

7

Steps

To WAN Optimization

*A resource for aligning WAN performance
with the needs of today's distributed enterprise*



Fast WAN. Fast Apps. Fast Business.

7 Steps to WAN Optimization

1. Treat Your WAN as a Strategic Business Asset
2. Discover Exactly What Traffic Is on Your WAN
3. Control and Protect Business Applications
4. Accelerate Business Traffic With Less Bandwidth
5. Strengthen Security with Traffic Management
6. Control the Forces That Impact VoIP Quality
7. Extend the Value of MPLS to the Edge of Your WAN

Introduction

Welcome to the Seven Steps to WAN Optimization. If you're reading this booklet, it is assumed you are an IT executive at an organization that depends on a wide-area network to connect users, systems, and applications across many different locations. We know that competition for your time is intense, so we've tried to keep this guide relatively brief, focusing on key issues and solution strategies. In Appendix A, we have listed additional resources for investigating any of the seven steps in more detail.

One last point before we get started. This booklet was

produced by Packeteer, Inc., the global leader in WAN application traffic management. But you'll notice that there is no mention of Packeteer products in any of these seven steps. Instead, much of the content here — some of which is just good common sense — is based on Packeteer's experience with thousands of customers worldwide. Think of this as a “best practices” guide that leverages the real-world experience of your peers in other IT organizations.

So, without further delay, let's get started on **7 Steps to WAN Optimization**.

Step 1: *Treat Your WAN as a Strategic Business Asset*

The words “WAN” and “strategic” are seldom used in the same sentence. That's because the WAN is often considered just another form of network plumbing. Very slow and very expensive plumbing, ordered from service providers that refer to your WAN links as “pipes.” So what's so strategic about all this?

Simply stated, your WAN is an incredibly strategic asset. It connects your business and supports the critical applications that power your business. Today's network-based economy has accelerated the evolution of the distributed enterprise — often comprised of a headquarters, several regional facilities, and dozens or even hundreds of smaller offices located around the world, all working together as one integrated business entity. The connectivity that enables users to share information and applications powering business is the foundation of the WAN's infrastructure.

Although the WAN's infrastructure is important, it's likely to be exponentially slower than your LAN infrastructure. For example, today's bandwidth-rich

LANs are often built with 100 Mbps or even Gigabit Ethernet, while most WAN links operate at T1 speeds (about 1.5 Mbps) or lower. That translates to a speed advantage of 60-100X or more for the LAN. This disparity often causes performance bottlenecks at the LAN/WAN boundary. Imagine 60 lanes on a highway converging into one lane without any warning — that's a problem your WAN has to deal with every day.

It stands to reason, then, that when your WAN doesn't perform well, your business does not perform well. Employee and overall business productivity suffers when remote users are forced to compete for WAN bandwidth and wait for access to business-critical applications. In fact, according to a recent Computerworld research project, 64 percent of multi-site enterprises surveyed rated application response time as their top WAN challenge.

In the past, the natural response to WAN performance issues was to “throw bandwidth” at the problem. That is perhaps the simplest attempt at a solution, but certainly not the most

effective. There are three critical problems with this approach.

First, in the absence of any intelligent controls, the amount of traffic traversing WAN and Internet links will always increase as available bandwidth increases. Eventually, performance starts to suffer, and you're right back where you started.

Second, TCP/IP treats all traffic equally, so ERP users will compete equally with Web surfers and P2P users for the added bandwidth with no guarantee of performance improvement for anyone. In fact, the most aggressive applications — which are normally the least business-critical — tend to consume the bulk of the additional resources.

The third problem is that "throwing bandwidth" at congested links results in significantly higher WAN service costs. As traffic consumes the additional resources and performance problems re-emerge, the demand for bandwidth will continue to increase.

IDC, a leading global IT research firm, discovered that many large enterprises already spend more than \$20 million each year on WAN circuits. In organizations of that size, even a 5 percent increase in WAN link upgrades translates to \$1 million each year in additional operational costs — with an unpredictable return on investment. WAN costs are inevitable, and while

it may not be practical to reduce these costs, it is certainly possible — and necessary — to control costs in the future with an effective WAN

optimization plan.

So this takes us to the first step toward optimizing the WAN — treat your WAN as a strategic business asset. As a starting point, that means the operation of the WAN should be aligned with the operations of the business. The issue of IT and business alignment should not be taken lightly. The Yankee Group, a leading IT research and consulting organization, recently acknowledged the critical importance and urgency of this issue, stating that, "CIOs

The operation of the WAN should be aligned with the operations of the business.

and IT executives who have not begun the process of correctly aligning the IT infrastructure with business processes and requirements will soon be at a competitive disadvantage."

