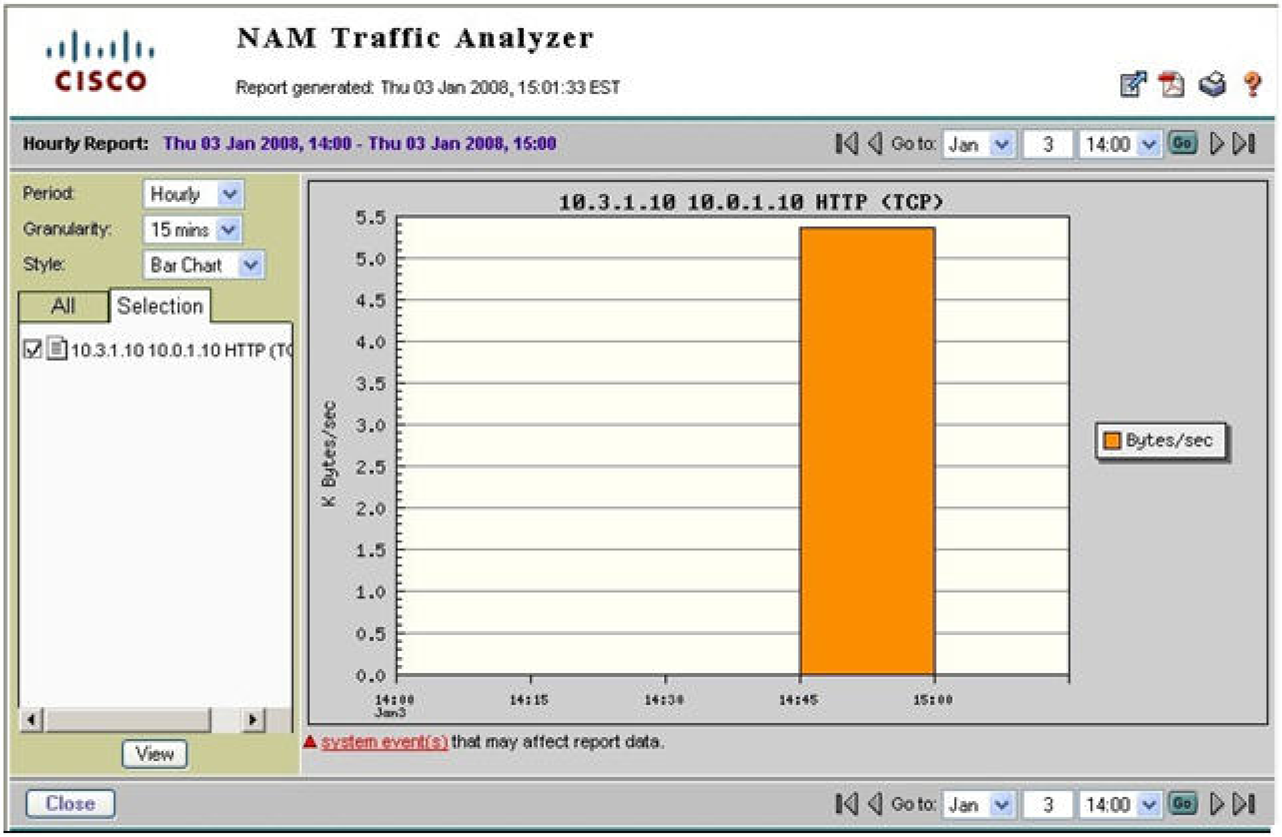

NAM Traffic Analyzer

Report generated: Thu 03 Jan 2008, 15:01:33 EST

Hourly Report: Thu 03 Jan 2008, 14:00 - Thu 03 Jan 2008, 15:00

Period: Hourly

Granularity: 15 mins

Style: Bar Chart

☑ 10.3.1.10 10.0.1.10 HTTP (T

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

Step 11. [WAN Segment] Check which WAN link is used at the branch site for this conversation.

1. Go to Monitor > Conversation > Application Hosts.
2. Change data source profile to focus on Data Center WAN links NetFlow Data Exports NDE-DC-WANLinkIn/NDE-DC-WANLinkOut/NDE-DC-SPB-IN/NDE-DC-SPB-OUT).
3. Filter by user IP Address (10.0.1.10).

The conversation appears on NDE-DC-SPB-IN and on NDE-DC-WANLinkOut., as shown in Figure 9-96.

