White Paper

centricity™ SOFTWARE

Maximizing Application Service Quality for Today's Distributed Network Environments

Moving Beyond Network-based Quality of Service (QoS)

Executive Overview

Although incredibly powerful as business platforms, the Internet and Internet Protocol (IP) networks are inherently "best effort" delivery systems. They lack the consistent and reliable performance required for the effective delivery of business applications. This is especially true in today's environments where customers and users access a variety of applications over hybrid local and remote networks.

Vendors have sought to improve Quality of Service (QoS) for these networks by adding bandwidth management features such as packet prioritization (802.1p/Q, DiffServ) and bandwidth reservation (RSVP) to their networking hardware products, or by introducing new QoS appliances. While such network-centric approaches can improve the delivery of network ***packets,*** *limitations in today's devices clearly demonstrate that adding network hardware is not enough—application service quality is more than just a network problem.*

Solving today's application delivery issues means addressing and answering real world challenges at a system level. These challenges include the unpredictable nature of the Internet, the encryption of network traffic, the small installed base of hardware that supports QoS, and the need to interact directly with users and customers to increase productivity, satisfaction, and ultimately, profitability. To be successful, application delivery systems must extend beyond the edge of the Local Area Network (LAN) to include the user. It is only at the end user that the success of an application can be measured. In other words, only by joining networks, applications and users together into a cohesive system can organizations truly deliver a high level of "Application Service Quality."

Centricity Software's Centerwise™ is a unique user-centric system that optimizes the application delivery process to provide the highest possible service levels to users and customers. It works with new or existing network applications using standard technologies and leverages current network infrastructure. Unlike traditional solutions that prioritize network packets after congestion has already occurred, Centerwise™ works at the user's desktop to:

- *Dynamically* ***prioritize user and application traffic*** *based on business objectives and policies, and*
- *Proactively* ***communicate network conditions to users*** *for increased productivity and reduced frustration*

Compatible with all major policy enforcement and bandwidth management standards, Centerwise addresses application issues directly at the client, and fuses the user, the applications, and the network together into a ***system.*** *With Centerwise, organizations can for the first time deliver consistent and reliable Application Service Quality (ASQ), even across remote, heterogeneous and hybrid networks, such as the Internet.*

This white paper identifies the current application delivery methods, discusses real-world challenges and solutions, reviews the current network-based product offerings, and outlines Centricity Software's new approach to delivering Application Service Quality.

New Application Delivery Models — New Challenges

Today's applications run over a dynamic mix of well-controlled enterprise networks, secure VPNs, and the wide-open Internet. Business and Information Technology (IT) managers are currently working hard to determine the combination of enterprise and service provider resources that best meets the requirements of their applications and users.

Businesses, driven by competition to be more cost-effective and responsive to customers, are turning to application delivery over 'hybrid' networks. Hybrid networks, in the form of intranets, extranets, and the Internet, combine both private and public infrastructure to deliver a greater degree of design flexibility, and at the same time reduce management costs through outsourcing.

Using these hybrid networks for application delivery creates risks and challenges that businesses must address. Outsourcing the network means giving up direct control, without giving up the need for reliability and predictability. The expected service level is not reduced simply because applications are now delivered over service provider and outsourced networks. In fact, because of the unpredictable nature of hybrid networks, implementing application control and management is more important than ever. IT and business managers must find ways of leveraging the public infrastructure to reduce costs while maintaining, or even improving, the service levels experienced by users and customers.

Application Delivery in the Real World

Successfully supporting business objectives and consistently delivering high levels of application service is a challenging, and *critical,* endeavor. More important than creating a patchwork of partial solutions is considering application delivery from a systems perspective. Rather than shifting bottlenecks from one place to another, managers must answer the question: "What are the key elements needed to create an effective solution?" The resulting system must improve the overall business results, and deliver the service levels required by applications, users and customers.