But what does aligning IT (and the WAN) with business really mean? At the highest level, it means that IT organizations are working closely with the various lines of business to understand their performance objectives and priorities, then determine how best they can align their investments, operations, and performance with the strategic objectives of the company.

In terms of the WAN, it means that the applications supporting business operations and priori-

ties should be carefully controlled to ensure optimal performance and response time for users.

For example, if certain financial data must be captured and compiled at the end of every month, then specific remote users of ERP and financial systems should be given high priority and guaranteed bandwidth during the last few days of each month to ensure that transactions are completed on time.

This strategic approach to managing WAN resources will help optimize its performance and value to your business. This approach also lays the foundation for the next step in WAN optimization.

Step 2: Discover Exactly What Traffic Is on Your WAN

As a strategic business asset, the WAN exists to support the applications and associated information that power the business — everything from e-mail to e-commerce to ERP. Unfortunately, as noted earlier, WAN congestion frequently disrupts the performance of these applications. This issue was highlighted in the results of a recent survey Packeteer conducted with Network World

magazine. The survey revealed that more than 80 percent of large enterprises surveyed experienced application performance degradation

that undermined business productivity and customer service.

That's a very real and very significant problem. Equally disturbing is one of the primary causes of the problem — a lack of visibility into WAN application traffic. The Network World survey revealed that more than 75 percent of IT managers did not have adequate visibility into their WAN application traffic.

If WAN optimization is the goal, it must begin with having visibility into application traffic running across the WAN. It is important to understand how resources are used before taking steps to optimize them. 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.

Many IT organizations have tried to identify applications by monitoring TCP ports. Some of these ports are well-known — such as ports 20 and 21 for FTP traffic, port 25 for SMTP, and port 156 for SQL databases. But the evolution of the Internet and Web applications

More than 75 percent of IT managers lack adequate visibility into WAN application traffic.

has made simple port associations inadequate.

For example, port 80 is usually associated with HTTP Web traffic, but it could also be used for streaming media, peer-to-peer music downloads or more. Even if the port 80 traffic is HTTP, is it Siebel? Oracle? Or casual Web surfing? The need to know will become increasingly important for three reasons.

First, the number of Web-based business applications is constantly increasing, so distinguishing between each application type and its performance characteristics is imperative in ensuring that each application is managed appropriately based on its relative importance.

Second, many business applications consume up to 5X more bandwidth when deployed over the Web. Therefore, these critical applications — combined with other types of Web traffic — could create significant WAN congestion and performance degradation.

Third, and perhaps most important, what you don't know about how your WAN bandwidth is being used can also have a significant financial impact on your

business. For example, 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 30 percent — is being consumed by non-business, recreational traffic (P2P, Internet videos, etc.). In such cases, this means that 30 percent of an annual WAN services budget — for example, about \$6 million of the annual \$20 million budget we cited earlier — generates no revenue or value to the business. That is unacceptable.

To address these challenges, the corporate WAN should be equipped with tools that provide more granular visibility into WAN traffic — especially at the application level. This requires deep packet inspection of inbound and outbound traffic to create a more complete profile of WAN traffic from Layer 2 through 7. This visibility enables a better understanding of how much bandwidth is being used by specific applications, and what servers, office locations, and users they are associated with. Just as important, you can also discover how each application is performing, and whether the

cause of performance problems is the network, an application, or a server. Having this information prevents the finger-pointing that inevitably occurs between application and network groups when the source of a perform-

ance problem is unknown. But it all starts with visibility — a necessary foundation for effective application traffic management and a prerequisite to the next step toward WAN optimization.

Step 3: *Control and Protect Business Applications*

One of the points made earlier is that once IT organizations gain granular visibility into WAN traffic, they often discover that a significant percentage of their bandwidth is consumed by non-essential applications. This includes everything from P2P applications like KaZaA, to music videos, to Internet radio, to general Web surfing. During the past five years, this type of traffic has become problematic for many distributed enterprises, causing network congestion and decreasing user productivity. Looking forward, it's a safe bet that the problem will worsen in the absence of any WAN optimization strategy.