Figure 9-96. Viewing Conversations on the Data Center WAN Link

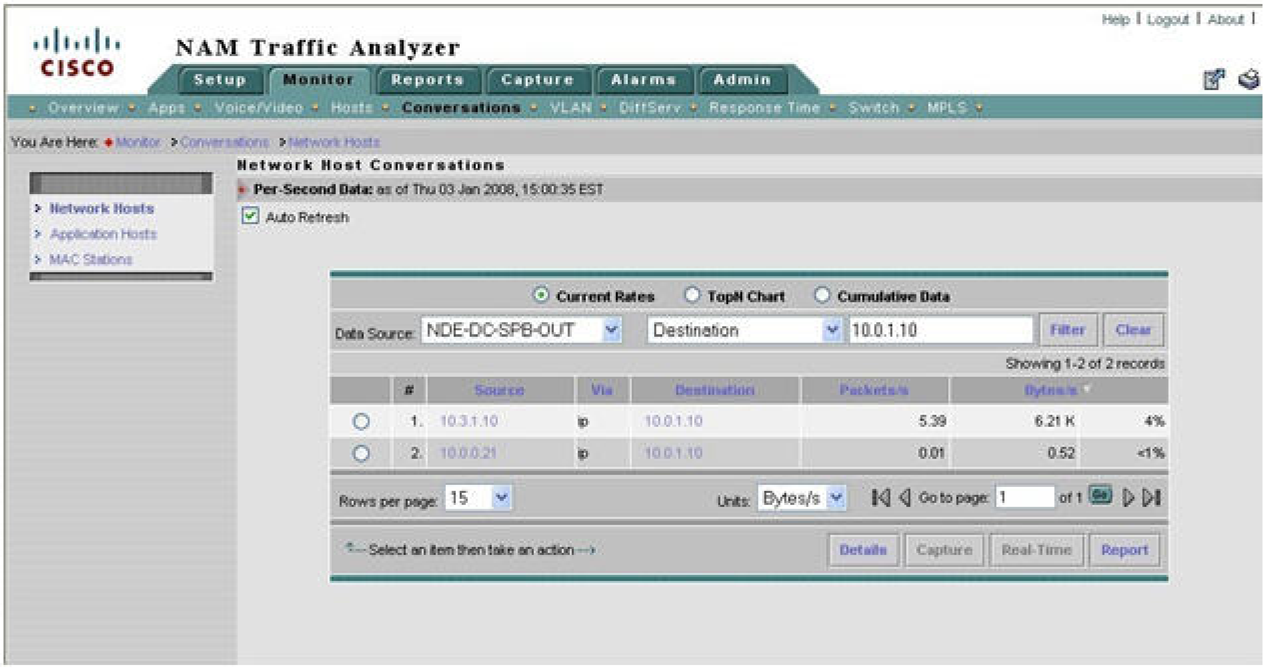

You Are Here: Monitor > Conversations > Network Hosts

Network Host Conversations

Per-Second Data: as of Thu 03 Jan 2008, 15:00:35 EST

☑ Auto Refresh

● Current Rates ○ TopN Chart ○ Cumulative Data

Data Source: NDE-DC-SPB-OUT | Destination | 10.0.1.10

Showing 1-2 of 2 records

	#	Source	Via	Destination	Packets/s	Bytes/s	
○	1.	10.3.1.10	ip	10.0.1.10	5.39	6.21 K	4%
○	2.	10.0.0.21	ip	10.0.1.10	0.01	0.52	<1%

Rows per page: 15 Units: Bytes/s Go to page: 1 of 1

Select an item then take an action

We established that the data center-to-branch traffic for this conversation is using SP-B instead of SP-A, as it is supposed to. The problem is traced to a typo in the load-balancing policy configuration of the multihomed data center router.

Here are the resolution actions:

- **Short term:** Correct the load-balancing policy at the data center router and ensure that all business-critical applications are routed through SP-A.
- **Long term:** Recommend using PfR for dynamic load balancing policy to ensure that business-critical applications are forwarded on best path.

9.25 Use Case 2: Conversation Analysis

This use case describes how to use NAM traffic volume, response time history reports, real-time monitoring to decide where to deploy WAAS and QoS and to validate WAN and application optimization benefits.

The use case describes an implementation that was deployed in a test lab. The use case provides an example of how the WAN and application optimization solution might work in an real-world scenario. However, the use case is not intended to reflect actual performance or behavior in any environment.

9.25.1 Objectives

Use NAM traffic volume and response time history reports and real-time monitoring to:

- Make decisions where to deploy WAAS and QoS policies (before WAN and application optimization deployment – Scenario A)
- Validate WAN and application optimization benefits (during deployment – Scenario B)

NAM TopN reports are used to target specific conversations and generate history reports for the required analysis.

Scenario A assumes that WAAS is not deployed yet while Scenario B assumes that WAAS is deployed.

9.25.2 Assumptions

1. NAM-2 is deployed at the data center Catalyst 6000 Series distribution switch.
2. NAM-2 uses SPAN or VACL to monitor the server segment on DATA PORT 1.
3. NAM-2 uses SPAN to monitor the WAN segment on DATA PORT 2

9.25.3 Use Case Example

Scenario A: A multinational enterprise is thinking about deploying WAN and application optimization and needs to identify the sites that would benefit most from WAN and application optimization.

Scenario B: The multinational enterprise in Scenario A started a pilot WAN and application optimization deployment. During the pilot, WAAS was deployed at 09:45 AM on November 15 to optimize a remote site that was experiencing poor performance. The customer uses NAM to demonstrate that WAAS improves application service delivery, thereby justifying WAAS deployments in other remote sites.

9.25.4 Use Case Workflow

9.25.4.1 Scenario A

Step 1 [Scenario A] Create a history report for top applications. This will identify the top applications running on the network. Even though this does not provide a breakdown per branch site, one can see if there is scope for improvement using WAN and application optimization.

