

CLIENT Orbital Data

PROJECT White paper

OBJECTIVE

Create a white paper to present the business case and strategic advantage of automated WAN optimization

PROJECT OVERVIEW


Managing a wide-area network (WAN) is far more complex than managing a local-area network. Because internet protocols were designed many years ago, performance degrades rapidly when clients and servers are located hundreds or thousands of miles apart.

The obvious solution would be to add bandwidth, but without optimization, more bandwidth usually provides little improvement. There are several WAN optimization methods available, but these techniques require constant fine tuning based on current network conditions – adding higher cost and more complexity to network management.

Orbital Data offers a different approach – technology that automatically analyzes WAN traffic, and applies the best technique in real time without the need for fine tuning.

Because WAN bandwidth can represent one of an organization's largest ongoing IT expenditures, this automated approach offers a strategic advantage in improving productivity and cutting complexity.

This white paper explains what causes poor WAN performance, why adding bandwidth can make the problem worse, and presents the strategic case for automated WAN optimization.



.....
*See below to read this white paper,
A Strategic Case for Automated
WAN Optimization.*

A Strategic Case for Automated WAN Optimization

*How to gain faster throughput,
improve application response,
and simplify your infrastructure*

CONTENTS

Executive summary:

An improved WAN, but at what cost? 3

Your WAN:

A key strategic asset, yet so unpredictable 4

The networking dilemma:

How your WAN conspires against you 6

Adding bandwidth:

What telecom companies don't want you to know 8

Getting more from your network:

A menu of WAN optimization strategies 9

No more tuning:

An automated approach to WAN optimization 10

The Orbital Data AutoOptimizer:

Simple, scalable WAN optimization 12

About Orbital Data 14

Executive summary:

An improved WAN, but at what cost?

Your WAN may be a critical strategic advantage, but if poor throughput and slow application response causes user frustration and poor productivity, that advantage vanishes rapidly.

Add bandwidth? It's certainly tempting. But although it seems counter-intuitive, adding expensive bandwidth usually makes a bad problem even worse.

While some optimization solutions promise significant improvement, complex configuration challenges can leave you with a solution that requires major infrastructure changes and constant fine tuning. In fact, you may find yourself deploying several different optimization techniques to suit changing network traffic patterns – adding yet another level of cost and complexity.

Automation: The future of WAN optimization

In this white paper, you'll learn about an exciting new approach – *automated* WAN optimization that ends the need for continuous monitoring. This approach offers:

- **Complete transparency:** Drop this solution right into your existing network with no configuration issues.
- **No fine tuning:** Leverage the best optimization techniques to suit specific network conditions automatically in real time.
- **Impressive results:** Organizations now enjoy predictable network performance with up to a 200-fold improvement in throughput and application response.

Now you can get LAN-like performance from your WAN and take advantage of more predictable network administration at an ROI that makes good business sense.

Your WAN:

A key strategic asset, yet so unpredictable

A Forrester report, “The WAN Traffic Compression Market in 2005,” reveals that corporate WAN performance is one of the top five issues confronting enterprise IT infrastructure managers at \$1 billion-plus companies. That’s not surprising when you consider that WAN operations can easily represent a company’s most significant ongoing IT budget item.

A slow and sluggish WAN was once merely annoying. But in today’s global IT environment, predictable WAN performance is critical to meeting multiple strategic objectives that include:

More effective collaboration

If half of your engineering team is in the U.S., with the other half in India, poor WAN throughput and slow application response can add to product development delays. Even over a 45 Mbps link, the transfer time for a complex engineering design file can easily grow way out of proportion to your expectations based on performance over a LAN. That slow throughput can wipe out any advantage you planned for locating half of your team in India in the first place.

Enhanced digital asset management

Media companies and others generate large digital files, which works fine when everyone’s under the same roof. But today’s teams may be separated by thousands of miles.

