The 4 Essentials of WAN Optimization

How to Boost Network Application Performance across the Enterprise

PACKETEER®

The 4 Essentials of WAN Optimization

How to Boost Network Application Performance across the Enterprise

Accelerating WAN applications like SAP, Oracle, PeopleSoft, VoIP, and email is more complicated than just adding bandwidth, compression, and probes. The following handbook lays out four comprehensive steps to eliminate performance problems once and for all and optimize WAN application performance across the enterprise. We'll show you how to identify network problems, avoid common pitfalls, and select the best solution. And our easy-to-use checklist will help guide you through the process.

Table of Contents

Introduction

Welcome to *The 4 Essentials of WAN Optimization*. If you face network traffic bottlenecks, poor application performance, and compromised business productivity, then read on. Managing WAN application performance can be quite a challenge. Productivity drops and frustration mounts when performance turns inconsistent, unpredictable, and slow. Increased network demands can take performance from adequate to unworkable. A single event—deploying a new application or relocating servers—is all it takes to trigger a decline.

Common WAN performance issues include:

- Repeated bandwidth upgrades increase costs yet fail to permanently increase performance.
- Surges from recreational traffic cause urgent, interactive applications to struggle.
- Enthusiasm for VoIP (Voice over Internet Protocol) fades when callers routinely face stutter and static during peak network usage.
- Intranet applications at a main data center offer easy access but poor performance.
- ERP performance plummets whenever an employee syncs email at a branch office.

Sound familiar? You're not alone! Thankfully, this handbook lays out a four-part strategy to maximize application performance and optimize your WAN. With the right solution, you'll be able to:

- Determine which applications, users, and locations consume the most available bandwidth.
- Ensure network priority for critical applications, such as SAP, Oracle, Citrix, VoIP, and email.
- Block and contain undesirable recreational and malicious traffic.
- Apply compression and acceleration to maximize available bandwidth.
- Improve and protect the performance of urgent and critical applications.
- Contain important but less urgent traffic (such as large email attachments).

Defining the Problem

Recent changes in application and network environments have wreaked havoc on application performance. Increasing traffic, diverse performance requirements, and a capacity mismatch between local and wide area networks have prompted the decline in performance of business-critical applications. Traffic growth stems from trends in applications, networks, and user habits, including:

- **Increased application traffic:** An explosion of application size, user demand, and media richness.
- **Recreational traffic:** Abundant traffic resulting from recent trends in Internet radio, MP3 and P2P downloads, instant messaging, Web browsing, interactive gaming, and more.
- **Voice/video/data network convergence:** One network that supports voice, video, and data with their variety in bandwidth demands and performance requirements.
- **Webification:** Applications with Web-based user interfaces, typically consume 5 to 10 times more bandwidth than thick clients.
- **Distributed applications:** Enterprise applications that run over the WAN or Internet instead of being confined to a single location.
- **Server consolidation:** Combining data centers and reducing the number of application servers forces previously local traffic (high bandwidth, low latency, and low cost) to traverse the WAN or Internet (low bandwidth, high latency, and expensive).
- **Security:** Worms, viruses, and denial-of-service (DoS) attacks—the top source of network congestion (according to a recent Network World survey).