Figure 9-97. Top Applications

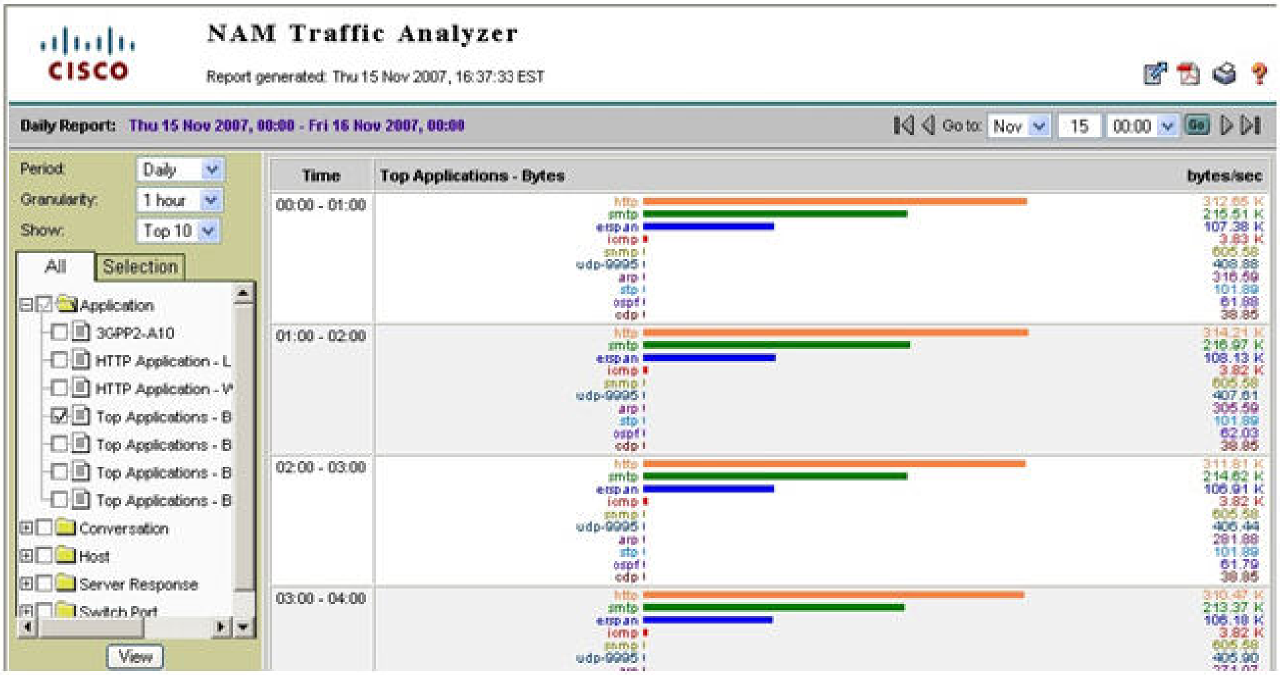

NAM Traffic Analyzer

Report generated: Thu 15 Nov 2007, 16:37:33 EST

Daily Report: Thu 15 Nov 2007, 00:00 - Fri 16 Nov 2007, 00:00

Go to: Nov 15 00:00

Period: Daily

Granularity: 1 hour

Show: Top 10

Time	Top Applications - Bytes	bytes/sec
00:00 - 01:00	http	312.65 K
	smtp	215.51 K
	erspan	107.38 K
	icmp	3.83 K
	snmp	605.58
	udp-9995	408.88
	arp	316.59
	stp	101.89
	ospf	61.88
	cdp	38.85
01:00 - 02:00	http	314.21 K
	smtp	216.97 K
	erspan	108.13 K
	icmp	3.82 K
	snmp	605.58
	udp-9995	407.61
	arp	305.59
	stp	101.89
	ospf	62.03
	cdp	38.85
02:00 - 03:00	http	311.81 K
	smtp	214.62 K
	erspan	106.91 K
	icmp	3.82 K
	snmp	605.58
	udp-9995	406.44
	arp	281.88
	stp	101.89
	ospf	61.79
	cdp	38.85
03:00 - 04:00	http	310.47 K
	smtp	213.37 K
	erspan	106.18 K
	icmp	3.82 K
	snmp	605.58
	udp-9995	405.90

Step 2 [Scenario A] Create a history report for top conversations (including application layer – see). This will help to identify top application conversations and show the volume of traffic (see Figure 9-99). Top application conversations will show up to the top 50 conversations. Based on this information, the user can deduce the most active branch sites and applications in terms of volume.