Some Quality of Service (QoS) management techniques, like packet prioritization, can provide incremental performance improvements, but still not satisfy the bandwidth needs of demanding applications such as Voice over IP (VOIP) and streaming media. Application delivery management should also coordinate with network and server management systems. For example, it makes little sense for a 'QoS' solution to give priority handling to packets if they are being sent to a server that is already overloaded. The effectiveness of any Application Service Quality system will be largely dependent on the ability to successfully address the following real world challenges:

- **Hybrid Networks—**Applications are now commonly deployed over heterogeneous networks that are actually combinations of both enterprise and service provider networks. In these application delivery environments, no single organization owns the entire end-to-end network. Consistent

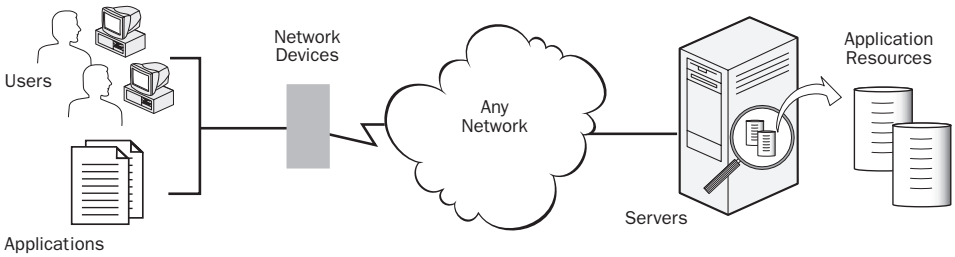

Improving Application Service Quality requires a user-centric approach that extends all the way to the desktop.

Today's Application Delivery Models

Intranet Model	Traditional method of delivering applications; the enterprise owns and controls all resources.
Capabilities:	Effective for supporting employees at the corporate headquarters and remote offices.
	Secure, private environment.
	Enterprise controls end-to-end performance.
Limitations:	Very costly to connect business partners, customers, and remote/mobile employees.
Extranet Model	Extends the intranet model by enabling user access via the Internet, servers and network services are often outsourced.
Capabilities:	Cost effectively extends access to business partners, customers, and remote/mobile employees.
Limitations:	Internet access raises security issues.
	Erratic and unpredictable Internet performance.
Internet Model	The model on which most new business-to-business and business-to-consumer applications are based. Most elements are outsourced to service providers.
Capabilities:	Fast deployment.
	Low-cost Internet communications.
Limitations:	Internet access raises security issues.
	Erratic and unpredictable Internet performance.
	Enterprise has no direct control over most network and application resources.

service level management is difficult even within a single network, but with multiple networks it becomes almost impossible using network-centric solutions alone.

- **Encrypted Data and Remote Users—** Virtual Private Networks (VPNs) continue to proliferate rapidly as organizations deploy them to secure and protect the data of their increasingly mobile workforces. When VPN traffic is encrypted using techniques like IPSec, network-centric bandwidth management approaches may lose part of their functionality. Encrypted traffic prevents the use of the OSI Layer-4 and Layer-7 packet information necessary to prioritize traffic based on application. Traditional management solutions can still prioritize encrypted packets based on destination, but cannot offer the critical benefits of prioritizing by application. As the need for encrypted LAN traffic increases, IT professionals must find solutions that offer application-based prioritization even in secure environments.

- **Avoiding Infrastructure Upgrades—**A major barrier to the widespread deployment of many QoS technologies is the need to replace or upgrade the networks, the applications, or both. For example, most routers and switches currently deployed do not support key QoS management standards such as DiffServ and RSVP. To implement either solution, these devices must be upgraded or replaced. Likewise, few currently deployed network applications, either commercial or custom, are enabled to take advantage of these emerging QoS standards. To support network based QoS, these applications must either be re-written, or replaced.

Network managers are all too familiar with the cost of upgrading an installed base of technology. The ideal solution must minimize changes required to the existing hardware and software, and avoid imposing any additional or ongoing maintenance costs.

- **Keeping End Users Informed—**Even with the most effective application delivery systems, users will experience problems. Users don't expect perfection, but they do expect to be notified and advised when exception conditions occur. Traditionally, a call to the help desk was the only course of action for end users. Today's solutions must be proactive, and reduce the burden on help desk staff, as user and customer populations expand exponentially and become increasingly mobile and remote.

The ideal solution must minimize the changes required to the existing hardware and software

The Current State of Quality of Service (QoS) Products

Using expensive WAN and Internet bandwidth efficiently and cost effectively is a major challenge. While network performance can often be improved by adding bandwidth, it is virtually impossible for an organization to provide enough bandwidth to guarantee that network contention will never occur. Instead, the availability and access to this expensive resource must be controlled and provisioned.

