

Case Study

Transforming Network Capacity Planning from an Art to a Science

Cisco IT Case Study / Network Capacity Planning / Cisco IOS NetFlow: This case study describes Cisco IT's network capacity planning process and its internal deployment of Cisco IOS NetFlow and third-party solutions within the Cisco network, a leading-edge enterprise environment that is one of the largest and most complex in the world. Cisco customers can draw on Cisco IT's real-world experience in this area to help support similar enterprise needs.

"Having tools that allow us to identify the applications consuming bandwidth is absolutely indispensable. A granular view of what's happening in our network has allowed us to forecast our need for additional WAN links and budget effectively several quarters out."

-- Joe Silver, Cisco IT Project Manager

Background

Of all the issues faced by enterprise companies in managing their networks, capacity planning is one of the most important. More an art than a science until recently, network capacity planning is all about balancing the need to meet user performance expectations against the realities of capital budgeting.

WAN bandwidth is expensive. Many companies—and Cisco Systems® is no exception—attempt to control costs by acquiring the minimum bandwidth necessary to handle traffic on a circuit. Unfortunately, this strategy can lead to congestion and degraded application performance.

A WAN circuit running at 80 percent of capacity is too full. Even a circuit that averages 60 percent of capacity may well peak at 95 percent of capacity for several periods during a day, reducing user productivity and negatively affecting business activities. Many IT organizations order new circuits (which can take anywhere from 30 to 90 days to deploy) when a circuit operates at 60 to 80 percent of capacity.

As recently as 2000, Cisco® relied almost exclusively on Simple Network Management Protocol (SNMP) to monitor overall WAN bandwidth utilization. Measuring overall traffic, however, does little to characterize network traffic, which is essential to deciding if additional capacity is warranted. Without knowing what types of traffic are using the network, it is impossible to know if quality of service (QoS) parameters for applications such as voice or video support target service levels. Complicating the challenges of traffic characterization is the reality that many new applications use a range of dynamic ports. These dynamic ports may also be used by several different applications within the enterprise.



Challenge

Through the late 1990s, Cisco operated only 140 Frame-Relay-based WAN sites in the United States. Bandwidth capacity was sub T-1. However, in 2000, WAN utilization began to increase rapidly, doubling every 12 to 18 months, degrading application performance, and affecting business operations.

Driving bandwidth consumption were voice over IP (VoIP, or Internet telephony) and video on demand (VoD), which share the network with more conventional uses, including e-mail, Internet access, and PC backups. Frequently, Cisco IT engineers found that traffic congestion on some network links had significantly reduced user productivity.

Though IT knew that traffic was increasing exponentially and that actual usage was not in line with expectations, it did not have access to the level of detail necessary to understand the true nature of problem. This made it almost impossible to make informed decisions about bandwidth upgrades.


“Cisco had no clearly established proactive capacity planning process,” says Keith Brumbaugh, Cisco IT Global Network Engineer. “We tended to implement upgrades in reaction to internal customer complaints rather than solid data. And in the early 2000s we were finding ourselves overwhelmed by the traffic from new applications such as voice and video. We were particularly concerned, because VoIP is sensitive to latency and jitter. Poor voice performance in our environment could detract from its business value.”

From experience, the network capacity planning team knew that a few applications can consume most of the WAN bandwidth on a given network segment. Further, the team knew that with visibility into the top 10 applications, along with the top 10 traffic pairs, it was possible to accurately identify and characterize 70 percent or more of network traffic. Any one of these top 10 applications can occupy 10 percent or more of a segment’s bandwidth—and when analysis extends beyond these applications, it shows that consumption quickly fades. In fact, any application not in the top 10 probably uses less than one percent of available bandwidth. “Almost half of Cisco IT’s WAN backbone traffic consists of data backup such as database syncs, server-to-server, PC backup, and SnapMirror, which is used to back up engineering data,” says Brumbaugh.

QoS: An Aside

A surprising number of IT staff at large enterprises believe that networks built with high-capacity switches, multi-gigabit backplanes, and high-speed LAN and WAN links should never need QoS management. They believe that the more bandwidth available, the less QoS is needed. All networks have congestion points where data packets can be dropped—WAN links where a larger trunk funnels data into a smaller one, or a location where several trunks funnel data into fewer outputs (Figure 1). Applying QoS does not create additional bandwidth. Rather, it helps smooth the peaks and valleys of network circuit utilization. QoS provides more consistent network performance from the point of view of users.

Figure 1. Congestion Point Examples



From a capacity planning standpoint, deploying QoS uniformly across the network protects important real-time voice and video applications—guaranteeing bandwidth and/or low latency—from occasional traffic spikes that can affect performance. Because of this measure of protection, Cisco IT planners believe that QoS settings must be deployed globally on all appropriate network devices in order for capacity planning to be fully effective.

Cisco planners also discovered that while most capacity planning occurs at the circuit level, it is also desirable where possible to plan within individual classes of service. It is possible to oversubscribe one or more classes of service without reaching utilization levels that would affect a circuit's overall performance. It is especially important to do this type of planning when using WAN technologies such as Multiprotocol Label Switching (MPLS), virtual private networks (VPNs), or Asynchronous Transfer Mode (ATM). Carriers, in addition to charging for a circuit, also charge for these classes of service. Managing the bandwidth levels of these individual classes of service ensures proper application performance without overspending for a particular class of service.