Figure 9-98. Conversation Report Creation Dialog

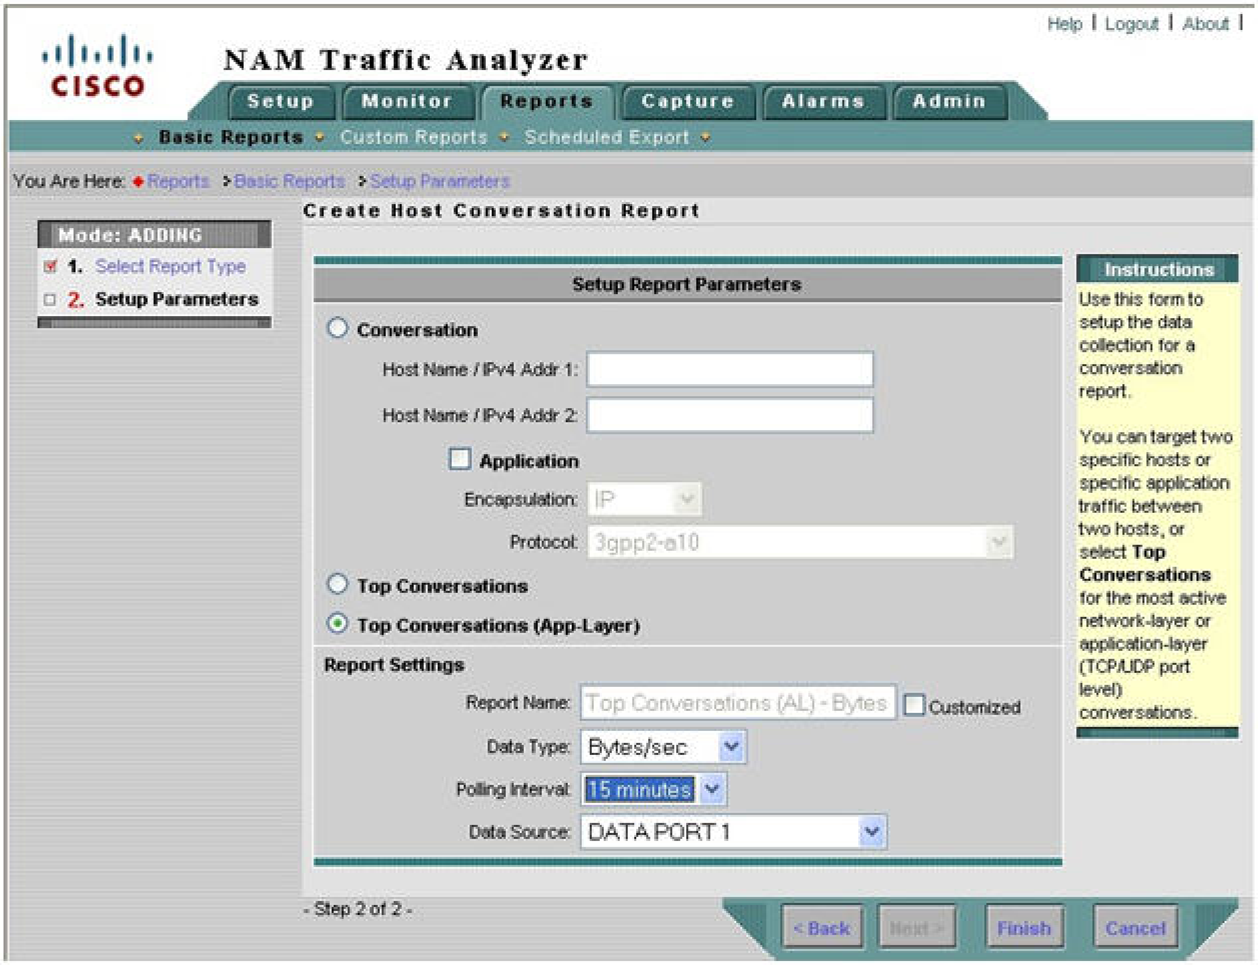

Create Host Conversation Report

Setup Report Parameters

○ Conversation

Host Name / IPv4 Addr 1:

Host Name / IPv4 Addr 2:

□ Application

Encapsulation: IP

Protocol: 3gpp2-a10

○ Top Conversations

● Top Conversations (App-Layer)

Report Settings

Report Name: Top Conversations (AL) - Bytes □ Customized

Data Type: Bytes/sec

Polling Interval: 15 minutes

Data Source: DATA PORT 1

- Step 2 of 2 -

Instructions

Use this form to setup the data collection for a conversation report.

You can target two specific hosts or specific application traffic between two hosts, or select **Top Conversations** for the most active network-layer or application-layer (TCP/UDP port level) conversations.