Unfortunately, there is a limited set of controls inherent in a TCP/IP network that can address this problem. In fact, TCP is less concerned with traffic control and more preoccupied with trying to make sure packets reach their destination. TCP certainly has good intentions; it will use as much available bandwidth as it can to deliver its payload and retransmit packets that may have been lost along the way.

But there are a few problems.

Because of its unwavering commitment to sending and resending packets, TCP is often a major contributor to the WAN congestion problem. Second, TCP regards all traffic as being of equal value — so unwanted recreational traffic gets the same treatment as business-critical applications. This creates an open competition for WAN bandwidth by any and all types of application traffic. It is costly and inefficient and does nothing to support the critical needs and priorities of the business.

A third problem is that TCP gradually sends bigger and bigger chunks of traffic until packet loss occurs. TCP will then back off, cutting in half the amount of information sent, but ramp up again when the opportunity arises. This creates drastic fluctuations in traffic, resulting in sub-par performance.

What is required is a way to regain control of the WAN and ensure that bandwidth is aligned properly with the priorities of the business. At the most detailed level, this involves allocating appropriate amounts of bandwidth to the right applications and the right users at the

right times to accomplish business goals.

Over the years, many technologies have emerged that have attempted to provide this type of policy control. For example, packet marking techniques can be employed using class-of-service or type-of-service (CoS/ToS) standards. These are useful because they take a proactive step toward preventing problems. However, while markings help determine which types of traffic take precedence on the carrier's networks, they don't offer much value in

terms of rate guarantees or 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 before connecting to the carrier network (see Step 7 for more on MPLS). Packet queuing is another com-

mon strategy for improving traffic control across the WAN. In this case, companies leverage the processing power of their routers to employ various standards-based queuing techniques. While these techniques can add value, queuing also has its limitations. One problem is that "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, by nature, is a reactive

approach that is implemented only after WAN congestion occurs. By then, users have already started to experience performance degradation.

Finally, queuing requires the router to take on additional CPU-intensive packet processing that places a significant performance burden on a device that already has a lot of responsibility.

What is needed is a specialized solution that offers a comprehensive set of traffic controls.

Queuing, by nature, is a reactive approach that is implemented only after WAN congestion occurs.

This includes the ability to assign priority to the types of traffic and applications that are most critical to the success of the business. Most likely, that means prioritizing critical back-office and customer-facing applications like ERP and CRM, over less-urgent applications like email and FTP. And it certainly entails placing a higher priority on business applications over Web surfing, music videos, and other recreational traffic.

However, effective traffic control does not end with prioritization — that's just the beginning. Prioritizing traffic on a crowded highway does nothing to relieve congestion, so effective traffic control also requires effective control of the bandwidth associated with these applications. This association is essential. For example, there are many types of ERP transactions that are time-sensitive and should be assigned not just a priority transmission status, but a guaranteed level of available bandwidth to ensure optimal response time whenever that type of transaction occurs.

But why stop there? To achieve WAN optimization, even more control may be appropriate in some cases. What if you could set a policy that assigned a spe-

cific priority level and a specific bandwidth level to a specific user at a specific location and time of the month to ensure that specific types of financial transactions occur on time, without disruption to business productivity? This level of control may not be needed, but the point is that it should be readily available in a purpose-built application traffic management system that can be deployed at critical links throughout the WAN. Such a system can also offer the same policy control to relieve or prevent congestion on the network. For example, policies can be implemented to limit the amount of bandwidth available for recreational traffic — or to block that traffic entirely. Or bandwidth for recreational traffic could be restricted until after business hours.

With the right traffic management system in place, businesses have total control over decisions involving bandwidth, applications, users, hosts, times, and locations. This flexibility gives you the power to effectively align WAN resources with the requirements of the business. It's another critical step toward WAN optimization. The question now is, what else can be done to further improve WAN performance?

Step 4: *Accelerate Business Traffic With Less Bandwidth*

Let's assume for a moment that you have applied the two previous principles of WAN optimization — gaining visibility into WAN traffic and applying strict controls on WAN bandwidth utilization. Now let's assume that, despite these improvements, your WAN is still not fully optimized from a performance standpoint. What next?

A logical reaction would be to increase the bandwidth capacity of the WAN links where congestion persists and performance is suspect. But that step may be

unnecessary. To further optimize the performance of critical applications over the WAN, consider the value of traffic acceleration technologies.

For example, increasing the speed of information delivery can be achieved through a common approach that has been around for years — compression. Already a commodity, compression is built into modern Web

browsers and is a useful tool for sending large files via email. But compression is emerging as an important tool for maximizing performance and value from existing WAN links.

