



WHITEPAPER

Performance First™

A Performance Mindset for Network Management Keeps Organizations Functioning at Optimum Levels

Over the last two decades, IT organizations have spent billions of dollars implementing fault management tools and processes to maximize network availability. While availability management is critical, infrastructure reliability has improved to the point at which 99.9 percent availability is commonplace. Given these improvements in infrastructure availability, companies are focusing more attention on performance management. By measuring how networked applications and services perform under normal circumstances, understanding how infrastructure and application changes impact performance, and isolating the sources of above-normal latency, IT organizations can ensure problems are resolved quickly, mitigate risk from planned and unplanned changes, and take measured steps to optimize application performance. In this whitepaper, you will learn why this shift is taking place and how a new management model, what NetQoS calls Performance First™, will empower you to advance to the next level in managing your network for application performance.

Most network engineers are very familiar with tools that report statistics on individual components such as links, routers, and servers. These infrastructure monitors have been around for a long time. Newer to the market are performance management appliances that monitor end-to-end service delivery, measuring the end-user experience, and providing a unique and comprehensive view of the enterprise network without the need to deploy remote probe devices or software agents on client desktops and servers. These performance management appliances not only measure how well response-time service level agreements (SLAs) are being met, they also help IT staff solve a wide variety of problems, leading to significant reductions in operating costs. Furthermore, given today's economic climate, it is more crucial than ever to make fully informed infrastructure investments, avoid unnecessary expenditures by optimizing the use of existing equipment, and ensure there is no inappropriate use of expensive network resources.

With the ubiquity and increasing importance of networked applications, IT organizations have begun to shift their approach to network performance management. The first concern for network engineers and operations managers is no longer just infrastructure availability, but also how well applications perform over the network. Reviewing the evolution of network management will help explain why this is happening and what it means.

Managing Network Availability

For most of the past two decades network engineers have focused on managing availability. Even today most of the established products from network and systems management vendors are designed to tell network managers whether or not key infrastructure components are working: Is that router up? Is this link down? Does the server need to be rebooted? Historically, network engineers needed to ask these questions frequently. Often routers did malfunction, links did fail, and hardware or operating system failures required servers to be restarted. And when the network infrastructure failed, operations ground to a halt.

Today, however, most enterprise and service provider networks operate with 99.9 percent uptime or better. What drove the change? One reason is that technology is much more reliable. In addition, more and more mission-critical applications today are networked, and managers have realized that network downtime and slow-downs can be excruciatingly expensive to the business. For instance, in a brokerage firm, application downtime can cost as much as six million dollars per hour. This high penalty has made it worthwhile for network managers to invest in redundant servers, hot-swappable components, replication systems, and other high-availability options to guarantee uptime.

Managing Network and Application Performance

While the reliability of networks has improved, the demand placed upon today's complex networks has caused performance issues to increase dramatically. Compounding the challenge is the now common expectation from today's technologically savvy end users of a ubiquitous and instantaneous network. Drivers of this increase in network traffic volume and complexity, and the need to monitor how well applications are delivered over the WAN include:

» **Data center consolidation.**

Enterprises and government agencies alike are migrating applications and data to central locations to save real estate, infrastructure, power and personnel costs, and improve manageability and security.

» **Increased number of remote users.**

Branch offices and telecommuters are proliferating as companies grow, merge, and expand globally; the exponential growth of e-commerce transactions continues unabated.

» **The rise of voice and video traffic.**

Voice and video traffic are increasing rapidly and both are particularly dependent on the quality and consistency of network delivery.

» **Legacy applications.**

"Chatty" applications designed to run over local area networks often do not perform well when deployed over the WAN.

» **Software as a service (SaaS).**

More organizations are choosing applications, such as Salesforce.com, that are delivered over the Internet as a service, in place of internally-hosted, client-server applications.

» **More complex applications including Service-Oriented Architectures (SOA).**

Distributed architectures and the use of Web services as a means to develop reusable software for rapid application delivery and easier maintenance introduce network traffic between the various application tiers and infrastructure components. Enterprise applications may have ten or more networked tiers with each tier representing a possible failure point or bottleneck.

» **Virtualization.**

Server virtualization is a powerful technology for consolidation, flexibility and cost savings, but concerns about visibility, bandwidth contention, performance, and overall management need to be addressed for critical applications to migrate to virtualized infrastructure, much less a cloud computing infrastructure, at a fast pace.

» **Outsourcing.**


In today's world of outsourced networks and managed services, ensuring service providers meet their SLA obligations and gaining visibility into the causes of degraded service is essential.

The convergence of increasing WAN use with improved device availability is leading network engineers to put performance first when it comes to managing their complex networks (see Figure 1). They see greater returns by shifting their focus from fault management—which is largely under control—to performance-based management, concentrating instead on how the network is affecting service delivery.

To put performance first and understand where there is potential for cost savings and efficiencies on the network, IT professionals must be able to:

- See how the network is being used, including which applications and users are the most active, and when.
- Measure how critical applications and services are performing around the clock and around the world.
- Determine when and where there is inappropriate and wasteful use of network resources.
- Identify the root cause of slow-downs and fix them before they impact business results.

Figure 1



Over the past two decades, efforts by IT organizations and infrastructure vendors have reduced infrastructure downtime to a minimum. The most important trend in wide area networking today is the growing need to manage application performance.

The Case for a New Management Approach

A Yankee Group report—aptly titled “Performance is the New Mandate for Network Management”—refers to a study on enterprise application management. The study found enterprises report an average productivity drop of 14 percent when experiencing performance degradations in Oracle, SAP, PeopleSoft, Siebel, and custom .NET and J2EE applications. In addition, a survey of 176 IT professionals conducted by Ashton, Metzler & Associates uncovered the fact that over a quarter of network operations centers (NOC) do not meet their organizations’ current needs. In order to fulfill current and emerging requirements, NOCs are being driven to do a better job of managing application performance, to implement more effective IT processes, and to be able to troubleshoot performance problems faster. Clearly “Performance First™” is becoming the new standard for network management.