Figure 9-99. Top Conversations

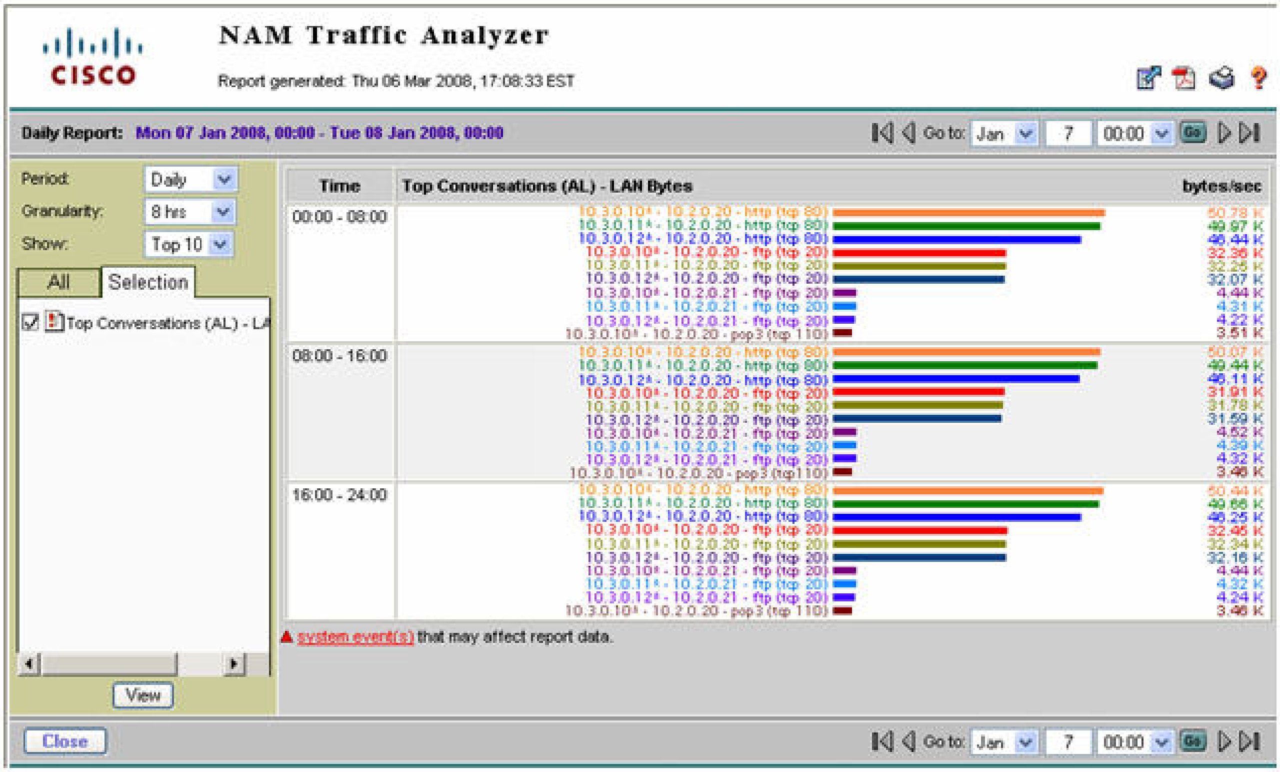

NAM Traffic Analyzer

Report generated: Thu 06 Mar 2008, 17:08:33 EST

Daily Report: Mon 07 Jan 2008, 00:00 - Tue 08 Jan 2008, 00:00

Go to: Jan 7 00:00

Period: Daily

Granularity: 8 hrs

Show: Top 10

All | Selection

☑ Top Conversations (AL) - LA

Time	Top Conversations (AL) - LAN Bytes	bytes/sec
00:00 - 08:00	10.3.0.10 - 10.2.0.20 - http (tcp 80)	50.78 K
	10.3.0.11 - 10.2.0.20 - http (tcp 80)	49.97 K
	10.3.0.12 - 10.2.0.20 - http (tcp 80)	46.44 K
	10.3.0.10 - 10.2.0.20 - ftp (tcp 20)	32.36 K
	10.3.0.11 - 10.2.0.20 - ftp (tcp 20)	32.25 K
	10.3.0.12 - 10.2.0.20 - ftp (tcp 20)	32.07 K
	10.3.0.10 - 10.2.0.21 - ftp (tcp 20)	4.44 K
	10.3.0.11 - 10.2.0.21 - ftp (tcp 20)	4.31 K
	10.3.0.12 - 10.2.0.21 - ftp (tcp 20)	4.22 K
	10.3.0.10 - 10.2.0.20 - pop3 (tcp 110)	3.51 K
08:00 - 16:00	10.3.0.10 - 10.2.0.20 - http (tcp 80)	50.07 K
	10.3.0.11 - 10.2.0.20 - http (tcp 80)	49.44 K
	10.3.0.12 - 10.2.0.20 - http (tcp 80)	46.11 K
	10.3.0.10 - 10.2.0.20 - ftp (tcp 20)	31.91 K
	10.3.0.11 - 10.2.0.20 - ftp (tcp 20)	31.78 K
	10.3.0.12 - 10.2.0.20 - ftp (tcp 20)	31.59 K
	10.3.0.10 - 10.2.0.21 - ftp (tcp 20)	4.52 K
	10.3.0.11 - 10.2.0.21 - ftp (tcp 20)	4.39 K
	10.3.0.12 - 10.2.0.21 - ftp (tcp 20)	4.32 K
	10.3.0.10 - 10.2.0.20 - pop3 (tcp110)	3.46 K
16:00 - 24:00	10.3.0.10 - 10.2.0.20 - http (tcp 80)	50.44 K
	10.3.0.11 - 10.2.0.20 - http (tcp 80)	49.66 K
	10.3.0.12 - 10.2.0.20 - http (tcp 80)	46.25 K
	10.3.0.10 - 10.2.0.20 - ftp (tcp 20)	32.45 K
	10.3.0.11 - 10.2.0.20 - ftp (tcp 20)	32.34 K
	10.3.0.12 - 10.2.0.20 - ftp (tcp 20)	32.16 K
	10.3.0.10 - 10.2.0.21 - ftp (tcp 20)	4.44 K
	10.3.0.11 - 10.2.0.21 - ftp (tcp 20)	4.32 K
	10.3.0.12 - 10.2.0.21 - ftp (tcp 20)	4.24 K
	10.3.0.10 - 10.2.0.20 - pop3 (tcp 110)	3.46 K

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

Step 3 [Scenario A] Set the monitoring for response time to 24hrs or 7 days and use real-time views to show TopN ART metrics and identify conversations with problems.

Figure 9-100. TopN Average and Maximum Transaction Time Conversations

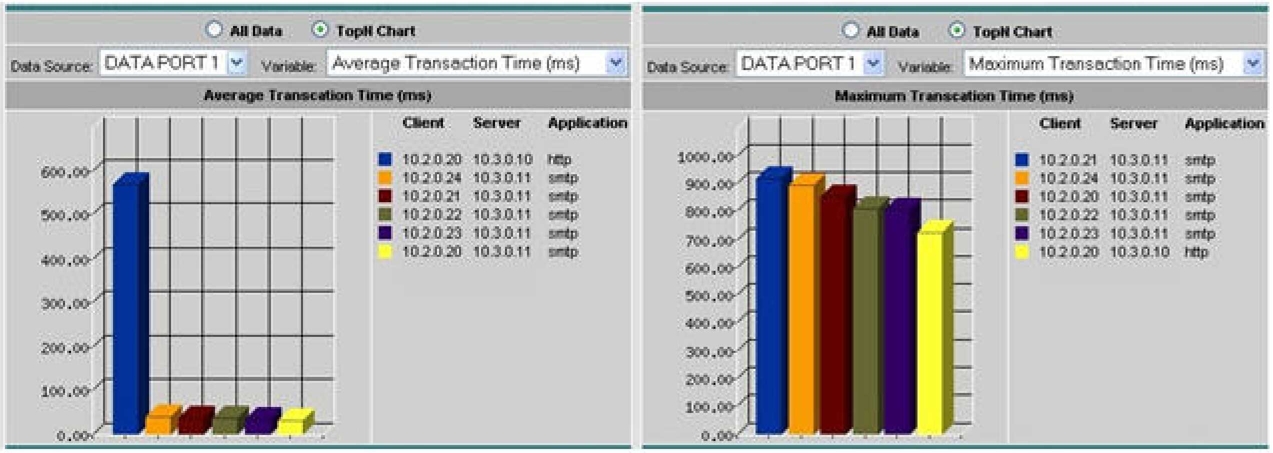

○ All Data ● TopN Chart

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

Average Transcation Time (ms)