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.

An intelligent solution 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. Without a doubt, compression can be a very effective component of a WAN optimization strategy.

Compression is emerging as an important tool for maximizing performance and value from existing WAN links.

Having said that, it is important to understand that compression technology has limited value as a standalone approach — despite claims made by some vendors. Compression must be applied in context of a complete WAN optimization strategy, and it must be applied only after issues of traffic visibility and control have been resolved.

Here's why. Compression reduces the amount of bandwidth required by a transmission, thus freeing resources for other applications.

In reality, applying compression alone is very similar to blindly throwing bandwidth at problems. Without first having adequate visibility and control, it is impossible to know or influence how

extra bandwidth is utilized. In many cases, extra bandwidth is quickly consumed by non-urgent or non-business traffic. This, of course, contributes to network congestion and raises the same problems that existed before compression was applied. In contrast, when compression is applied to a WAN that is already tightly monitored and controlled, there can be tremendous benefits. Why? Because the best use of the newly available bandwidth can be determined based on IT's holistic view of users, applications, network resources, and business priorities.

In that context, compression's benefits are sustainable, and significant savings are available by avoiding unnecessary WAN upgrades.

Step 5: Strengthen Security With Traffic Management

What does security have to do with WAN optimization? Well, actually quite a bit. As discussed earlier, in today's distributed enterprise, the WAN provides a vital communications channel for all users. If any WAN link is compromised by a security breach or attack, performance suffers and network integrity is jeopardized.

Many enterprise organizations implement robust perimeter security systems that consist of firewalls, intrusion detection systems, anti-virus solutions, access control systems, and more. Despite this defense, attacks still occur — thousands every day. Although most attacks are stopped at the border, some penetrate the perimeter and wreak havoc on business efficiency.

So let's assume an attacker has just discovered a hole in the perimeter and is about to flood the network with malicious traffic that will disrupt all IT systems and business as a whole. Let's also assume that this is a "zero-

day" attack; that is, it's a new virus for which no patch or security solution is immediately available.

Such attacks are often characterized by a significant increase in connections or attempted connections by one or more hosts. The attack may start slowly at first, but when enough machines

Policies can be implemented to constrain bandwidth availability for suspicious traffic.

become infected, the amount of traffic generated causes congestion at bandwidth-constrained WAN links. As a virus or worm

spreads, the resulting surge in traffic shuts down network services and corrupts vital business data.

When this happens, IT realizes that it's facing a full-fledged security breach. Plenty of questions emerge: Is this a known virus or a blended threat that uses multiple modes of infection? Is there a signature or patch posted? Which machines are infected? Can I shut down the source of the traffic? How do I get my network back on track, while trying to inoculate the infected hosts?

Although this scenario may sound hopeless, an effective WAN traffic management system can play an important role in containing the problem and protecting the network until a final resolution can be implemented.

The protection offered by a traffic management system is based on the visibility and control capabilities discussed earlier. In terms of visibility, any new, unidentified traffic can be automatically classified as such by the traffic management system. IT managers can quickly associate this new unknown traffic with data on network utilization, top users, top hosts, etc., to acquire a composite view of how a malicious application is impacting overall traffic flow and network performance.

Once the source of the problem is understood, traffic control comes into play. Earlier, we talked about creating and enforcing policies to ensure the optimal performance of business-critical applications. In a

similar way, policies can be implemented to constrain bandwidth availability for the suspicious traffic, limit its access to critical resources, or block it entirely. This capability enables effective containment of zero-day attacks and helps to ensure network integrity and business continuity until a final fix is available from a security vendor.

It's also worth noting that, in this particular example, the impact from a zero-day attack can be further minimized if the network is already properly configured with traffic management policies. With such policies in place, WAN resources are already carefully allocated and controlled to provide optimal performance of business-critical applications. This approach to WAN optimization not only ensures network QoS, but also provides a proactive approach to reinforcing an existing security architecture. Sound bandwidth control policies effectively starve the attack by leaving less bandwidth available for unsanctioned traffic.

Step 6: Control the Forces That Impact VoIP Quality

Of course, change is a constant — your work is never done.

Even now, new technologies are changing the way WANs are architected and managed.

Convergence, for example, is a term that's been discussed for nearly a decade, but now it's

taking hold as an increasing number of businesses begin to consolidate voice, video, and data traffic over a single IP

infrastructure. According to a recent META Group study, more than 40 percent of all enterprise customers surveyed will migrate to voice-over-IP (VoIP) by 2007.