How long will it take to send a 30-gigabyte file from Hollywood to New York? Who knows? It could take 30 minutes or 30 hours because file transfer times can be as unpredictable as the weather. But with more predictable WAN transfer times, organizations can enjoy more predictable scheduling.

Improved regulatory compliance

Protecting data is more than just a good business practice – it’s required by a complex and intricate set of regulations that make it difficult to remain in compliance when your network slows down.

A major financial institution, for example, might be required to regularly back up 30 gigabytes of financial and customer data from dozens or hundreds of locations to one central data center every business day. But with fixed nightly backup windows, slow throughput creates serious bottlenecks.

That means more risk. Even losing a few hours worth of transaction data can mean the loss

of hundreds of thousands of dollars in recovery costs and increased regulatory scrutiny – but those risks that can be mitigated by improved WAN optimization.

Simplified infrastructure management

Consolidating servers, eliminating multiple platforms, and centralizing applications can bring significant savings.

However, you may find that remotely hosted applications behave quite differently over your WAN than they do on your LAN. Poor application response time can quickly erode productivity gains and cost savings. But by eliminating application response bottlenecks, you can take full advantage of simplifying your infrastructure.

The networking dilemma: How your WAN conspires against you

It's tempting to view your WAN link as if it were merely some pipe in a plumbing system, but unlike a water faucet, WAN traffic does not flow from one end to another at predictable rates. While there are many reasons you might be experiencing unpredictable time-outs and sluggish application response, most of the time the challenges result from the nature of your traffic.

At one time, the bulk of WAN traffic was simply email and FTP file transfers. But WANs are now called on to support more sophisticated technologies such as remotely hosted web applications, video conferencing, and VoIP – any of which can impact performance for several of reasons:

Latency

Over your LAN, latency is measured in mere fractions of a millisecond. However, over a WAN link that spans halfway around the world, you could see latencies of 200 to 300 milliseconds or more. This causes havoc in application response time.

Applications that were designed with LAN-like latencies in mind behave very differently when exposed to the extended latency of a WAN link. For some applications, just remotely opening a file might require dozens of exchanges between client and server. And from a latency standpoint, adding bandwidth doesn't usually offer any improvement. It just results in wasting even *more* bandwidth.

The nature of TCP

TCP (Transmission Control Protocol) protocols date back to the early 1970s when email was the dominant application and WAN circuits were measured in mere kilobits. It was not designed for effective long-distance transfer of large files or for efficient use of remotely hosted applications.

Under TCP, data flow is confined to the endpoints of the network. But problems can occur inside the network – and with TCP, each endpoint has little visibility of link conditions. So TCP adjusts data flow very slowly. The sending node simply guesses how fast to send data, gradually increasing the rate until packets are no longer acknowledged, then backs off.

You can visualize this if you think about how you might add more hot water while in your shower. Because of latency (the time it takes for the hot water to get from your water heater to the shower head), you need to make gradual changes to avoid getting scalded.

This “slow start” method of ramping up transmission means that TCP problems increase with the square of the distance, so a 200-mile link will demonstrate only one-fourth of the performance of a 100-mile link. And a 400-mile distance offers just one-sixteenth the performance of a 100-mile link.

Congestion

Easing network congestion can clear the way for more effective throughput, but that doesn't always translate into increased speed. For example, while you could add a few lanes to a highway, the maximum speed of any car on that highway remains the same.

What's more, today's network congestion issues are far more likely to be a result of what's happening at your endpoints than from conditions inside the network itself. If you think about the frustration drivers face when ten lanes of busy traffic merge into one, you can visualize how network transition points can be vulnerable to congestion and cause throughput and application response issues.

Chatty applications

The degree to which applications send packets back and forth from endpoint to endpoint makes a big difference in network performance. In wide area networking, many small transactions consume a disproportional amount of resources and time as one larger transaction.

WAN latency imposes a heavy penalty on web applications that require dozens or hundreds of time-consuming round trips between endpoints to complete a requested operation. And if another bandwidth-hogging transaction happens to occur at the same time, response times multiply rapidly.