By default, IP networks are a best-effort packet delivery service that must delay or discard packets when the network becomes congested. Even well-engineered IP networks can experience congestion as networks support increasing numbers of users, applications become more bandwidth intensive, and usage patterns become more unpredictable.

Most current solutions are network-centric, hardware-based attempts to impose order upon chaotic network traffic by queuing and discarding packets. In effect, these solutions are simply another 'best-effort' component of the network and can sometimes even contribute to bandwidth problems. The fundamental assumption of the 'best effort' approach, that network chaos is inevitable, has resulted in efforts to mitigate, rather than prevent, traffic congestion.

Two major categories of QoS products have emerged to deal with bandwidth management and network performance issues.

Monitoring and Analyzing Products

These products monitor the traffic within networks, and in some cases, collect end-to-end performance metrics such as application response times and throughput. These products identify application flows based on layer-2 through layer-7 criteria. They also produce reports that show the amounts of bandwidth being used by users and applications. Monitoring and analysis products simply identify problems, rather than solve them. Because they cannot affect bandwidth usage, they should be viewed as just the first step in managing QoS.

Control Products

The second category of QoS management products goes beyond monitoring to exercise control over network performance. They use a variety of techniques to control and prioritize packet flows when bandwidth demands exceed the available supply. These solutions employ bandwidth management techniques such as packet prioritization and bandwidth reservation. They are typically implemented in routers, switches, and bandwidth management appliances. Because they are implemented as network devices, these products have only a partial view of the overall application delivery environment. Their network-centric approach to QoS management controls packets already in the network, but fails to address application delivery at a system level that includes the user and application. And, control is limited when data encryption is used in the network or VPN.

To more fully understand the present state of QoS technology, it is useful to review the control-based products currently on the market. We can further define these approaches as:

- Router and switch-based prioritization
- Traffic shaping using 'QoS appliances'
- Bandwidth reservation
- Content management

Approaches to Improving 'Quality of Service' for Applications

Router or Switch-based Data Prioritization (802.1p, DiffServ)

Mechanism: Queuing based on predetermined algorithms, weighted fair queuing, or class based queuing. Generally operates on packets that have been previously marked for priority.

Pros: Leverages existing hardware

Effective when congestion is low or infrequent

Cons: Manages congestion in a 'reactive mode.' Significant congestion causes packet loss and retransmissions. (Degradation for TCP/IP apps will be more severe than for UDP/multimedia applications.)

Difficult to configure and administer in multi-vendor or multi-version environment, especially as new applications are added or network dynamics change.

Application recognition is possible, but limited due to overhead at a single point. Typically requires hardware and/or software upgrades.

Provides point-to-point, rather than end-to-end prioritization. Limited in encrypted environments to Layer-2 and Layer-3 prioritization.

Bandwidth Management (QoS) Appliances

Mechanism: Network device filters and shapes traffic based on packet type or other attributes. Marks packets with priority information, based on policy. Similar to router-based QoS.

Pros: Typically plug and go installation.

Simple solution for simple problems.

Cons: Except for TCP traffic, provides only reactive management.

Additional point of failure, potential bottleneck.

Requires vendor updates to support each new application and protocol.

Limited in encrypted environments to Layer-2 and Layer-3 prioritization.

Router or Switch-based Bandwidth Reservation (RSVP)

Mechanism: Application initiates reservation request. Request is routed from end-to-end.

Pro: Could become the de facto standard.

Cons: End-to-end implementations are required but rarely available, especially through the Internet.

May appear to be end-to-end, but devices which don't support the protocol will be silent, creating a 'false positive' reservation.

Limited in encrypted environments to Layer-2 and Layer-3 prioritization.

Content Management (Load Balancing, Caching)

Mechanism: Moves content closer to LAN. Improves availability of content.

Pros: Improves availability of content/traffic by moving content closer to LAN.

Cons: Does not improve (or address) bandwidth or delivery issues through the LAN, such as bandwidth contention.

Less appropriate for dynamic content or transactions.

System-wide Policy Enforcement (e.g., Centerwise™ system)

Mechanism: Control Points calculate and manage bandwidth allocations.

Agents receive and enforce policy on the client devices (PC's).

Pros: Manages resources across the organization, not just per user.

Can prioritize by application even in encrypted environments.

Decentralized approach scales very easily.

Effective with existing applications in existing networks.

Cons: Requires software load on client.