On bandwidth-rich corporate LANs, introducing voice traffic is relatively risk-free. The challenge is to ensure that VoIP performs effectively across bandwidth-constrained WAN links.

A single VoIP call typically requires at least 30 Kbps of bandwidth for acceptable voice quality. This includes 8 Kbps for a VoIP packet using a G.729 codec algorithm, plus about 20 Kbps in IP overhead bandwidth. To determine total WAN capaci-

ty needed to support VoIP, estimate the number of simultaneous calls that may take place across each link.

Although it is possible that WAN links need to be upgraded, it is equally possible that current links are appropriately sized and

just need to be optimized.

Again, leveraging application-layer visibility into WAN traffic behavior could reveal ways to

limit recreational traffic and reassign bandwidth to VoIP, avoiding costly WAN upgrades in the process.

Compression should not be applied to VoIP because codecs used to transmit voice traffic over the WAN (such as G.729) already compress the packets. Additional compression could compromise voice quality.

However, compression could be applied to the IP headers associated with the voice packets, as well as other selected data traffic, freeing up more bandwidth for VoIP and providing another way to avoid unnecessary WAN link upgrades.

VoIP calls can be assigned dynamic QoS policies that are automatically implemented.

Beyond bandwidth requirements, there are three other network issues that must be addressed: latency — the end-to-end delay in delivering the voice stream from the speaker's mouth to the listener's ear (this should not exceed 150 ms each way); jitter — the unpredictable and variable delays in the delivery of each voice packet (acceptable jitter is in the 20-30 ms range); and packet loss — the dropping of voice packets caused by network congestion (it should not exceed 1-2 percent).

The good news is that these types of problems, including the congestion that can contribute to latency, can be resolved through the same traffic management technologies that are applied to critical data applications. In many ways, voice is

just another business application, and it deserves the same careful treatment as ERP, CRM, or any application that powers your business. So, for example, all VoIP calls can be assigned dynamic QoS policies that are automatically implemented when a call is initiated.

Likewise, when the call is over, the bandwidth can be re-assigned to another application.

The traffic management principles outlined here for VoIP readiness also apply to videoconferencing and streaming video applications (such as on-demand learning in which video is used). Adhering to these principles will help move IT a step closer to WAN optimization, even across a converged voice/video/data network infrastructure.

Step 7: *Extend the Value of MPLS to the Edge of Your WAN*

Multi-protocol label switching (MPLS) has been generating well-deserved hype over the past few years. Like VoIP, its adoption rate is growing among distributed companies that are looking for ways to control communication costs without compromising performance quality. In fact, according to a Webtorials survey of nearly 200 IT organizations, the departure from frame relay is accelerating rapidly and next-generation MPLS will soon dominate the WAN landscape.

Carrier edge routers are not well-suited to application-level traffic management.

MPLS is a value-added WAN service that uses explicit and pre-defined network paths to transport traffic across the carrier's complex backbone network. It's compelling to businesses because MPLS is intended to provide multiple classes of services (CoS) and true QoS in an IP network.

However, like any emerging technology, MPLS has its drawbacks. Two specific issues relate to WAN optimization.

The first involves understanding the limits of MPLS in terms of its breadth and depth. Most major carriers have spent a great deal of time and money preparing their massive networks to enable MPLS services. These networks typically provide national or international coverage, but still their breadth of coverage is limited, which means that their QoS services are also limited. This is because

MPLS networks do not extend to the boundary of a corporate network.

There is, instead, an intermediary access network that connects a corporate office to a MPLS service. It is over these "first mile" or "last mile" networks where traffic congestion is often a real problem — one that a MPLS service provider cannot resolve effectively.

Carriers try to manage the problem using edge routers to process traffic over the last mile. But these routers are busy doing routing and are not well-suited for application-level traffic management, so their

ability to apply granular controls and policies based on Layer 2-7 data is quite limited. Carriers are typically unable to identify specific types of traffic or manage on a per-session or per-user basis. That is why the "depth" of a MPLS service is also limited. But comprehensive visibility and control are essential to achieve WAN optimization and end-to-end QoS.

Fortunately, the problem can be solved with the right traffic management system installed on the WAN. Based on the visibility, control, and acceleration capabilities already discussed, it's a simple matter to minimize congestion from the enterprise edge over the access network with traffic-shaping policies and, at the same time, apply necessary DiffServ, ToS, and MPLS packet markings to ensure that traffic is mapped to the proper CoS across the carrier's MPLS network.