For example, if a large FTP transfer is initiated at the same time that a chatty application needs more throughput, the FTP transfer will tend to crowd out the application-related traffic. While more bandwidth might seem to be the logical answer, the issue isn't the size of your pipe, but whether you're filling it efficiently. If not, then adding bandwidth is not going to make you any better off. It may even make a bad problem worse.

Adding bandwidth: What telecom companies don't want you to know

Telecom companies are lining up to sell you additional bandwidth. After all, they stand to make many thousands of dollars more in revenue per month if you decide to upgrade.

If you call them, they might ask you how fast a connection you need, as in whether you need a 15 Mbps or 45 Mbps link. But what the telecom companies don't want you to know is that they don't really sell speed – they sell capacity. Because of inherent latency in a long-distance network, a packet of data won't travel more quickly from New York to London over a so-called “faster” connection.

For example, if the top floor of an office building is suffering from low water pressure, replacing the one-inch piping with two-inch piping isn't going to solve the problem. To increase pressure, you'll need a more powerful pump. Without it, the fattest pipe in the world won't make any difference.

Making a bad situation worse

Paradoxically, adding bandwidth can actually *decrease* the performance of your WAN – especially in the absence of any WAN optimization technology.

Without intelligent controls in place, WAN traffic simply expands to fill the available bandwidth. When throughput problems reemerge, you'll be back where you started – except that your monthly recurring costs will be a lot higher. And because some applications are notorious bandwidth hogs, they just consume more bandwidth whenever they see more of it available. Ironically, this just exacerbates the problem.

Think of it this way: If a highway administration adds two extra lanes to a four lane freeway, capacity theoretically increases by 50 percent. But what if drivers took note of the wider road and decided that it seemed open enough to straddle two lanes? If only a few drivers did so, the result would be only three effective lanes – creating 25 percent *more* congestion.

Before you add bandwidth, evaluate your WAN traffic to see if a simpler, less expensive WAN optimization technique will be enough to resolve throughput and application response issues. And if you do decide to add bandwidth, make sure you deploy the right WAN optimization techniques to make sure you don't pay for more bandwidth than you really need.

Getting more from your network: A menu of WAN optimization strategies

If adding bandwidth isn't the cure-all that most telecom companies would have us believe, then what techniques *can* be used to get LAN-like performance from a WAN link? There are many approaches available from a variety of vendors, but most rely on flow control, compression, or acceleration. These techniques can be useful, but reported problems include:

- **Poor application and/or network transparency:** Adding devices may require changes to your applications or end nodes, or make packets unrecognizable to your network monitoring tools.
- **Complex management challenges:** Considerable configuration and parameter tweaking may be needed, especially those that require explicit configuration of tunnels between sites.
- **Significant infrastructure changes:** New devices may require changes to routing, switching, and security infrastructures.
- **Inefficient overlays or tunnels:** Any solution that relies on overlay tunnels (or other forms of encapsulation) can create network transparency issues.

A dynamic management problem

In most cases, a specific WAN optimization technique is best suited to deal with a one type of issue. But in the real world, WAN traffic is dynamic, with changing needs based on the changing characteristics of the traffic.

At one point in time, perhaps a user needs to access a web-based application. At the same time, another user might begin the transfer of some exceptionally large engineering design files. While it's theoretically possible to optimize your WAN to that specific scenario, doing so on a real-time basis is nearly impossible. There are just too many applications to consider – and no two are alike in the techniques that can best be used to speed them up.

The dynamic management of WAN optimization is a lot like piloting an aircraft. To do so manually, you need to make constant adjustments to factor in wind speed, wind direction, and the power needed to maintain your desired altitude and speed. Because that can be tiresome, an autopilot can be used to reduce the burden on the pilot.