Targeted at B2B market, currently supports only Microsoft* Windows desktops.

Router and Switch-Based Prioritization

Prioritization within routers and switches involves two steps—packet classification and priority queuing.

Packet classification is based on information in packet headers or in the application payload. Layer-2 and Layer-3 header information, (usually MAC and IP addresses) can be used to identify traffic flowing between specific hosts within a network, and can also contain priority code points used to classify traffic. Many products support classification based on one or more of the industry-standard priority marking schemes, including IEEE 802.1p/Q, and Diffserv/IP ToS (Type of Service).

After classification, packets are placed in an output queue based on their relative priority or on the amount of bandwidth they are allocated. A queuing algorithm, such as weighted fair queuing or class-based queuing, then determines the transmission order of the queued packets.

Router and switch-based prioritization techniques can manage the low congestion levels encountered in LANs supporting only a small traffic load. However, such prioritization does not effectively address the high degree of congestion that occurs at network bottleneck points such as WAN links. And, it does nothing to prevent congestion, instead reacting to congestion only after it occurs. This results in increased latency as queue lengths increase. It can also result in packet loss and retransmission that delays traffic, reduces the efficiency of networks, and causes even more congestion.

Priority queuing has a very limited scope of operation. It is not an end-to-end bandwidth management technique, and only affects the performance of outbound traffic on individual ports of routers or switches. To achieve end-to-end bandwidth management with routers and switches, prioritization must be implemented at each network hop.

Consistent performance can only be achieved when each router or switch implements similar queuing mechanisms. The queuing mechanisms being implemented by various vendors, including weighted fair queuing and class-based queuing, differ significantly in their capabilities and effectiveness. Managers of multi-vendor networks should be aware that a single priority queuing technique might produce different results from vendor to vendor. And, even in single-vendor networks, implementation differences between hardware platforms and software release levels require frequent upgrades and re-configuration efforts.

Bandwidth Management Appliances

Self-contained bandwidth management appliances typically combine monitoring and bandwidth management capabilities with product- and protocol-specific policy management capabilities. QoS appliances are most often used at the LAN-to-WAN boundary, where network congestion most frequently occurs. Some appliances attach directly to WANs, but most are installed behind WAN routers.

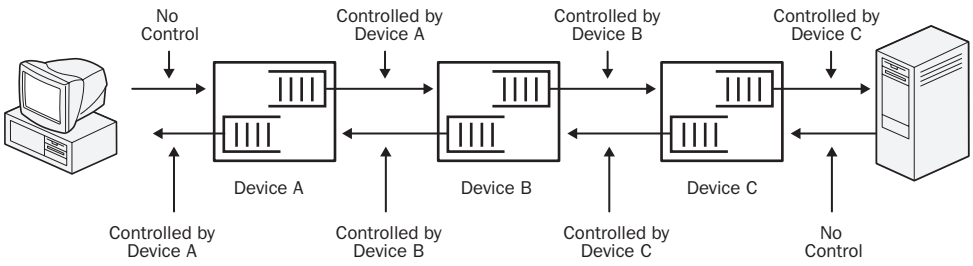

Router-based QoS is only effective on a per hop basis.

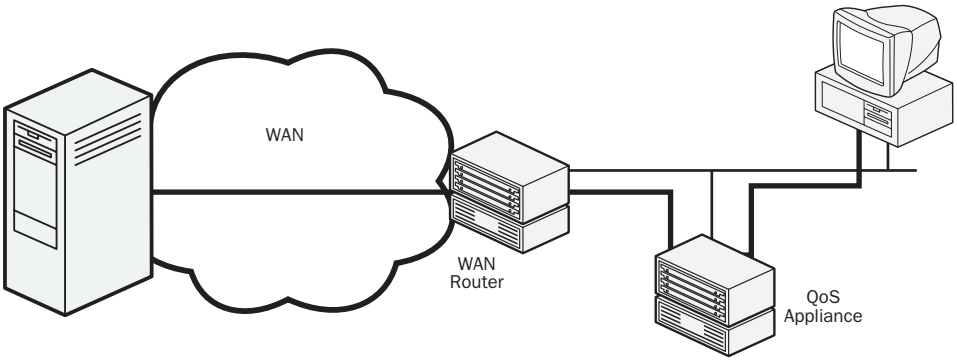

QoS appliances can become network bottlenecks under a heavy traffic load.