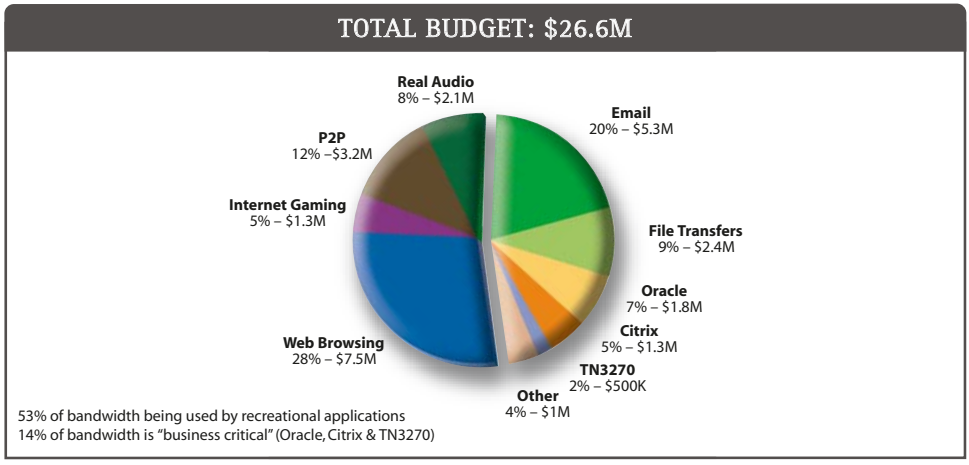

A 2003 IDC report found that the typical large U.S. enterprise spends $26.6 million annually on WAN circuits. This illustration shows a snapshot of one such company's top applications running across the WAN and the percentage of bandwidth each consumes—only $3.7 million supports critical applications while more than $14 million sustains recreational traffic.

Avoid the Common Pitfalls of WAN Optimization

In response to increasing network and business demands, many organizations are employing triage and stopgap-tactics often adding insult to injury. Here are just a few of the common pitfalls to avoid:

Give Me More Bandwidth

It's logical enough: To accommodate more traffic, just add more bandwidth. Yet all too often network managers expend much of their budget on bandwidth upgrades without the desired results. Without control over network resources, less urgent, bandwidth-intensive applications monopolize added resources and leave critical applications in the lurch. In essence, you're adding fuel to the fire instead of allocating those additional resources where they're needed most.

I'll Just Compress Everything

A compression-only approach doesn't fare much better. Without application-aware control features and proper management, compression may enhance the wrong applications—as in the case of just adding more bandwidth.

Probes Can Tell Me About It

In general, probes provide low-level network information about protocol underpinnings and connectivity troubleshooting but don't address the problem or help diagnose it. When flow information on a per-application or per-database basis is needed, a probe only delivers packet rates on a per-interface or per-device basis. In other words, you won't get the level of detail needed to differentiate SAP traffic usage rates, for example, or the ability to actually fix the problem.

I Can Piece It All Together

Pulling together a piecemeal solution ultimately adds more complexity and cost—especially if the components come from different vendors deploying different technologies. Often these technologies don't mesh or, together, don't add up to a workable, scalable solution. It's better to anticipate increasing network demands and plan accordingly.

1 Visibility—Gain Insight into your WAN
2 Control—Take Command of Network Resources
3 Acceleration—Keep Business Moving
4 Management—Manage from a Central Location

The All-in-One Advantage

Enterprise IT architectures are becoming more complex and are often globally distributed, increasing the pain of WAN application performance degradation and business disruption, according to a recent IDC report. The report went on to conclude that "The WAN optimization market is moving toward a more mature, complete application management solution that includes compression, QoS, monitoring, and IT service management."

Only a comprehensive solution—that monitors, controls, accelerates, and manages WAN application performance—can tackle today's demanding network and deliver top performance tomorrow. These four essential elements build on each other to deliver intelligent network management that proactively responds to—and heads off—present and unforeseeable network performance problems.

"Enterprise IT architectures are becoming more complex and often globally distributed, increasing the pain of WAN application performance degradation and business disruption...The WAN optimization market is moving toward a more mature, complete application management solution that includes compression, QoS, monitoring, and IT service management."

IDC Worldwide WAN Optimization Management 2004-2008 Forecast—Nov 2004

THE 4 ESSENTIALS OF WAN OPTIMIZATION

The 4 Essentials of WAN Optimization build on each other to create an intelligent network that optimizes WAN application performance at every location and anticipates what may arise now or in the future.

- **Gain *Visibility*:** Identify which applications traverse the network, what portion of the network they consume, how well they perform, and where delays originate.
- ***Control* Bandwidth:** Policy-based bandwidth allocation manages application performance over the WAN and Internet. Flexible control policies protect critical applications, pace greedy traffic, limit recreational usage, and block malicious traffic.
- ***Accelerate* Traffic:** Get more out of your bandwidth by allowing more data to flow through constrained WAN links, freeing bandwidth for the critical applications that need it most. Also push more traffic over long, high latency links (e.g. satellite) to fully utilize links.
- ***Manage* Centrally:** Consolidate reporting and management to a central location, using the reporting and policy management system of your choice.