	Client	Server	Application
🟦	10.2.0.20	10.3.0.10	http
🟧	10.2.0.24	10.3.0.11	smtp
🟥	10.2.0.21	10.3.0.11	smtp
🟩	10.2.0.22	10.3.0.11	smtp
🟪	10.2.0.23	10.3.0.11	smtp
🟨	10.2.0.20	10.3.0.11	smtp

○ All Data ● TopN Chart

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

Maximum Transaction Time (ms)

	Client	Server	Application
🟦	10.2.0.21	10.3.0.11	smtp
🟧	10.2.0.24	10.3.0.11	smtp
🟥	10.2.0.20	10.3.0.11	smtp
🟩	10.2.0.22	10.3.0.11	smtp
🟪	10.2.0.23	10.3.0.11	smtp
🟨	10.2.0.20	10.3.0.10	http

For example, Figure 9-100 shows that the conversations with the highest average transaction time are all between networks 10.2.0.0/8 and servers 10.3.0.11 and 10.3.0.10.

Step 4 [Scenario A] Set up response time history reports for conversations identified by Steps 2 and 3.

Figure 9-101. Average Transaction Time Historical Report

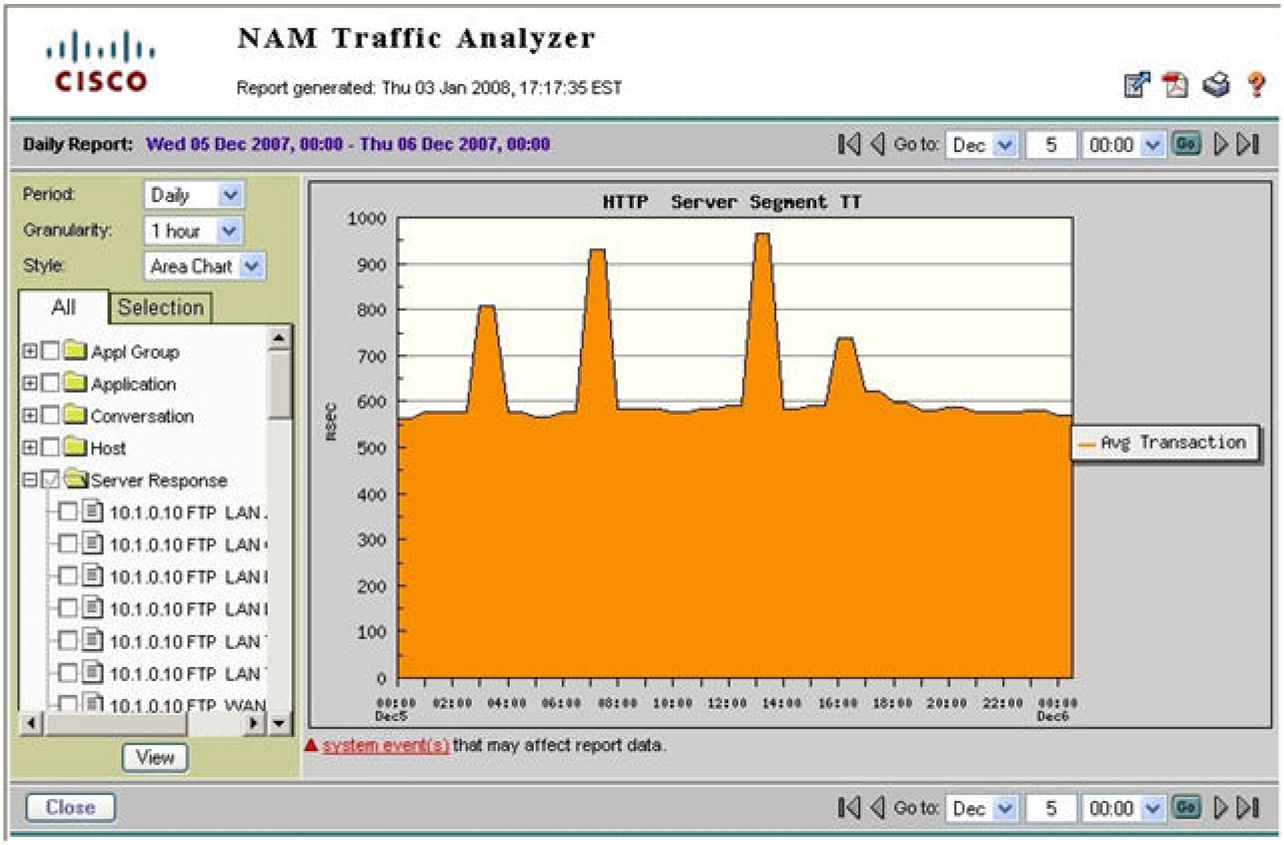

NAM Traffic Analyzer

Report generated: Thu 03 Jan 2008, 17:17:35 EST

Daily Report: Wed 05 Dec 2007, 00:00 - Thu 06 Dec 2007, 00:00

Go to: Dec 5 00:00 Go

Period: Daily

Granularity: 1 hour

Style: Area Chart

All | Selection

- ⊞ □ Appl Group
- ⊞ □ Application
- ⊞ □ Conversation
- ⊞ □ Host
- ⊟ ☑ Server Response
  - □ 10.1.0.10 FTP LAN
  - □ 10.1.0.10 FTP LAN
  - □ 10.1.0.10 FTP LAN
  - □ 10.1.0.10 FTP LAN
  - □ 10.1.0.10 FTP LAN
  - □ 10.1.0.10 FTP LAN
  - □ 10.1.0.10 FTP WAN