These QoS appliances typically use a traffic-shaping algorithm along with one or more priority queuing algorithms to control the rates of application flows and to allocate bandwidth.

These QoS appliances can operate across multi-vendor networks and generally don't require changes to existing infrastructure. A key limitation, however, is the lack of scalability. Because all managed traffic must flow through the appliance, these devices become bottlenecks when traffic levels rise. Ongoing maintenance is also an issue because they must be periodically updated to recognize new applications (such as Napster*-like variants), and new protocols. In addition, using a QoS appliance means adding a potential point of failure to the network.

Router and Switch-Based Bandwidth Reservation (RSVP)

Some applications, such as streaming media and Voice over IP (VOIP), are very sensitive to delays in delivery and require dedicated bandwidth to provide high quality sound and video. Bandwidth reservation addresses this requirement by setting aside the required bandwidth at the time that the applications are started. The industry standard for bandwidth reservation is the Resource Reservation Protocol (RSVP).

With RSVP, applications reserve bandwidth by sending a reservation request into the network. RSVP-enabled routers along the data path then reserve the amount of bandwidth requested for the duration of the connection.

RSVP is a potentially powerful bandwidth management tool, but has some significant limitations. First, most routers and switches currently deployed in networks do not implement RSVP. And, because RSVP increases the amount of processing overhead in routers and switches, its scalability is limited and virtually eliminates its usefulness over the Internet. Finally, RSVP signaling requires capabilities that very few applications currently possess. Few businesses or technology managers have the budget or time to upgrade networks and applications to support RSVP. Given these constraints, RSVP is best suited for use in combination with other Application Service Quality technologies.

Content Management Systems

Caching and load balancing are content management techniques that can have a positive effect on improving the delivery of relatively static content over the Internet.

Caching reduces the effects of network delay and outages by staging replicated content closer to the user. When a client retrieves data from a server, a caching device located near the client can cache the information so that subsequent requests for the same data are redirected to the cache rather than all the way back to the server. Requests are redirected by dedicated appliances or by redirectors built into routers or switches.

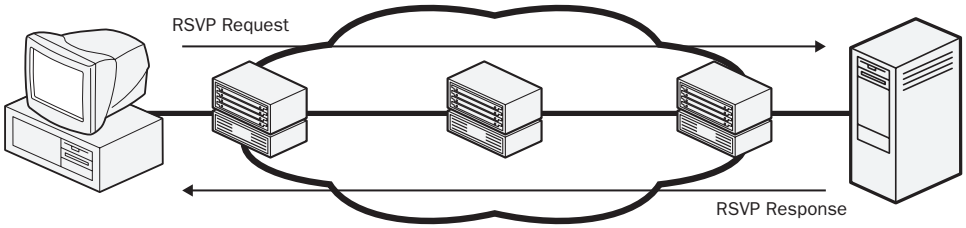

RSVP must be supported by all routers along the data path.

Load balancing works by optimizing the use of server bandwidth, redirecting service requests to the server that is best able to handle the request. Load balancing may be implemented in dedicated appliances or integrated into routers and switches. Both caching and load balancing redirect client service requests based on a variety of criteria including source and destination address, TCP/UDP port number, or application content. However, these products are limited by their inability to access the application content of sessions that are encrypted end-to-end (as with Windows 2000 IPSec). This can be a severe limitation when design criteria require caching or load balancing based on OSI Layer 7 elements, such as URLs or cookies. In addition, they do not address bandwidth contention and congestion.

While content management technologies can be useful for improving the delivery of certain types of content, they are not a complete approach to application delivery.

The Centerwise Customer-Centric Software System

The Centerwise software system takes a broad view of the entire application delivery process to provide the highest possible service levels, even in today's hybrid environments. It encompasses not only the technical aspects of application delivery and monitoring, but also provides the human components critical to the effectiveness of any network application. Rather than simply monitoring traffic or prioritizing packets, Centerwise fuses applications and users with networks to deliver assured application service levels. This approach addresses performance and bandwidth issues at the source-the users' desktop, and uniquely provides visibility from the user all the way through to the application and server, even across the Internet. This total visibility allows Centerwise to dynamically prioritize any application for any user based on available bandwidth, communicate real-time with users, and help resolve customer problems.