Other validations of the Performance-First paradigm appear in the ever-growing number of whitepapers, analyst reports, and trade magazine articles that tout headlines such as: “The Real Value of Network Visibility,” Aberdeen Group, or “Blueprint for Application Performance Management,” Network World Magazine, and “The Performance Management Mandate,” Kubernan.

A New Model for Network Management: Performance First™

The Performance-First model flips the status quo on its head. It inverts the traditional, bottom-up device monitoring approach and begins with top-down visibility into overall performance of applications running

over the network. In this approach, infrastructure availability and utilization are no longer the sole gauges of network health. After all, it makes little sense to focus 100 percent of network management efforts on the 0.1 percent or less of the time there are hardware or software infrastructure failures. Fault management is necessary, but not sufficient.

The Performance-First approach is driven by the fundamental purpose of the network infrastructure—to transport data from one end of the system to the other as rapidly as possible. The more efficiently data flows at the transport layer the better the application performance. Latency is the best metric to use when deciding how to optimize the network, plan new infrastructure rollouts and upgrades, and identify the severity and pervasiveness of problems. This approach recognizes that between the limits of the network and application infrastructure being “up” or “down”, there is a wide spectrum of performance variation. It is not uncommon for availability status indicators in the NOC to be all green while the help desk phones are ringing off the hook with users complaining about slow response times.

Measuring the latency of key applications running over the network and identifying where there is opportunity for improvement enables IT organizations to focus on the most important issues: making informed infrastructure investments to support business demands, delivering consistent, acceptable end-user response times, and resolving problems faster. IT organizations that successfully make the transition to a Performance-First management approach typically receive high marks from the lines of business they serve.

Characteristics of the Performance-First Organization

According to an Aberdeen Group report, “The Value of Network and Application Visibility: Improving the Usability of Performance Data,” enterprises with best-in-class network and performance management solutions spend less to manage network and application performance, and have an 85 percent success rate in preventing application performance issues, as compared to 32 percent of all others. To achieve this level of success, organizations with a best-in-class solution incorporate the following network management disciplines into their standard operating processes:

Maintain visibility into network traffic across the entire IT infrastructure. Network traffic is monitored on all WAN segments, to and from the data center and remote locations. Network engineering staff have 24/7 visibility into link utilization, traffic composition and patterns, active users and applications, class of service traffic distribution, and routing paths.

Monitor the health of all key infrastructure components. New and updated devices are automatically discovered (sometimes incorporated in a federated CMDB) and monitored for key metrics such as availability, processor and memory utilization, interface utilization, packet errors, and discards.

Monitor the delivery of key data, voice and video applications to end users. Critical applications are catalogued and mapped to data center infrastructure components. Service level objectives are established for each, and actual performance is tracked against these targets using the appropriate mix of passive and active measurement techniques. Latency is measured in each segment of the application delivery infrastructure, including the client side, the WAN link, and in each of the various application tiers (Web, database, etc.) in the data center. Key applications receive prioritized network access through the use of quality of service (QoS) policies and traffic acceleration and optimization devices. Summary SLA compliance and trending reports are generated and distributed to all interested parties.

Establish baselines for normal performance so deviations are easily detected. Automated baselines, derived from monitoring production application flows, establish the normal performance for important applications at various times of the day, day of the week, and week of the month to accommodate fluctuations across business cycles. These adaptive baselines are used to trigger early warning alerts of performance degradation to operators. Baselines are preferred over traditional static thresholds, building more intelligence into the “manage by exception” mantra.

Use well-defined procedures for troubleshooting problems and investigating the root cause of degraded or interrupted service. Early detection technology identifies anomalies in key metrics such as response times, traffic flows, and device utilization, which may indicate an emerging performance or security problem. Alarms that are linked to fault tracking and ticket or workflow management systems alert operations staff of problems. Relevant performance metrics and events are correlated to establish the scope and impact of an incident. Root cause investigations are automatically launched when service starts to degrade, and escalation procedures are documented and followed, including the gathering of appropriate performance metrics at each step. When necessary, user and application behavior at the time of an incident can be recreated for debugging purposes using stored packets. Key Performance Indicators (KPIs) such as mean time to repair (MTTR) and mean time between failures (MTBF) are tracked over time to establish a cycle of continuous improvement.

Use standard change management processes to control, verify, and measure the impact of planned changes. The performance of infrastructure components and their impact on application service delivery are measured before and after any configuration changes are made. Changes that cause an unexpected degradation in application performance are quickly backed out. Often, lab resources are available to test proposed changes and their impact on performance before production deployment.