The second issue to be aware of with MPLS has to do with its CoS options. As noted earlier, one of the advantages of MPLS

is its ability to support multiple CoS options, each with different price/performance levels. You may prefer a "gold" service level for priority traffic such as VoIP or ERP and may prefer "silver" or "bronze" service levels for less-critical traffic such as email.

But there is a risk that the CoS you choose may not be the optimal choice for a particular application — it could be too much or too little. To know for sure, you must have visibility into the performance of the specific application in question and be able to measure its end-to-end response time. This information allows you to track end-to-end performance of critical traffic and determine whether you have applications mapped to the appropriate MPLS classes.

MPLS will likely be a dominant WAN technology for many years to come, but to ensure maximum ROI and WAN optimization, it is imperative that the issues of last-mile control and optimal alignment of applications and CoS are addressed effectively.

Summary: The Next Step Is Yours

These seven steps to 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.

In addition, by following these steps, it is possible to gain substantial cost savings. These savings typically come as a result

of slowing or avoiding costly and unnecessary WAN link upgrades by making better use of the bandwidth currently supporting the business.

Now that you've had a chance to think a bit more about WAN optimization — as well as its strategic importance to the applications, systems, and users throughout your company — the next step is yours. If you'd like to learn more about how to make WAN optimization a reality, please take a few minutes to review the resources available to you in Appendix A.

Appendix A: Additional Resources

In the interest of brevity, each of the Seven Steps to WAN Optimization was presented at a fairly high level. To learn more about any of these steps, visit www.packeteer.com. Or, if you prefer, visit any of the URLs below for additional information.

An Introduction to Application Traffic Management

<http://www.packeteer.com/resources/?docid=90>

1. Treat Your WAN as a Strategic Business Asset

<http://www.packeteer.com/resources/?docid=50219>

2. Discover Exactly What Traffic Is on Your WAN

<http://www.packeteer.com/resources/?docid=50020>

3. Control and Protect Business Applications

<http://www.packeteer.com/resources/?docid=50021>

4. Accelerate Business Traffic With Less Bandwidth

<http://www.packeteer.com/resources/?docid=50120>

5. Strengthen Security With Traffic Management

<http://www.packeteer.com/resources/?docid=50377>

6. Control the Forces That Impact VoIP Quality

<http://www.packeteer.com/resources/?docid=1863>

7. Extend the Value of MPLS to the Edge of Your WAN

<http://www.packeteer.com/resources/?docid=133>

About Packeteer, Inc.



Packeteer, Inc., (NASDAQ: PKTR), incorporated in 1996, is the pioneer and global leader in Application Traffic Management systems for wide-area networks (WANs). Packeteer systems consist of a family of intelligent hardware appliances, each configurable with patented software technology that delivers a comprehensive set of visibility, control and acceleration capabilities to enterprise customers and service providers.

The Packeteer system automatically identifies and classifies hundreds of specific types of applications and enables IT managers to partition WAN bandwidth and apply strict bandwidth management and QoS policies to the critical applications that power the business. Likewise, Packeteer empowers IT managers with the ability to limit or block malicious traffic. Once policies are in place, Packeteer also enables IT managers to apply state-of-the-art compression technology to “create” more

available bandwidth and accelerate information delivery across the WAN. Lastly, Packeteer provides system-wide management tools for policy configuration and control, comprehensive performance monitoring, and enterprise-wide reporting.

For enterprise customers, Packeteer enables IT organizations to optimize their WAN infrastructure by effectively aligning application and network resources with business priorities and providing measurable cost savings. For service providers, Packeteer systems enable the deployment of application-aware, managed network services that provide QoS and expand service provider revenue opportunities.

With headquarters in Cupertino, California, the company's products are sold through more than 100 resellers, distributors, service providers, and system integrators in more than 50 countries worldwide.

Contact Information:

Packeteer, Inc.

10201 N. De Anza Blvd. | Cupertino, CA 95014
408.873.4400 | 800.697.2253
info@packeteer.com | www.packeteer.com

Notes

Notes

©2004 Packeteer, Inc. All rights reserved. Packeteer and the Packeteer logo are registered trademarks of Packeteer, Inc. in the United States and other countries. All other company trademarks are the property of their respective owners. No part of this document may be reproduced, photocopied, stored on a retrieval system, transmitted, or translated into another language without the express written consent of Packeteer, Inc.