1 Visibility

Gain Insight into your WAN

Before you can take steps toward application performance optimization, you must understand application traffic running across the WAN and how network resources are affected. Visibility into your network is a crucial component to gaining this invaluable insight. Yet do you really know, for a fact, about everything that is running on your network? Most IT managers do not, due to the complexities of network traffic.

Typically, WAN visibility has been provided by network monitoring tools that focus on Layers 2-4 and generate statistics, graphs, and reports showing network utilization, top talkers, top listeners, etc. This information is definitely useful, but it does not provide a granular view into application-layer traffic. It doesn't reveal who is using which applications, and how much bandwidth is being consumed along the way.

Visibility lets you:

- Automatically detect and classify hundreds of business and recreational applications.
- Identify top applications, users, servers, branch offices, and Web destinations.
- Analyze bandwidth usage, response times, the impact of configuration changes, and sources of delay.
- Track response times and break them down into time spent on the network and server.
- Set standards for service levels and track their compliance.
- Monitor conditions of interest, then, when thresholds are crossed, automatically take action to correct, document, and/or notify someone of the problem.
- Measure, graph, and/or export more than 100 metrics describing usage, availability, efficiency, response times, errors, and diagnostics.

What You Don't Know Can Hurt You

What you *don't* know about your WAN bandwidth can hurt you. When true application-level classification and monitoring capabilities are deployed on WAN links, many IT organizations discover for the first time that a significant chunk of their total WAN bandwidth—often more than 50 percent—is being consumed by non-business, recreational traffic (peer-to-peer, Internet videos, etc.).

The need to know will become increasingly important as Web-based applications increase—making it necessary to distinguish between each application type and performance characteristics as well as manage the high bandwidth needed by business applications. Simple IP address or static port schemes fall short. Port associations are rendered ineffective when trying to identify applications by monitoring TCP ports—such as port 20 for FTP traffic, port 25 for SMTP, and port 156 for SQL databases. For example, port 80—typically associated with HTTP Web traffic—can also be used for streaming media, P2P music downloads, or more. Even if traffic is HTTP...is it Siebel? Oracle? Casual Web surfing?

Get Down to the Application Level

To address these challenges, visibility must be granular especially at the application level. Rich traffic classification provides deep packet inspection of inbound and outbound traffic, detects dynamic and migrating port assignments, differentiates applications using the same port, and uses Layer 7 application indicators to identify applications. The result is a better understanding of bandwidth use and application performance by specific applications, servers, office locations, and users.

What Do You Know about your WAN?

- Which applications are running on your wide area network? Which use the most resources?
- How much of your bandwidth budget is consumed by critical versus recreational traffic?
- Do applications meet your users' expectations? Do they meet committed service levels?
- Which users and which branch offices are the top consumers of a particular application?
- How does each MPLS class of service perform? Does performance match service class and cost?
- Do remote locations get all the bandwidth they pay for? Do they need all of it? Are they using it efficiently?

TOP 10 CLASSES

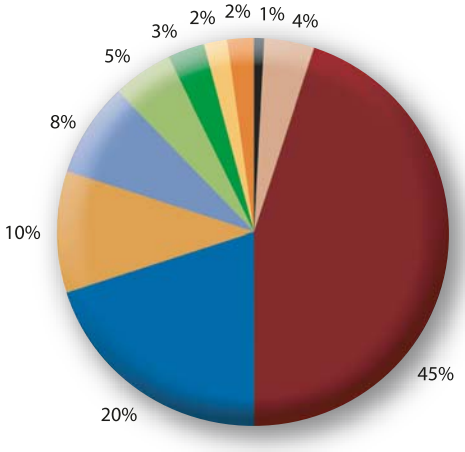

CLASS NAME	AVERAGE RATE (BPS)	(%)
1./Inbound/BitTorrent	230.4K	45
2./Inbound/WebBrowsing	102.4K	20
3./Inbound/DCOM	51.2K	10
4./Inbound/FTP	40.9K	8
5./Inbound/SMTP	13.4K	5
6./Inbound/POP3	25.6K	3
7./Inbound/Ares	10.2K	2
8./Inbound/Citrix	10.2K	2
9./Inbound/SAP	512K	1
All other classes	20.5K	4

Top 10 applications consuming bandwidth over a one week period.