Solution

Categorizing Network Traffic

Cisco IT began its efforts toward improving the capacity planning process by categorizing network traffic into three types:

- **Legitimate, business-related traffic** – Companies build their networks to accommodate legitimate, business-related traffic. If a link is at capacity and all traffic is legitimate, then a network upgrade may be necessary. A factor influencing this decision is that some legitimate traffic, such as backups, file transfers, or VoD replication, can be scheduled outside of peak utilization hours. The ability to make scheduling changes can often postpone the need for an upgrade. “When we first implemented a new application to back up user PC hard drives across the network, we seriously underestimated the impact it would have on the WAN—especially smaller branch WAN links,” says Joe Silver, Cisco IT Project Manager. “While backups were done incrementally, the initial backup was always large—and when we first deployed the application, they were all initial backups. After looking at the performance problems, and realizing they were created by a legitimate application that would eventually stop needing so much bandwidth, we decided to avoid WAN upgrades. Instead, we asked



the application developers to schedule all initial backups after hours. When they did, the problem was solved.”

- **Inappropriate traffic** – Traffic in this category can include everything from recreational Internet downloads to viruses, worms, or distributed denial of service (DDoS) attacks. Capacity planners have discovered that it is not important or even desirable to eliminate recreational traffic entirely, until it begins to significantly affect bandwidth availability and compete with the top 10 applications. Investigating and eliminating inappropriate traffic may postpone the need for bandwidth upgrades while improving performance for business-related activities. “At one point, one of our larger offices was running into performance problems, which we eventually traced to one person who was uploading and downloading a tremendous number of files at work,” says Brumbaugh. “Because this was not business-related, we talked directly with that individual about Cisco’s policy regarding non-work-related behavior. The performance problems cleared right up.”
- **Unwise traffic** – Harder to describe than inappropriate or legitimate traffic, unwise traffic can result from how and where business-related applications are used. Backups or database synchronizations performed at inappropriate times or over inappropriate segments of the network are obvious offenders. Traffic consuming significant bandwidth during peak hours that can be safely moved or rescheduled is unwise traffic. Determining which traffic fits this category is the responsibility of the capacity planning engineer. In many cases, applying standard QoS configurations automatically slows unwise traffic by marking it “scavenger-class” and not allowing it to impinge on other traffic during hours of peak use. Interestingly, unwise traffic is not necessarily scavenger-class traffic. During traffic analysis, capacity planning engineers can choose to reschedule, eliminate, or categorize traffic as scavenger-class. “Clients were complaining about WAN performance between our Irvine office and headquarters in San Jose,” says Silver. “After the analysis, we determined that the circuit was congested with SnapMirror backup traffic. SnapMirror backs up servers in the data center and is not considered mission-critical. Working with the backup team, we decided to categorize backups as scavenger-class, which allowed them to be throttled back during times of congestion. We avoided a bandwidth upgrade, and overall WAN performance improved immediately.”

Capacity planning should provide volume and content traffic information to network architects, designers, and operators—making it possible to size the network accurately while meeting business requirements. Capacity planning should also provide management and finance executives with the data required for budgeting and forecasting by pointing out which connections are approaching saturation and will require an upgrade.

Sizing and Utilization Guidelines

Cisco planners believed it was vital to establish sizing and utilization guidelines to serve as baselines for managing network capacity (Tables 1 and 2). The planning team found that initial sizing guidelines based on headcount were appropriate for most Cisco field sales offices. However, bandwidth generalizations were not always appropriate for some engineering and extranet sites and Internet POP locations. These types of locations required evaluation on a site-by-site basis. Equally important, planners realized that while guidelines were important, they did not eliminate the need for an engineering analysis to ensure that the bandwidth solution was appropriate for each location.



Table 1. Sample Initial Sizing Chart

Headcount	Primary WAN Bandwidth	Secondary WAN Bandwidth
1–10	1.5 Mbps	None
11–40	1.5 Mbps	1.5 Mbps
41–100	3 Mbps	3 Mbps
101–150	4.5 Mbps	4.5 Mbps
151–200	6 Mbps	6 Mbps
201–500	45 Mbps	6 Mbps
501+	155 Mbps	45 Mbps

Table 2. Sample Utilization Threshold Chart

Primary/Backup WAN Bandwidth	* Percentage Utilization (Watch/Analyze)	** Percentage Utilization (Upgrade)
*** 100/100	60%	80%
**** 100/50	40%	50%
***** 50/50	40%	50%

- * 15-minute average threshold exceeded 10 percent or more during local business hours (monthly)
- ** 15-minute average threshold exceeded or equaled 20 percent of local business hours (monthly)
- *** Primary circuit handles 100-percent of the traffic until failure, when the backup takes over
- **** Primary circuit handles 100-percent of the traffic until failure, when the backup takes over with 50 percent of the capacity of the primary
- ***** Primary and backup circuits load-share until a failure occurs


Whether provisioning bandwidth at a new location or deciding when to upgrade an existing circuit, Cisco planners were acutely aware that their decisions must be cost-effective. Though established guidelines were usually appropriate, planners found that in some locations, higher-bandwidth circuits were less expensive than lower-bandwidth solutions. In these cases, capacity planners based decisions on cost, rather than the bandwidth actually required.