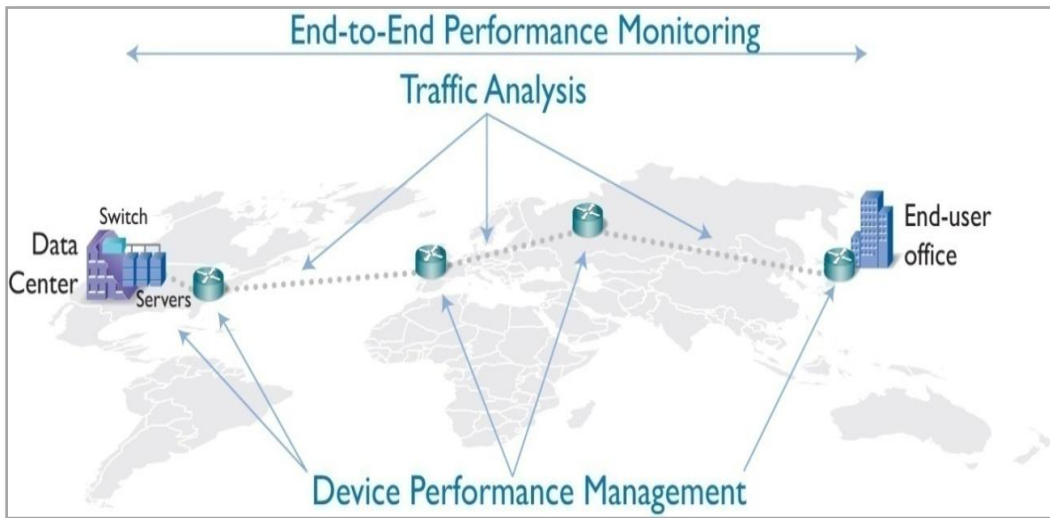
Plan infrastructure capacity requirements. Trend data for application performance and infrastructure utilization is analyzed and correlated with anticipated demand provided by business units. Usage-based cost analysis gives visibility into the cost of bandwidth and computing resources, for example by business unit, branch office, and application. This information ensures infrastructure upgrades and additions are planned and provisioned in advance, thus allowing for best vendor pricing to be negotiated.

Provide application service delivery reports to all constituents. Reports are automatically generated and electronically delivered to IT managers, business unit managers and executives. The reports are organized by business construct such as functional area, branch, geography, or application service. Summary views of service delivery over short- and long-term periods provide visibility into developing trends, candidate areas for additional infrastructure investment, and alignment with business priorities. Performance data is treated as a data warehouse, enabling "what-if" analyses and in-depth usage, trending, and business analysis.

The Critical Components of Network Performance Management

Managing application delivery across modern enterprise and service provider networks is a complex undertaking. The Performance-First approach is based on the premise that you must capture and analyze data from applications, devices, and the network itself to get an accurate and comprehensive understanding of how an organization is supporting application delivery. See Figure 2.

Figure 2



A Performance-First approach starts with measuring response times and quality of service end-to-end to get an overall view of application delivery, and employs other key performance metrics and analyses as needed, including traffic flows, device performance, and even application replay to reproduce problem circumstances.

» **Application Response Time Monitoring.**

Track, measure, and analyze application performance for all user transactions from end-to-end for insight into the end-user experience and the source of any latency issues.

» **UC (Unified Communications) Quality of Experience.**

Monitor call and video quality and network-based call setup, and measure the impact of convergence across all application performance.

» **Traffic Analysis.**

Visualize and analyze the composition of network traffic on specific links. This yields the information needed to redirect or reprioritize application traffic, detect anomalous behavior, or add capacity.

» **Device Performance Management.**

Employ “application aware” routing capabilities such as Cisco CBQoS, IP SLA, and NBAR and to poll network infrastructure components to isolate the source of problems such as a busy router or a server memory leak so that corrective action can be taken.

» Long-term Packet Capture and Analysis.

Store detailed packets for analysis before, during, and after incidents, without needing to recreate the problem.

Initiating a Performance-First Approach: Establishing a Baseline

Whether troubleshooting a bottleneck, monitoring a new application rollout, or upgrading the network infrastructure, a Performance-First management approach starts by understanding and establishing an overall performance baseline. What is normal performance for an application or a group of users? How does “normal” change during busy and off-peak business cycles? Which applications and users are experiencing poor performance? Where is the increased latency occurring? Is it in the network, server or application itself? What impact did the new application have on other application response times? Did the infrastructure upgrade deliver the performance boost expected? Response time monitoring gauges how well the network is delivering services to the end user and provides the best overall view of what is happening on the network.

Beyond the Baseline

Once end-to-end performance metrics have been captured and the source of latency isolated, further analysis becomes more focused. Long-term packet capture and analysis speeds troubleshooting with “back-in-time” packet-level detail and stream reconstruction to replay the actual application and user behavior at the time problems occurred. Traffic analysis enables network engineers to understand the composition of traffic on specific links where latency is higher than normal or expected.

If the source of latency is isolated to an infrastructure component—a busy router or a server memory leak, for instance—network managers need device performance management capabilities to poll the device in question and isolate the problem so corrective action can be taken.

If latency cannot be attributed to the network or the server infrastructure and can be shown to be isolated in the application itself, the network team is armed with the proof that will eliminate the typical finger-pointing between IT infrastructure and application teams.

For enterprises dependent on high quality and reliable UC performance, specific metrics such as call quality, handset details, and mean opinion score (MOS) are needed to measure UC quality of experience. UC performance also needs to be tied to an overall network performance context to ensure that VoIP and video are not choking out other networked applications and vice-versa.

Network Traffic vs. Vehicle Traffic: An Analogy

Managing network traffic is similar to managing vehicular traffic in a metropolitan area. An end-to-end performance view of the network is similar to the view a helicopter has of traffic in a city. The helicopter pilot can observe generally how well traffic is flowing and where there is congestion.

However, he can't always tell if the problem is because of a slippery road surface, a stalled automobile, a bad traffic light, or an unusually high number of slow-moving trucks from a nearby construction site. Closer inspection of traffic elements and road conditions is required to determine the cause.

On a network, a full-spectrum view of performance provides the first piece of information required for effective management: Is application performance normal? Is it better or worse? State-of-the-art monitoring technology can also determine where latency is abnormal by timing packet flows across the network and server infrastructure, and through intelligent base-lining to determine a typical performance profile. With this information, it is possible to determine quickly which user locations and applications are affected, which WAN links are bottlenecked, or which server in a particular tier of the application infrastructure is performing poorly.