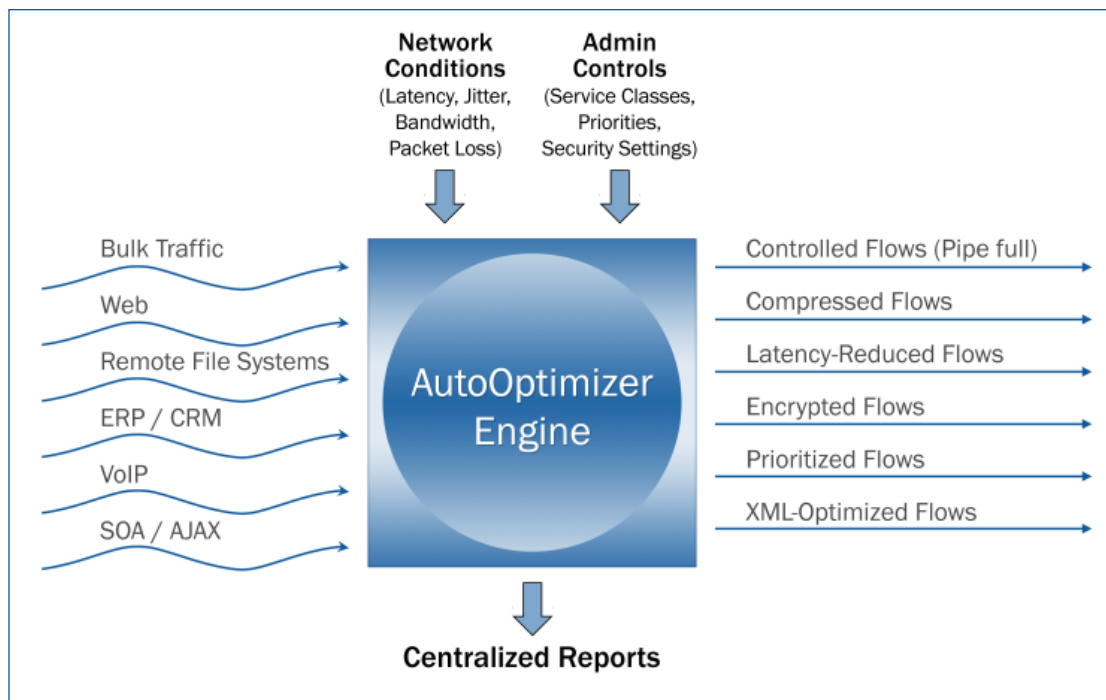
And that's the promise of a new approach – automated WAN optimization. Deploying it is like putting your WAN on autopilot.

No more tuning: An automated approach to WAN optimization

The challenge of optimizing a complex WAN environment can cause operational costs to soar – far outstripping initial capital costs. Making manual adjustments to compensate for changing network conditions and applications is an option that simply doesn't scale. And the problem is compounded further when more than one WAN optimization technique is deployed.

In fact, a study by Ashton, Metzler & Associates showed that 71 percent of IT executives experienced at least one unanticipated side effect with at least a "moderate" impact from installing a WAN optimization product. One out of eight said the impact was "very significant."

But a new automated approach can prevent these problems. The cornerstone of this technology is an engine that automatically and dynamically deploys appropriate WAN optimization techniques based on current conditions and data flow. This diagram shows how it works:



The right optimization techniques at the right time

The basic framework for automated optimization is to detect anomalies and changes based on current traffic patterns, use an algorithm to determine the techniques for optimal performance, and apply them – all in real time.

Needs are measured and the right optimization tool is selected from a variety of options, including TCP flow control, long history and link-level compression, service class policies, and protocol-specific techniques such as CIFS acceleration. What's more, the automated engine can distinguish between bulk-transfer and interactive applications to better manage latency sensitivity.

This improves the end user experience for everyone. A remote user accessing a Microsoft Word file needs it to open quickly. But a large data replication task would best benefit by reducing the footprint of the data on the WAN link.