In addition to offering an immediately effective solution, Centerwise provides a forward-thinking, extensible design. With its open-ended architecture and application programming interfaces, Centerwise delivers a robust platform for customization and future functionality. As they become necessary, options such as server load balancing, caching and cache redirection, and usage-based billing and auditing for service providers can be incorporated by Centricity Software, its customers, and its business partners.

While network-centric management solutions operate within the network, Centerwise delivers a layer of services at the user's desktop, between the application and the network. This unique position allows Centerwise to manage how application traffic *enters* the network based on business-centric policies, rather than the traditional approach of trying to manage traffic already *on* the network. Using patent pending technologies and the industry-standard WinSock API, Centerwise:

- Monitors network activity, implements business-centric policies to improve application delivery, and adjusts dynamically to changing network conditions
- Works with existing applications and infrastructure today, while also complementing current network-based QoS efforts, for robust integration into any environment

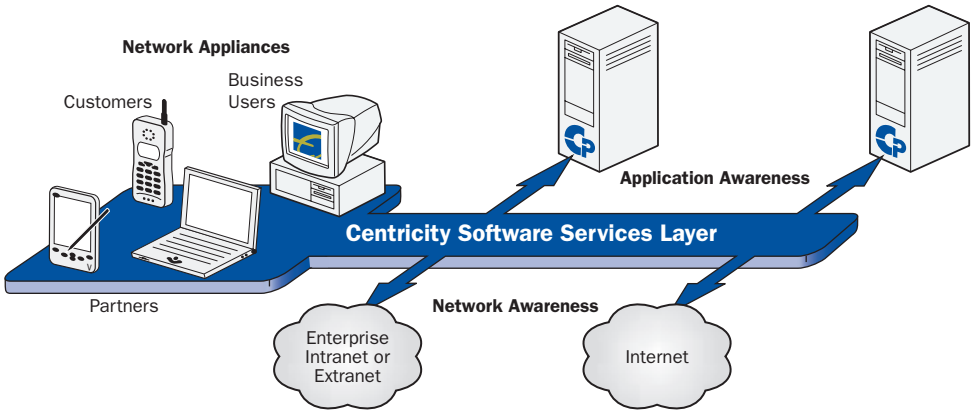

The Virtual Services Layer provides the basis for understanding applications and the network from the client's point of view.

- Delivers meaningful support services directly to the user or customer in real-time, using the Virtual Help Desk

Policy-based Bandwidth Management

Centerwise supports business-centric policies based on users and applications, and provides real-time adjustments to these policies in response to changes in the delivery environment. System administrators define priorities in a central location on the network. These policies are communicated across the network to cooperating Centerwise software agents residing on each client. Each agent then interprets and applies the policies and priorities locally. As network conditions and user behaviors change, Centerwise automatically adjusts bandwidth allocations to maximize utilization of available bandwidth.

With Centerwise, businesses have the flexibility to create the types of policies best suited to their environment. They can create user-specific policies that provide a single user, perhaps the CEO, with priority access to bandwidth. They can also build group policies that apply to a logical group of users, such as the accounting staff. And, they can build policies for a particular application or application group, limiting the amount of combined bandwidth allocated to FTP, E-mail, and HTTP, for example.

In addition to managing policies, Centerwise contains bandwidth tuning algorithms that dynamically adjust allocations in response to changes in the application delivery environment. Centerwise collects network usage data at the desktop, which tuning algorithms systematically compare with the initial bandwidth policies established by the system administrator. As changes in user behavior occur, Centerwise adjusts the original allocations, reallocating bandwidth from idle to busy workstations automatically. Valuable bandwidth never sits idle, and responses to network change take place in real-time, without requiring system administrator intervention.

A real-time traffic analysis and reporting facility automatically identifies users and shows their current bandwidth allocations. Business, application and network managers can use this information to help determine the policies needed to meet their performance goals.

Virtual Help Desk for Real-time Support

Today's demanding users want real-time feedback and support when problems arise. In the new application delivery models, these users might be demanding customers who require immediate attention, or mobile workers who are difficult to reach by conventional support methods. Meeting the support expectations of such users requires tools beyond the traditional help desk, and staying competitive means minimizing support costs whenever possible.