If the cause of congestion is determined to be a network issue, traffic analysis can be used to drill down to view details on the links in question. Monitoring network traffic flows is roughly analogous to observing vehicular traffic on a stretch of road passing through a specific intersection, providing specific information about the type and volume of traffic for problem resolution and planning. Analyzing network flows on an individual link determines how much traffic is HTTP, SAP, or streaming audio for example, and also provides a measure of volume by protocol, host, and conversation. Additionally, capturing every packet and making them available for replay is like having a camera at an intersection that monitors all traffic, but also flags the red-light runners for quick identification.

Without the end-to-end performance (helicopter) view to direct the focus of diagnostic efforts, traffic analysis on a link-by-link basis in a very large network with thousands of links is at best, inefficient. Furthermore, attempting traffic analysis by placing probes at each router is an expensive and ongoing management burden, similar to trying to locate a traffic slowdown somewhere in the city by sending 500 cars to look at each of the 5,000 roads to find the problem. This approach may eventually work, but it is a huge waste of resources and time.

Cisco IOS[®] NetFlow and IP Flow Information Export (IPFIX) technology, embedded in today's enterprise routers and switches from the major network equipment manufacturers, makes traffic flow analysis possible without expensive probe deployments. By utilizing the existing router and switch infrastructure, IT staff can harvest and analyze NetFlow data to provide comprehensive flow volume and composition data for each link in the enterprise. This may be used for troubleshooting, capacity planning, and measuring the impact of changes.

Device-level performance is helpful for capacity planning and fault management (a broken traffic light). Several performance details for each device can be checked with SNMP device-level statistics such as the percent busy for a server or router processor, memory utilization, packet errors, and discard rates.

Automobile traffic can also be affected by factors other than cars and trucks, so to complete the analogy, imagine that UC applications are like street cars. They share the same basic infrastructure and have to cross intersections and obey stop lights, but they have some unique characteristics that require special attention. For UC, this means creating and monitoring baselines specific to UDP traffic as well as gathering statistics down to the handset.

The Benefits of the Performance-First Paradigm

Putting performance at the forefront of network management positions IT organizations to benefit from numerous outcomes:

» **Prove the performance of applications running over the network.**

Too often IT managers are unable to provide objective measurements of performance against application SLAs. They have no way of knowing how well their service provider is meeting its performance targets. Frequent user complaints may be unfounded, but can't be disproved.

» **Deliver consistent application performance and measure it.**

You can't manage what you don't measure. Without real-time visibility into end-user response times, traffic flows, and infrastructure health, it's impossible to manage application performance proactively.

» **Mitigate the risks from planned changes and unexpected events.**

How many times each week in an IT group is the question asked, "What changed?" More application

outages and “brown-outs” are caused by planned changes with unintended results than any other cause. It is critical that the impact of these changes is discovered and isolated immediately.

» **Make more informed infrastructure investments.**

When infrastructure managers make uninformed upgrade decisions, the cost can be high. Often the anticipated results don't materialize, ROI is negative, and performance problems persist. Infrastructure utilization metrics alone are inadequate; knowing the exact source of latency and the composition of network traffic can often present alternatives to expensive upgrades, such as changing QoS rules or rescheduling data-intensive applications such as backups. Without before-and-after measurement of WAN optimization and application acceleration efforts, deployment of these technologies and techniques will be trial-and-error, often resulting in exhausted budget and investments that don't yield the intended results.

» **Work collaboratively and more effectively.**

Performance management in large WANs requires laser-like precision because in today's distributed network applications there are thousands of hardware and software elements affecting end-user response times. Network managers need tools that give them real-time global visibility and historical information to optimize the network infrastructure for application performance and work with peer groups to plan for changes.

» **Troubleshoot problems faster to reduce MTTR.**

When a performance problem does arise, what is the typical process to resolve it? Who gets involved? How long does it take on average to fix problems? How much of this time is spent finger pointing between IT functional groups? Knowing the source of availability and performance problems means the right technicians can be assigned immediately, and having detailed diagnostic information means they can fix problems quickly.

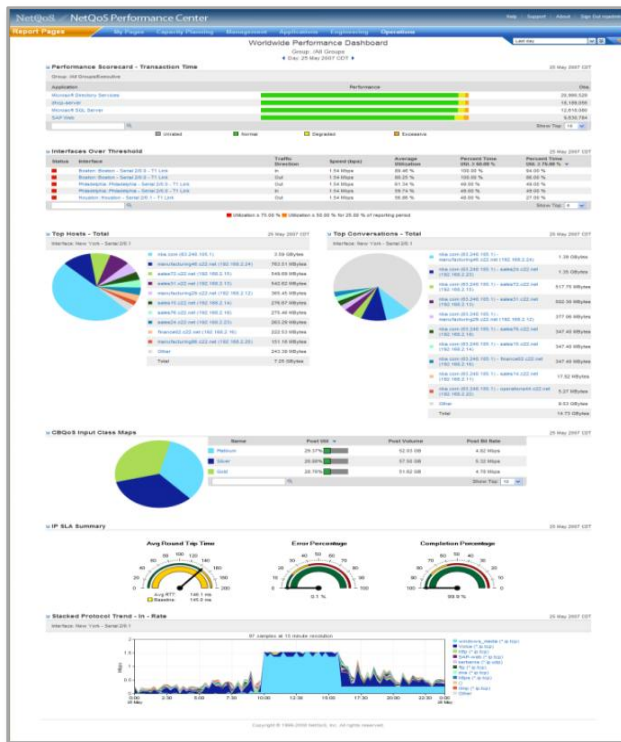
The NetQoS Performance Center

Recognizing the absence of tools to support the Performance-First management paradigm, the NetQoS founders set out to fill the void. With top-down performance analysis spanning enterprise views of the end-user experience down to deep-packet inspection, the NetQoS Performance Center suite of products provides global visibility into all of the core metrics needed to sustain and optimize application delivery. The NetQoS Performance Center offers best-in-class functional modules that scale to support the world's largest and most

complex networks, leverage embedded, industry standard instrumentation as data sources, and eliminate the need to deploy remote network probes or desktop and server agents.