View

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

Close

Go to: Dec 5 00:00 Go

Step 5 [Scenario A] Set up conversation volume history reports for conversations identified by Steps 2 and 3 monitoring the server segment (DATA PORT 1).

After setting up the reports, data is collected for a specific period (e.g. 1 week) before it is examined by the IT department. Based on the generated reports, applications and branch sites that experience long response times can be identified and a decision can be reached where to deploy WAAS.

9.25.4.2 Scenario B

During WAAS deployment, the benefits of WAAS can be validated by NAM as follows:

Before WAAS is turned on:

Step 1 [Scenario B] Set up conversation volume history reports for conversations in step 5 but this time monitoring the WAN segment (DATA PORT 2).

Both sides of WAAS are now being monitored by NAM-2.

After WAAS is turned on:

Step 2 [Scenario B] Review response time history reports for conversations set up in step 4. Applications optimized by WAAS should experience improved response times (See Figure 9-102).

Figure 9-102. Conversation Transaction Time Before and After WAAS

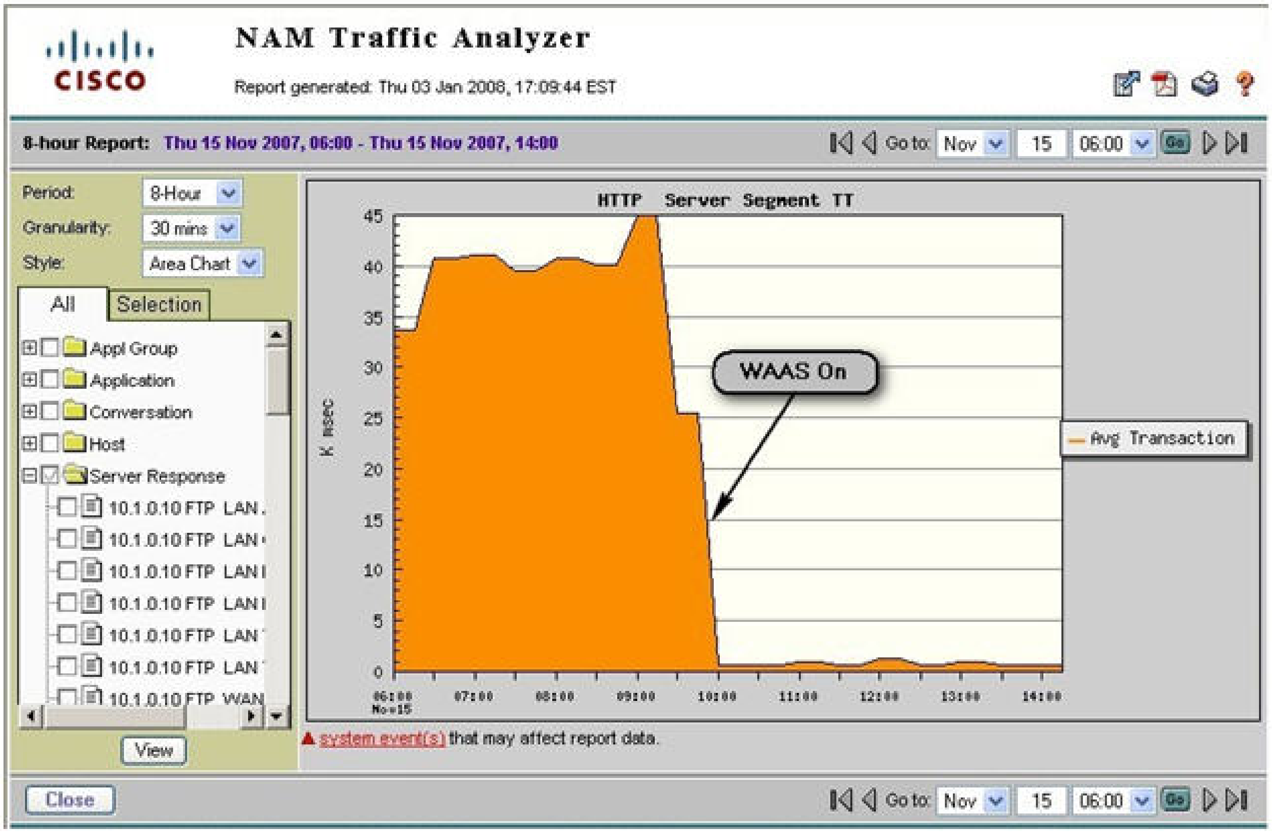

NAM Traffic Analyzer

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

Go to: Nov 15 06:00 Go

Period: 8-Hour

Granularity: 30 mins

Style: Area Chart

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

Step 3 [Scenario B] Review conversation volume history reports for conversations set up in step 5 and 6. Applications optimized by WAAS should experience reduced traffic volume on the WAN Segment (see Figure 9-103).