Take Command of Network Resources

Once granular visibility is achieved, then the intelligence is there to allocate resources according to business priorities. With the right set of traffic controls, the appropriate amount of bandwidth goes to the right application and the right user at the right time. Resource hungry P2P applications from music videos, to Internet radio, to general Web surfing—can be relegated to the end of the line while mission-critical applications—like SAP, Siebel, and Oracle move front and center.

Challenges of Controlling the WAN:

TCP Treats All Traffic Equally

The TCP/IP network is concerned with making sure packets reach their destination. TCP will use as much available bandwidth as it can to deliver its payload and retransmit packets that may have been lost along the way—actually contributing to WAN congestion. TCP also regards all traffic as being of equal value—creating open competition for WAN bandwidth. And TCP goes all out and then backs off—causing drastic traffic fluctuations.

UDP Doesn't Manage the End-to-End Connection

Unlike TCP, UDP sends data to a recipient without establishing a connection and attempting to verify that the data arrived intact. Because UDP doesn't manage the end-to-end connection, it doesn't get feedback regarding real-time conditions, and it can't prevent or adapt to congestion. Therefore, UDP can end up contributing significantly to an overabundance of traffic, impacting all protocols, including UDP and TCP. In addition, latency-sensitive flows, such as VoIP, can be delayed and rendered useless.

Packet Markings Don't Promote End-to-End QoS

Packet marking techniques using class-of-service or type of service (CoS/ToS) standards take a proactive step toward preventing problems. However, despite markings that help determine which types of traffic take precedence, they don't offer rate guarantees or granular associations with users and hosts that promote end-to-end QoS. Packet markings are often translated into MPLS (Multi-Protocol Label Switching) tags and marked as "premium" traffic when transported across the carrier backbone. But the primary performance problem is usually encountered at the LAN to WAN connection before connecting to the carrier network.

Packet Queuing—Too Little Too Late

Packet queuing leverages the processing power of routers to employ various standards-based queuing techniques, with limitations. "Queues" by definition involve packets waiting in line, resulting in delays, and worse, dropped packets. Dropped packets result in retransmissions, which waste more bandwidth and further degrade network performance. Queuing is reactive—implemented only after WAN congestion occurs. And queuing requires the router to take on additional CPU-intensive packet processing, placing a significant burden on an already overloaded device.

MPLS CoS is Not QoS

MPLS has become the standard for connecting distributed locations—adopted by organizations hoping to take advantage of different classes of service and ensure appropriate application performance. However, they often discover that placing key applications into premium service classes does not reap the expected benefits. The right traffic does not get placed in the right MPLS service class. Premium classes deliver sub-premium performance as they drown in copious non-urgent traffic; important applications are improperly assigned to only best-effort classes. Also, traffic gets hung up in a congested bottleneck just before each entry point to the provider's MPLS network with unmanaged traffic heading into a LAN (inbound) at an inappropriately high flow rate.

What Does Work: Application Traffic Management

Only applying a broad spectrum of tools and technologies can control application performance. Two of the key technologies are bandwidth management and QoS. Together, these technologies build a complete solution that includes explicit bits-per-second minimum and maximum bandwidth rates, relative priorities, the ability to precisely target the right traffic, inbound and outbound control, and features to address the deficits discussed in the sections on queuing and packet marking.

Control lets you:

- Protect the performance of important applications, such as SAP and Oracle.
- Contain unsanctioned and recreational traffic, such as P2P and Internet radio.
- Provision steady streams for voice or video traffic to ensure smooth performance.
- Stop applications or users from monopolizing the link.
- Reserve or cap bandwidth using an explicit rate, percentage of capacity, or priority.
- Detect attacks and limit their impact.
- Balance applications, such as Microsoft Exchange, that are both bandwidth-hungry and critically important, to deliver prompt performance with minimal impact.
- Allow immediate passage for small, real-time traffic such as Telnet.
- Provision bandwidth equitably between multiple locations, groups, or users.
- Monitor conditions of interest, then, when thresholds are crossed, automatically take action to correct, document,and/or notify someone of the problem.