The Tools

Cisco characterized, analyzed, and detected anomalies in network traffic flows using Cisco IOS® NetFlow technology, including NetQoS ReporterAnalyzer, their selected third-party reporting solution. This solution used the data captured by Cisco IOS NetFlow to report on network traffic..

Cisco IOS NetFlow (Figure 2) has become the primary network accounting and anomaly detection technology in the network industry. In fact, in 2003, Cisco IOS NetFlow Version 9 was chosen for a proposed IETF standard, the IP Flow Information Export (IPFIX). IPFIX defines the format by which IP flow information is transferred from an exporter, such as a Cisco router, to an application that analyzes the data.

Figure 2. Overview of Cisco IOS NetFlow



“Essentially, we turned on NetFlow with no negative impact to the network,” says Brumbaugh. “It didn’t create CPU or memory problems on routers, and collecting the data didn’t saturate our WAN links. We never used probes—they are intrusive, and there tend to be scalability issues. We selected NetQoS as our capacity planning application largely because it took advantage of the Cisco IOS NetFlow capability present on all our routers.”

“NetFlow also saved us time,” adds Silver. “Though the system does require some attention, it is minimal compared to what we had to do five years ago. Back then, our IP accounting system was very hands-on. It could take 20 hours to harvest data—20 very tedious hours—and the results were often poor. Now, we get detailed data in a matter of minutes.”

Reporting

Cisco regularly performs capacity planning on existing locations, though it is currently not practical to perform detailed ongoing analysis on every circuit. Reports generated by Cisco IOS NetFlow, plus size and utilization guidelines, help planners determine the circuits they must watch and where bandwidth augmentation is necessary. Global capacity planners are alerted proactively when circuits reach established thresholds via daily, weekly, and monthly reports that highlight circuits above the established “watch/analyze” threshold of 60-percent utilization for 10 percent or more of local peak traffic hours.

Results

When the capacity planning team determines that bandwidth augmentation is necessary, it provides a recommendation to Cisco IT’s network operations, where it is reviewed by IT network engineers and managers. Once it is determined that an upgrade is necessary, proceeding with the upgrade becomes a business decision. “I’d been seeing a steady increase in bandwidth utilization on our circuit to India over the past 12 months,” says Brumbaugh. “I knew that if it continued at that rate, the link would be saturated



quickly. Because I knew at a granular level what the traffic was, I could comfortably help build a business case for increasing bandwidth. It's great to work with real information." After going through Cisco IT's network design process, the architects and engineers in IT decided they needed to begin a major upgrade in the design of the Cisco India WAN and the backbone links connecting India to the rest of the world.

In addition to providing business decision makers with hard data and communicating with them more effectively, Cisco capacity planners can now prioritize and manage deployments better—delivering bandwidth before performance deteriorates and productivity decreases. A coherent planning process has also made it easier to understand the impact of application rollouts. The result has been an improved relationship with the application team—both groups can better plan and budget activities that affect each other's operations.

Lessons Learned

"We're guessing that many of the problems we experienced are shared by 99 percent of the enterprises out there," says Silver. "What we learned, very simply, is that attempts to plan network capacity without appropriate tools are inaccurate and expensive. We knew that we needed to improve both our tools and our processes—and when we did, we realized that the money we spent implementing NetFlow was low in relation to the amount we had been spending without it."

Next Steps

The Cisco capacity planning team plans to use Cisco IOS NetFlow technology to develop an even more detailed view of network traffic. In a MPLS environment, where the enterprise pays its bandwidth circuit providers differentially for different classes of service, it is important not to pay for service that is unnecessary. By monitoring and reporting on classes of service at a deeper level, Cisco expects to be able align what it pays for service with what is actually needed, saving significantly going forward.



For additional Cisco IT case studies on a variety of business solutions,
go to Cisco IT @ Work

www.cisco.com/go/ciscoitatwork

Note:

This publication describes how Cisco has benefited from the deployment of its own products. Many factors may have contributed to the results and benefits described; Cisco does not guarantee comparable results elsewhere.

CISCO PROVIDES THIS PUBLICATION AS IS WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some jurisdictions do not allow disclaimer of express or implied warranties, therefore this disclaimer may not apply to you.



Corporate Headquarters

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

European Headquarters

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

Americas Headquarters

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

Asia Pacific Headquarters

Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 317 7777
Fax: +65 317 7799

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

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

Copyright © 2004 Cisco Systems, Inc. All rights reserved. CCIP, CCSP, the Cisco *Powered* Network mark, Cisco Unity, Follow Me Browsing, FormShare, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, MICA, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, ScriptShare, SlideCast, SMARTnet, StrataView Plus, Stratum, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0505)

Printed in the USA