The NetQoS Performance Center portal provides integrated reporting from the NetQoS functional product modules, allowing network engineers, operations staff and IT managers and executives to view the role-specific metrics they need with real-time data and historical views presented in a single Web-based dashboard. See Figure 3.

Figure 3



The NetQoS Performance Center integrates performance metrics from end-to-end application response time, traffic analysis, device performance, UC quality of experience, and long-term packet capture and analysis.

To understand the full value of what the NetQoS Performance Center offers, it is important to understand the contributions of each product module.

Application Response Time Module – NetQoS SuperAgent®

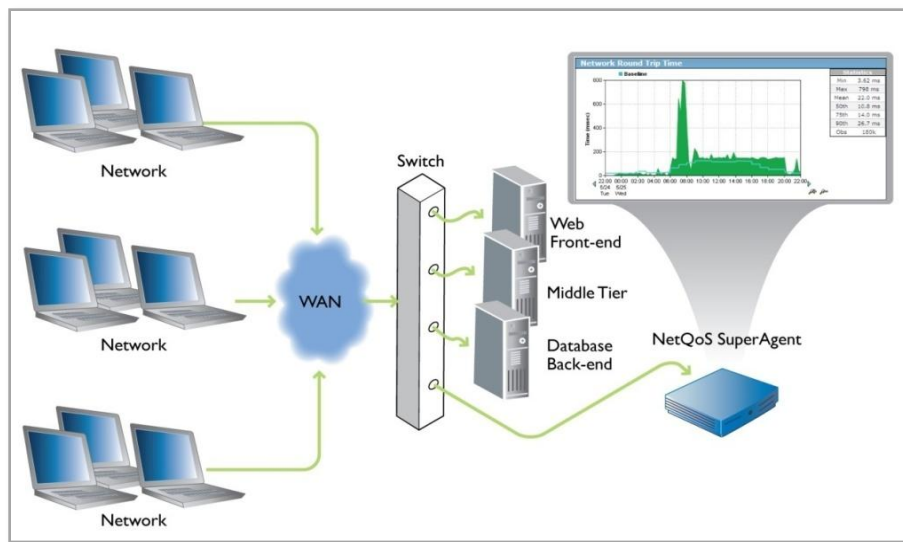
The most useful measure in establishing how end-to-end performance appears to the user is response time. NetQoS SuperAgent® monitors all the TCP application packets from the network into the data center and out

again. This provides a way to measure network round trip time, server response time, data transfer time, and much more.

SuperAgent is the first product in its class to capture this information passively without the use of server or desktop agents. Before SuperAgent, enterprises had to deploy software on application servers and end-user machines to get a view of application performance. For enterprises with hundreds of applications and tens of thousands of users, it was a management nightmare. As part of the efforts to keep data collection as efficient as possible, NetQoS has worked with Cisco to enable Cisco Wide Area Application Services (WAAS) devices and Cisco Network Analysis Modules (NAM) to send response time data directly to SuperAgent without additional appliances, probes or agents.

With SuperAgent, a single device deployed in the data center can report application performance for all users across all locations. The SuperAgent appliance is attached to a switch mirror or span port, or a network tap near the server farm. This location provides a way to effectively inspect all TCP packet headers and calculate end-to-end response time metrics based on the source and destination IDs. Latency for each network hop, application server tier, and application component can also be determined. Most applications today have multiple tiers, and incoming client requests are routed to be processed, piece-by-piece, by a combination of different servers such as a transaction, application, and database servers. This means the only place to effectively inspect and troubleshoot multi-tiered application transactions is from the data center. See Figure 4.

Figure 4



The data center is the only measurement point to capture end-to-end response time and the latency for each hop in the application infrastructure. SuperAgent taps into the switch mirror port to detect all TCP application packets.

SuperAgent measures the latency in the application, server, and network components against normal performance baselines that SuperAgent generates automatically. Baselines are established for each hour of the day, each day of the week, and each week of the month. Normal fluctuations between peak and off-peak transaction loads (nights, weekends, end of month, end of quarter, etc.) are accommodated. Alerts may be triggered when baseline performance is violated so that network managers know immediately, 24 hours a day, if end-user response times are acceptable, and if not, where latency is abnormal (application, server, or network) so that further diagnosis may be done.

Similarly, this information may be used very effectively to prepare for new application rollouts, additional users, and other changes. The Performance-First management approach enables IT staff to be much more proactive in managing the network for application performance.

VoIP and Video Quality of Experience Module – NetQoS Unified Communications Monitor