3 Acceleration

Keep Business Moving

Visibility and control take you most of the way, and traffic acceleration technologies will get you there in a hurry. Instead of increasing bandwidth capacity of WAN links where congestion persists and performance is suspect, compression and protocol acceleration provide a powerful approach to boosting performance. Acceleration is an important tool for maximizing performance and value from existing WAN links.

The Best Algorithm for the Data

Today's best WAN compression solutions typically employ several different compression algorithms; each algorithm is best suited for certain types of traffic. An intelligent compression solution will automatically select the best algorithm for a specific traffic type and search for ways to optimize the packet stream using that algorithm. It will also recognize traffic types that won't benefit from compression, such as streaming media or encrypted data. When applied correctly, it is possible to get 10X compression on certain file types, resulting in greater bandwidth availability and accelerated information delivery.

Compression with Control

Despite its advantages, compression technology has limited value as a standalone approach and works best as part of a complete WAN optimization strategy, with visibility and control as the foundation. Compression reduces the amount of bandwidth required by a transmission, thus freeing resources for other applications. Without adequate visibility and control, it is impossible to know or influence how extra bandwidth is used—often it is quickly consumed by non-urgent or non-business traffic. This contributes to network congestion and application performance problems.

When compression is applied to a WAN that is already tightly monitored and controlled, the benefits are significant. Newly available bandwidth can be allocated based on IT's holistic view of users, applications, network resources, and business priorities.

Compression lets you:

- Enjoy compression gains of up to 10X without loss of quality or data.
- Increase capacity and direct bandwidth gains to critical applications.
- Ease congestion on a saturated WAN link.
- Postpone or avoid bandwidth upgrades.
- Eliminate the burden of having to define and maintain compression tunnels, used to shrink, transfer, and restore traffic.
- Customize compression techniques for individual applications.
- Streamline repeated data, shrink transfer size, and/or reduce the number of packets.

COMBINING CONTROL AND ACCELERATION

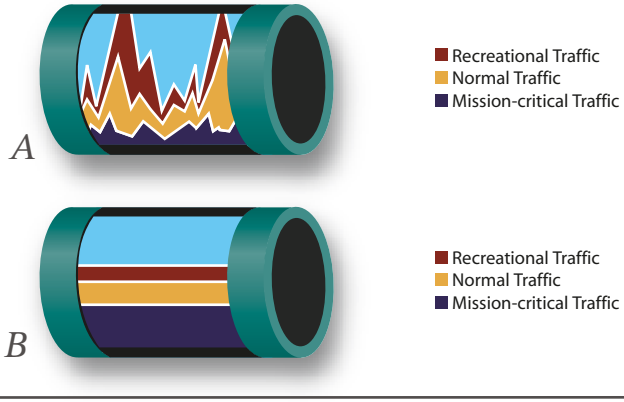

A. An unmanaged link shows that non–urgent traffic impacts the bandwidth available for critical applications.
B. A link with control (a managed link) ensures that mission-critical applications get additional capacity.

Get the Most out of your Connections

Protocol acceleration gets the most use out of your connections. It speeds up performance of transactions over WANs hindered by long distances and/or congestion. And it fully utilizes the entire link so that large files can be transferred quickly.

Overall, protocol acceleration enhances performance across the board. It:

- Speeds WAN delivery of Webified applications to the end user.
- Accelerates TCP over high latency links—ramping bandwidth and overcoming distance latency with Intercontinental WAN links, satellite links, and mobile data links.
- Saves bandwidth and response time by retransmitting lost data without penalizing other traffic.
- Fills pipes to full use and quickens critical, large flows when even very small latency amounts get in the way and high bandwidth flows cannot ramp up.
- Leverages a complete all-in-one solution to overcome unidentified performance problems, including poor performance of business-critical applications, MPLS deployments, VoIP projects, and infected PCs creating DOS attacks.