Figure 9-103. WAN Segment Conversation Traffic Volume

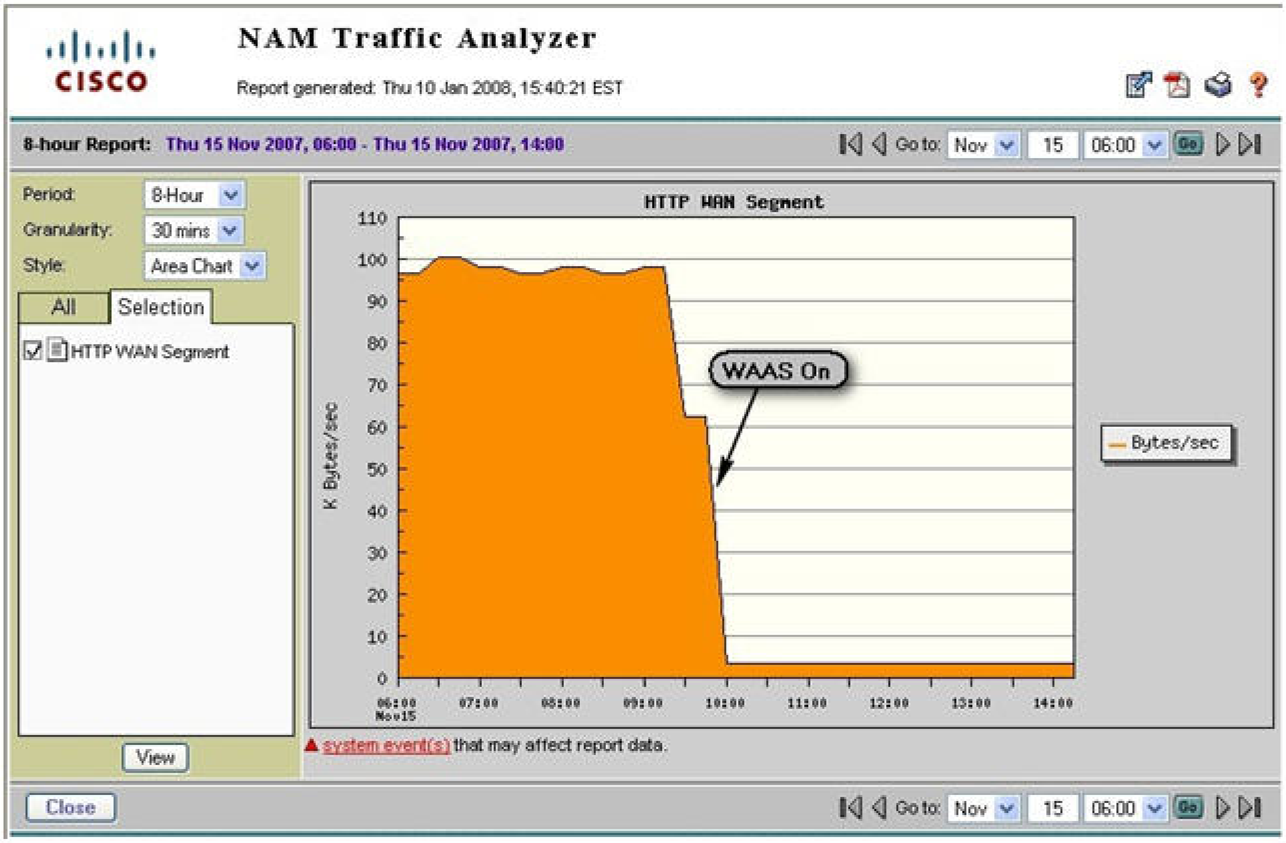

NAM Traffic Analyzer

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

Go to: Nov 15 06:00

Period: 8-Hour

Granularity: 30 mins

Style: Area Chart

☑ HTTP WAN Segment

system event(s) that may affect report data.

9.25.5 Deployment Caveats

The following deployment caveats apply to NetQoS and Cisco NAM.

- If server load balancing (SLB) or SSL offloading takes place at the distribution layer (for example, using ACE), the traffic seen at the WAN Segment will show the virtual IP address (VIP) of the SLB/SSL offloader. On the other hand, traffic in the server segment shows the real IP address of the server being used. This makes it difficult to correlate traffic on the WAN segment with traffic on the server segment.
- If a FWSM is present at the distribution layer, the WAN segment accounts DC incoming traffic before the FWSM blocks it. Similarly, the LAN segment accounts DC outgoing traffic before the FWSM blocks it. This can result in discrepancies in the volume of traffic in the LAN and WAN segments.
- The Catalyst 6000 Series has a limit of two SPAN sessions. However, each session can have multiple sources and destinations. For example, both NetQoS and NAM-2 can share a single SPAN session on the server segment.

On PFC3, IOS Release 12.2(18) SXE and later releases support 128 sources in each local SPAN/RSPAN/ERSPAN session; 64 destinations per session are supported for local SPAN. See Configuring Local SPAN, Remote SPAN (RSPAN), and Encapsulated RSPAN section of Catalyst 6500 Release 12.2SXF Software Configuration Guide.

- A SPAN session can have VLANs or ports as data sources, but not both in the same session. Therefore, one SPAN session cannot monitor both the WAN and server segments. An additional SPAN or VACL session is required.
- Only two NetFlow Data Export destinations are supported. If NetFlow data is exported to NAM-2 and NetQoS Reporter Analyzer, NetFlow data cannot be exported to other destinations. To work around this limitation, use of a flow replicator.
- NAM-2 and NetQoS have a 1Gb/s capacity limitation.