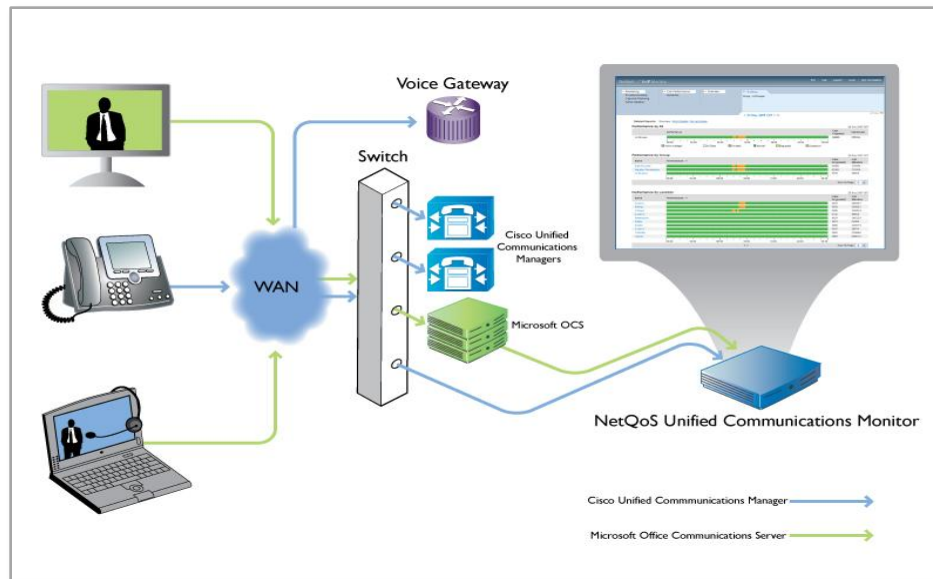
More than any other type of application, voice and video over IP (VoIP) require optimal network performance to deliver an acceptable quality of experience for users. NetQoS Unified Communication Monitor tracks the call quality user experience, provides alerts on call performance problems, and isolates performance issues to speed troubleshooting and reduce MTTR. With Unified Communications Monitor, the performance of the Cisco Unified Communications Manager (CallManager) IP PBX can be assessed by tracking, evaluating, and reporting on key metrics—all without deploying server agents or probes. Call setup metrics, such as Delay to Dial Tone and Call Failures, show how well a CallManager IP PBX is performing, helping to better monitor UC performance on the network.

Unified Communications Monitor also provides monitoring capabilities for Microsoft Office Communications Server (OCS) environments. Instead of a system-based monitoring focus for OCS, Unified Communications Monitor measures end-user quality of experience across the network, including video metrics such as frame loss and frozen video.

For every call, Unified Communications Monitor reports on Mean Opinion Score (MOS) and metrics associated with underlying network factors, such as packet loss and jitter. Unified Communications Monitor also breaks out performance data from the IP and Public Switched Telephone Network (PSTN) legs of calls that pass through voice gateways traveling to endpoints in the PSTN. This data highlights if it is the network that is responsible for less-than-optimal call quality.

Unified Communications Monitor is designed to function with the other capabilities of the NetQoS Performance Center so that VoIP and video quality of experience may be monitored while managing overall network quality of service from a single Web-based console. In addition to voice and video metrics sourced from the dedicated instrumentation and analytics, a proper UC dashboard should also include metrics to establish that non-UC applications are not choking out UC performance (and vice versa). Another best practice is to use regular IP SLA tests to test locations for UC readiness and continued performance.

Figure 5



Unified Communications Monitor passively monitors all traffic to and from the call manager using SPAN port access.

Network Traffic Flow Analysis Module – NetQoS ReporterAnalyzer™

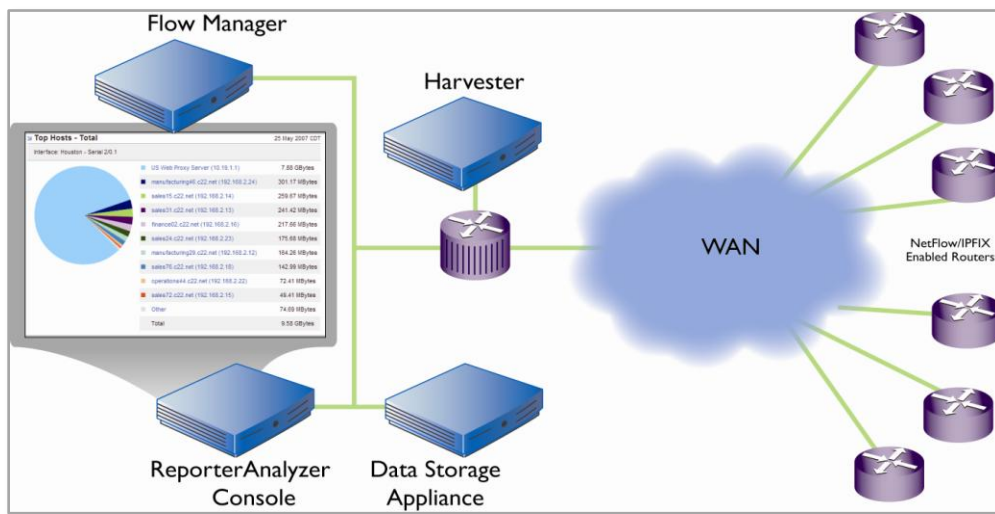
Having visibility into the composition of traffic on every link in the network—including which users and applications are consuming bandwidth—gives network engineers the information they need to plan effectively for new applications and additional users, and make smart decisions on router configurations, bandwidth allocations, and upgrades. If end-user response time degradation is caused by abnormal latency in a network segment, as determined by SuperAgent, analysis of that segment's traffic composition can quickly identify the source. The NetQoS traffic analysis module—ReporterAnalyzer™—can identify a remote branch office backing up its server during the daytime, users downloading music files, a denial of service attack, or a fast-spreading virus.

ReporterAnalyzer was the first product to harvest and report on enterprise-wide NetFlow data (see Figure 6) and is the “gold standard” in the industry. Today ReporterAnalyzer also harvests flow data from enterprise routers and switches manufactured by Nortel, Juniper, and others that have adopted the emerging IPFIX standard. By leveraging the flow data instrumentation in these devices, NetQoS makes it possible to deliver global visibility into every flow across the network without deploying probes.

ReporterAnalyzer is unique in its ability to report on 100 percent of flow traffic for the entire network and access more than a year’s worth of enterprise-wide data for trending and detailed planning. NetQoS provides a complete view of application traffic while many other vendors only offer Top-10 reporting solutions, reporting traffic on only the most-used ten applications on an interface while omitting hundreds or even thousands of other applications.