For the Microsoft Word file, the engine would apply CIFS read optimization and application buffering to help the application load the remote document in minimal time. For the data replication task, compression and adaptive flow control would assure that the replication data fills the link's bandwidth effectively.

With continuous WAN "health checks," IT managers can rest easy knowing that an automated engine is compensating for any anomalies.

Significant improvements with a quick ROI

Those who have tested this automated approach see up to a 200-fold improvement in throughput and application response. Instead of using just 5 percent of its capacity, your T3 link can actually approach its full 45 Mbps maximum.

With a strategy that requires no fine tuning, automated WAN optimization assures the best performance with the least IT effort. There are no changes required to the network or applications – even as the network increases in scale – offering a substantial reduction in cost and a rapid ROI.

The Orbital Data AutoOptimizer: Simple, scalable WAN optimization

An “industry first,” Orbital Data’s AutoOptimizer is the only true automated solution for boosting WAN throughput and improving application response times.

At the heart of the system is the AutoOptimizer Engine, which dynamically applies the best performance boosting techniques depending on the application, data, and network conditions.

AutoOptimizer speeds up a broad range of enterprise applications such as Oracle Financials, Siebel CRM, and PeopleSoft HR. And you can expect significant improvements in email transfer speed, Web portal operation, and performance of disaster recovery solutions such as NetApp, SnapMirror, and NSI Double Take.

In fact, AutoOptimizer improves performance for virtually any TCP-based application, increasing WAN throughput by as much as 200 times and speeding up application response by as much as 400 percent.

With the Orbital Data’s AutoOptimizer, you can take advantage of:

- **Automated, dynamic management**

An intelligent, self-managing system makes sure that the right acceleration techniques are applied to meet changing conditions.

- **Extended reach to your entire enterprise**

AutoOptimizer architecture allows flexible choices for acceleration endpoints to meet the needs of remote users.

- **Higher transactional performance**

Performance improves through automated scheduling of interactive traffic ahead of bulk traffic, reduced compression, and improved retransmit logic.

- **Automated discovery and configuration**

Orbital Data units automatically discover each other and require virtually no configuration. Because they are transparent to the network, they can be dropped into your infrastructure with dynamic VPN capability and no need to maintain unit-to-unit mappings.

- **CIFS acceleration**

Windows file sharing is optimized by pipelining multiple CIFS messages over the WAN link and by accelerating file reads/writes and browsing of directories.

- **Complete transparency**

Because they are transparent to the network, Orbital Data units require no changes to applications, routers, firewalls, or clients. And you can use your existing network management tools just as you always have.

- **Enterprise-class scaling**

A unique design allows for 300,000 connections, 500 Mbps of throughput, and support for fan out of up to 300 remote branch offices.

Take advantage of a free trial

Get a personalized, practical perspective on how Orbital Data's AutoOptimizer can dramatically increase the performance of applications across your network. Try it for 30 days free with no obligation.

To see if you qualify for a free trial, or to learn more, call 1-800-280-3406 or email freetrial@orbitaldata.com. You can also visit us online at www.orbitaldata.com.

About Orbital Data

Orbital Data is the only company that can deliver LAN-like performance to WAN applications transparently, without costly network or software upgrades.

Using OrbitalData AutoOptimizer technology, enterprises can automatically boost WAN throughput and improve performance of distributed applications. Branch offices experience the same fast and predictable application response times as if they were on the corporate LAN.

With no need to add expensive bandwidth, companies can see a payback on their investment in just a few months and take full advantage of all of the strategic benefits of WAN connectivity. Orbital Data customers include Atlas, DataPath, Sony, Conexant, Petroleum Geo-Services ASA, Time Warner, Disney, and Yahoo!

Orbital | Data™

1900 South Norfolk Street
San Mateo, CA 94403

800 280 2406 toll free

650 227 7900 main

650 227 7911 fax

www.orbitaldata.com