Acceleration lets you:

- Increase the amount of data carried over the link.
- Enjoy compression gains of up to 10X without loss of quality or data.
- Gain additional useable bandwidth.
- Customize compression techniques for individual applications.

4 Management

Manage from a Central Location

The advantage of a comprehensive WAN optimization strategy is being able to manage your system from a centralized location—leveraging powerful monitoring, control, and acceleration with centralized information collection, analysis, and reporting. For large deployments, in particular, validating enhanced performance by evaluating reports from each appliance can be challenging. The right centralized system can scale to large deployments of more than 1,000 units and provide central, fast, and flexible analysis of application performance across the distributed WAN network. By aggregating metrics and creating organization-wide reports on performance, you gain the insight needed to control network expenditures and understand network ROI.

Centralized management lets you:

- Deploy network-wide partitions and policies across your system.
- Define partitions and policies as a percentage of link size or explicitly as bps.
- Distribute software upgrades.
- View a summary of the status of all locations.
- Disseminate traffic classes for new categories of applications or traffic.
- Publish definitions of events of interest.
- Respond quickly on an enterprise scale to network or application performance changes.
- Minimize the administrative overhead and total cost of ownership.
- Configure multiple appliances simultaneously.

Powerful Analysis Tools

Centralized insight into application performance throughout an organization—particularly at bandwidth-constrained WAN links—takes a proactive step in aligning application performance with business needs. Application and network managers can gather, sort, present, analyze, and share information and bandwidth utilization response times, network efficiency, service-level agreement compliance, hosts, top applications, and new applications. Essentially, you can create and manage network-wide application performance policies with no more effort than it would take to set policies on a single device.

In Summary

These ***4 Essentials of WAN Optimization*** provide a holistic view of optimizing network, application, and server resources across the distributed enterprise. They provide a strong foundation for you and your IT organization to gain maximum performance and value from your existing WAN infrastructure—and better prepare you to leverage new technologies like VoIP and MPLS. Also, you should stand to benefit from substantial cost savings, due to avoiding costly and unnecessary WAN link upgrades, by making better use of the bandwidth currently supporting the business.

Choosing the right WAN optimization solution can be a daunting task, but it doesn't have to be. Look for these criteria, and you'll be on your way to applications that run at peak performance:

☑ **All-in-one solution**
Point solutions add complexity and cost. Besides, they may not resolve the problem; or they may only do so temporarily. An all-in-one solution leverages interoperability and builds to form an intelligent network approach.

☑ **Visibility into what's happening on your WAN**
You can't make smart decisions about optimization until you've got true insight into how bandwidth and network resources are allocated.

☑ **Deep application classification**
It's important to know what's happening at the application level and to be able to automatically detect and classify hundreds of business and recreational applications.

☑ **Bandwidth usage control**
Use explicit rate, percentage of capacity, or priority to allocate resources according to business priorities.

☑ **Protection of business-critical applications**
Contain unsanctioned traffic and allocate resources where they're needed most.

☑ **Monitor areas of interest**
Automatically take action when thresholds are crossed. Correct, document, or notify before congestion occurs.

☑ **Performance acceleration with control**
Use compression and protocol acceleration, based on insight not impulse.

☑ **Increased capacity gains directed to critical applications**
Customize compression techniques for unique applications or locations. Enhance bandwidth with compression ratios of up to 10:1.

☑ **Centralized management**
Minimize administrative ownership and TCO while also gaining holistic insight and control over the entire network from a single location.

☑ **Detailed analysis and reporting**
Gather, sort, present, analyze, and share information about bandwidth utilization, response times, network efficiency, service-level agreement compliance, hosts, top applications, and new applications across a system.

About Packeteer, Inc.

Packeteer®, Inc., (NASDAQ: PKTR) is the global market leader in Application Traffic Management for wide area networks. Deployed at more than 7,000 companies in 50 countries, Packeteer solutions empower IT organizations with patented network visibility, control, and acceleration capabilities delivered through a family of intelligent, scalable appliances.

10201 N. De Anza Blvd. Cupertino, CA 95014
Tel: 408.873.4400 • Fax: 408.873.4410
info@packeteer.com • www.packeteer.com

PART NO. 1007.A 04/05