ReporterAnalyzer monitors the distinct performance of various virtual circuits—i.e., discrete data paths of varying speeds on the same physical link. This makes it possible to troubleshoot Multi-Protocol Label Switching (MPLS) networks and Virtual Private Networks (VPN), which fence off certain portions of the network.

Figure 6



ReporterAnalyzer leverages native instrumentation, such as Cisco IOS® NetFlow, to monitor flows for every host, protocol, and conversation across the network.

Device Performance Module – NetQoS NetVoyant®

While a Performance-First approach to network management is displacing fault-centric models, utilization statistics and analysis are still critical for comprehensively managing application delivery. NetQoS NetVoyant®, the device performance module of the NetQoS Performance Center, provides SNMP-based performance metrics for network infrastructure, devices, and services. It polls WAN and LAN components throughout the infrastructure to collect statistics such as CPU and memory utilization, component or service availability, packet discards and errors, and information about FECNs and BECNs, used to address network congestion during peak usage periods. NetVoyant provides the ability to import and report on custom Management Information Bases (MIBs).

NetVoyant support for embedded Cisco IOS capabilities allows network teams to fully leverage their existing infrastructure to gain better testing, visibility, and differentiated classes of service. NetVoyant eases configuration and provides complete reporting for Cisco IP Service Level Agreements (IP SLA) enabling more accurate and useful tests for jitter, latency and packet loss, which can be especially useful when rolling out VoIP. NetVoyant leverages Network Based Application Recognition (NBAR) metrics from the Cisco Catalyst® 6500 Supervisor Engine Programmable Intelligent Services Accelerator (PISA) to show live application traffic across the network. Analysis of these metrics allows operations staff to monitor the network infrastructure without the expense of RMON probes and provides critical guidance for configuring QoS policies. Reporting on Cisco Class-Based Quality of Service (CBQoS) traffic metrics delivers detailed, forensic views of how QoS policies impact each class of service— and the resulting link and device performance —throughout the network.

NetVoyant helps network engineers and operations managers solve problems quickly by singling out the causes of device performance issues. NetVoyant helps IT staff manage network and server infrastructure capacity with automatic calculations that compare

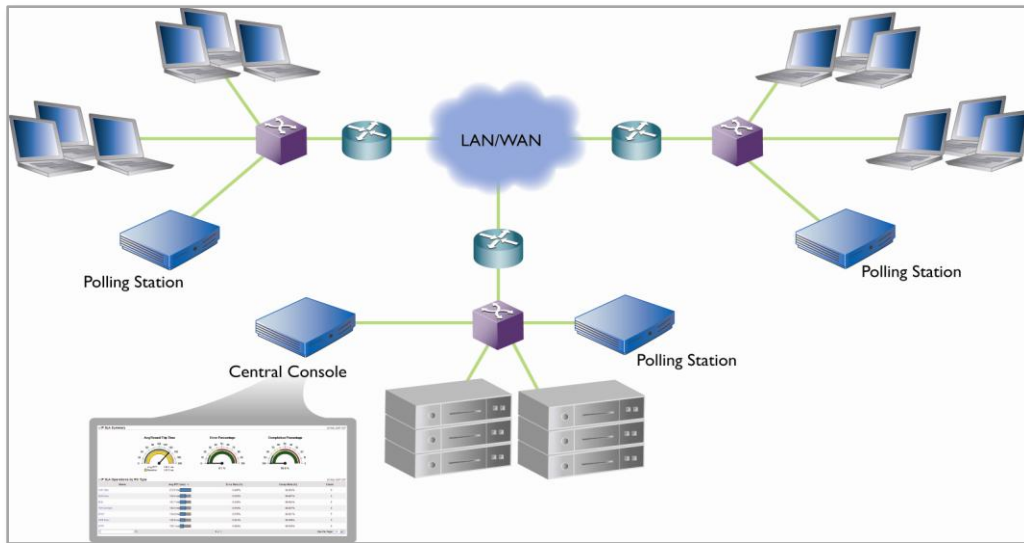
NetQoS Anomaly Detector: Performance-First Network Behavior Analysis

NetQoS Anomaly Detector, a component of the NetQoS Performance Center, analyzes traffic flows, response times, SNMP data, and other metrics to profile the network and provide early warning alerts of changes in behavior for any client or server, whether those changes are impacting performance or signaling a security breach.

Anomaly Detector leverages the data sources of the other NetQoS products, so there is no need for additional network instrumentation.

current performance values with normal baselines. In addition, NetVoyant trend reports help staff understand the growth characteristics of network utilization and device resource consumption and make better informed capacity and performance planning decisions.

Figure 7



NetVoyant polls infrastructure components and reports on their health.

Long-term Packet Capture and Analysis Module – NetQoS GigaStor™

Use of a single, integrated product suite that quickly pinpoints performance issues and presents all diagnostic information through an efficient, repeatable workflow can immediately improve service delivery and staff efficiency. For this reason, NetQoS combined the global, application response time monitoring and automatic investigations of SuperAgent with direct drill-down offered by GigaStor™, the long-term packet capture and retrospective network analysis module of the NetQoS Performance Center.