The Centerwise Virtual Help Desk provides a unique user support mechanism for the new hybrid network

The Centerwise Virtual Help Desk provides a unique user support mechanism for the new hybrid network. It detects changes in the application delivery environment that affect the availability of resources, and provides users and customers with useful, easy-to-understand information in real-time about these changes. The Virtual Help Desk identifies both local and remote network failures automatically, determining if a network link is broken, a server is down, or a particular application is unavailable–even if the problem is on the other side of the WAN.

After detecting a problem, the Virtual Help Desk distributes a message to affected users immediately, informing them of the cause of the problem, and offering suggested alternatives and workarounds until the problem can be fixed. To be able to address the site specific needs of users, these messages can be customized for each installation. The Virtual Help Desk operates automatically, requiring no effort on the part of busy network managers. And, the Virtual Help Desk can reach users and customers wherever they are located–on the local link, in a remote office, or across a B2B Extranet.

The Virtual Help Desk works proactively with users to keep them informed and satisfied, reduces the workload of the traditional help desk staff, and minimizes the impact of application delivery problems on profitability.

Value Propositions for any Environment

Confronted with still developing hybrid network models, and yet-to-be fully adopted QoS standards, network managers need application delivery solutions they can implement *now*. And, these solutions need to provide visibility into the application, even in IPSec and VPN environments. Centerwise is designed to serve as a stand-alone solution, as an immediately effective first step towards improved Application Service Quality, and as a user-centric complement to network-centric tools already in place.

As a stand-alone solution, Centerwise offers easy installation and immediate results today, using a standard API that already exists on virtually every network desktop. Through its use of the industry-standard WinSock API, Centerwise allows the application of powerful policies to control any standard application over any network. An “effective priority” algorithm works along with these policies to control the amount of bandwidth allocated to each user or application, even over networks that have no QoS management capability of their own. Centerwise is ideal for network managers that want results now, but don’t want to introduce non-standard control techniques onto their network, or adopt emerging standards that presently require costly upgrades.

Centerwise extracts additional value from existing solutions through its system level approach to application delivery. By working at the desktop, Centerwise provides greater management visibility into the application. Centerwise can classify and mark packets with DiffServ and 802.1p priority code points prior to their introduction onto the cable. Routers and switches don’t incur the overhead of packet classification, nor do they have to be policy-aware—an important consideration for networks that include legacy routers and switches. Similarly, Centerwise can perform RSVP signaling to reserve bandwidth on behalf of applications. Because the Centerwise agents, rather than the applications, initiate RSVP bandwidth reservation requests, existing applications themselves do not have to be modified. In IPSec and VPN environments, the Centerwise system provides a level of application visibility that traditional QoS solutions cannot offer. By applying control at the desktop, rather than on the network, Centerwise allows managers to integrate application-based policies into even secure, encrypted environments. Centerwise offers an ideal first step for businesses that may wish to implement DiffServ, 802.1p/Q, RSVP, or VPNs in the future, and represents a value-added extension for businesses already using these standards.

Summary

Today's network-based applications are being delivered by a new mix of enterprise networks and service offerings from ISPs, application hosting services, and ASPs. Management of the application delivery process is essential, but conventional management tools are not designed to fully address these complex and often unpredictable environments. The trend toward application delivery over hybrid enterprise and service provider networks requires a more comprehensive, more user-centric QoS strategy for Application Service Quality.

Centerwise offers a new customer-focused solution that optimizes application delivery to meet business objectives and improve user and customer relationships. Centerwise allows businesses to implement application management and control immediately using standard technologies, without the need to change or upgrade hardware or software. Centerwise leverages key network QoS technologies such as DiffServ and RSVP, and will continue to support emerging standards as they develop. With Centerwise, organizations can now effectively manage their environments even as they implement new application delivery models, and help ensure the success of their business as they move to the future.

Centricity Software

Centricity Software, based in Portland, Oregon, is developing and bringing to market software solutions to more effectively deliver and differentiate applications and services for corporations, B2B enterprises, and service providers. With the unique ability to control and resolve application and network issues based on business policy, paid level of service, or group and individual priority, Centricity Software's solutions put the customer first by providing users with the most effective personal application service available.

Visit us on the Web at

www.centricitysoftware.com

Centricity Software
4900 SW Meadows Road, Suite 400
Lake Oswego, OR 97035
tel 503.675.1200
fax 503.675.2796
toll free 888.675.3090

11/00 231004-001

centricity™
SOFTWARE