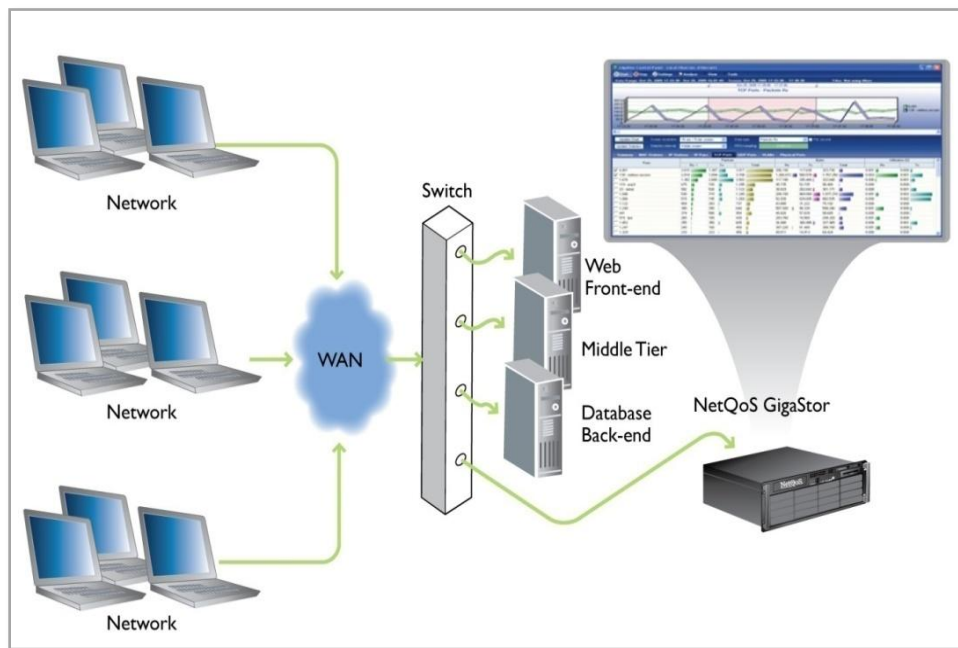
The intelligent baselines and thresholds in SuperAgent identify the relevant packet-level decodes so network engineers can immediately drill down to the associated packets on the GigaStor appliance to pinpoint the causes of performance issues instead of wasting time having to recreate problems and capture the right packets. There's never a need to deploy additional probes or wait for problems to reoccur—simply clicking the automatic investigation links reveal the relevant packet-level diagnostics and expert root cause analysis.

GigaStor provides a comprehensive selection of packet-level analytics needed to diagnose even the most complex performance issues. Multi-hop analysis tracks conversations with nanosecond resolution to help

troubleshoot issues caused by network congestion, fragmentation, and packet loss. Connection Dynamics views show system-wide conversation details that pinpoint the sources of latency and packet loss. Stream reconstruction views can rebuild complete Web pages with graphics and can reconstruct emails, documents, and instant messages for forensics, compliance, and security auditing. Finally, voice and video over IP application analytics provide detailed communication diagnostics and performance metrics for rapid troubleshooting.

GigaStor displays application-specific requests, identifies failed transactions, and pinpoints the sources of delay. Expert analytics report transaction details for hundreds of protocols including SQL, MSRPC, Citrix, POP3, SNMP, SMTP, Oracle, HTTP, FTP, DNS, DHCP, Telnet, SMB, and VoIP so you can determine the root cause of performance issues faster and reduce MTTR.

Figure 8



GigaStor collects and preprocesses all network packets for itself and SuperAgent, reliably capturing traffic from saturated gigabit and 10 GbE links while supporting up to eight monitored ports.

Summary

IT organizations can no longer manage networks in isolation from the applications they support. Traditionally, IT staff built their network management practices around infrastructure availability and fault management. Today, most networks are available more than 99 percent of the time and increasing user expectations for fast, trouble-free networked applications requires a shift from a device-centric to a performance-centric focus.

For Performance First to be viable, a rich set of performance metrics must be gathered and analyzed because today's complex networks present multi-dimensional challenges. Real-time metrics are required for day-to-day operational visibility into application delivery while historical data and trends are needed to support infrastructure investment decisions. Latency-sensitive voice and video traffic must coexist harmoniously with data applications but the performance metrics and methods of capture are quite different. Sophisticated analyses of time-synchronized, multi-source data are required to identify abnormal behavior patterns signaling rogue activity or inappropriate use of the network that would go undetected by conventional SNMP-based monitoring. Summary data is preferred for management and planning reports while network flow forensics and retrospective analysis of full packet payloads are often needed by network engineers to pinpoint the source of application problems.

Collecting, analyzing and reporting this wide array of performance data may seem daunting, but at NetQoS, delivering the fully integrated suite of management products to enable customers to embrace a Performance-First management approach is our mission. For more information, please visit us on the web at www.netqos.com or call 1.888.835.9575.

About NetQoS Inc.

NetQoS provides network performance management software and services that improve application delivery across the world's most complex networks. More than 1000 service providers, government agencies, and large enterprises – including half of the Fortune 100 – use the NetQoS Performance Center to monitor application service levels, troubleshoot problems quickly, and plan for change. Representative NetQoS customers include Chevron, Lockheed Martin, Reuters Group plc, American Express, Siemens, Boeing, Deutsche Telekom, NASA, and Barclays Global Investors. Headquartered in Austin, Texas, NetQoS has R&D centers in Austin and Raleigh, N.C., and regional sales offices in London and Singapore. For more information, visit www.netqos.com or call 877.835.9575

NetQoS Global Headquarters

5001 Plaza on The Lake
Austin, TX, 78746 United States
Phone: 512.407.9443
Toll-Free: 877.835.9575
Fax: 512.407.8629

NetQoS EMEA

1650 Arlington Business Park
Theale Reading, RG7 4SA United Kingdom
Phone: + 44 (0) 118 929 8032
Fax: + 44 (0) 118 929 8033

NetQoS APAC

NetQoS Singapore Representative Office
Level 21, Centennial Tower
3 Temasek Ave., Singapore 039190
Phone: + 65 6549 7476
Fax: + 65 6549 7001

Website: www.netqos.com

E-mail: sales@netqos.com

© 2009 NetQoS, Inc. All rights reserved. NetQoS, the NetQoS logo, SuperAgent, and NetVoyant are registered trademarks of NetQoS, Inc. in the United States and in other countries. ReporterAnalyzer is a trademark of NetQoS, Inc. Other brands, product names and trademarks are property of their respective owners.