

# Best Practices for Selecting WAN Optimization Solutions: Benchmarking Performance ROI

A Shunra Software White Paper

## Executive Summary

As businesses become increasingly dependent on the effective delivery of IT services over far-flung enterprise networks, technologies that optimize the performance of applications over the WAN have become increasingly attractive. In fact, this type of technology may be indispensable for companies operating on a global scale.

There are, however, many types of WAN optimization and application performance acceleration solutions from many different vendors. And IT budgets are not unlimited. So any decision-maker contemplating the acquisition of such a solution must somehow select the one that will deliver the greatest value, based on the company's specific IT environment and specific business requirements.

This white paper explains how to best make this critical technology decision. It provides an overview of best practices for defining requirements and comparing alternative solutions. Most important, it offers a metric - performance ROI - that IT organizations can use to select the WAN optimization solution that delivers the greatest business value.

Shunra enables these best practices by providing evaluation teams with a controlled benchmarking environment they can use to accurately compare solutions from various vendors. This benchmarking environment allows performance gains at the end-user desktop to be measured for both current and projected network conditions.

WAN optimization can ensure that remote users around the world have ready access to critical applications and services. But it's important to acquire the right solution for the job. This white paper outlines the principles and practices that will ensure that you do so.



## The WAN Optimization/Application Acceleration Solutions Market

Advances in WAN optimization and application acceleration have led to a proliferation of solutions, presenting potential buyers with both greater choice and greater confusion. In fact, vendors marketing these solutions use a variety of names to describe them – including “WAN acceleration,” “application acceleration” and “wide-area data services.” For the purposes of this paper, all of these diverse technologies will be referred to by the single, common term “WAN optimization.”

The following is an overview of the types of application acceleration and WAN optimization technologies available today<sup>1</sup>:

- **Data Reduction** is a collection of techniques aimed at reducing the amount of traffic that actually traverses the WAN. Implemented differently by various vendors, these techniques typically include data compression, caching, and terminal services – as well as recently introduced network memory solutions where only incremental changes in data are fetched across the network, while duplicate requests for the same data are served locally. These techniques are typically more effective for applications with static data and/or where the same data is accessed multiple times.
- **Latency mitigation** applies various algorithms to overcome the adverse impact that network delay has on application throughput and performance. They can include algorithms based on the network layer (such as TCP and CIFS acceleration) and application-specific algorithms for SQL, HTTP, Microsoft Exchange, SAP and/or Oracle, etc. These algorithms typically deliver greater benefits in environments where application performance is constrained by network latency, rather than by bandwidth.
- **Bandwidth management and QoS** prioritize certain network traffic to overcome the performance problems that occur when WAN links are over-subscribed. This way, traffic such as Internet browsing that is deemed less important – or where degradations in performance are considered more tolerable – takes a “back seat” to critical services such as transaction processing. These techniques are typically most appealing where a small number of high-priority applications must maintain their performance even under extreme network saturation conditions, while the performance of other applications can be allowed to deteriorate.

Most commercial WAN optimization solutions combine more than one of these techniques. It is the variety in these combinations – and the associated variation in their impact on different types of traffic – that make it difficult for decision-makers to determine where to best spend their limited budgets.

---

<sup>1</sup> This overview of WAN optimization technologies is based in part on the following document: Buyer’s Guide: Application Delivery Solutions, by Jim Metzler (2006). It is a vendor-sponsored resource which can be found online at <http://www.webtorials.com/abstracts/SilverPeak3.htm>

## Performance ROI: A Business-Driven Decision Metric

Faced with such diverse choices, IT decision-makers must carefully evaluate alternative solutions. A technical benchmark is insufficient for such an evaluation, because WAN optimization solutions vary greatly in cost and impact different applications in different ways. By the same token, it is difficult to calculate true ROI, since that would require assigning a concrete value to each second of reduction in an application's response time - which even if possible is not very practical.

From a business perspective, a more reasonable and practical metric is the performance gain achieved in proportion to a solution's cost - or "performance ROI." This metric can be applied uniformly across all solutions to determine which one best meet an organization's specific needs.

Several key principles are important for benchmarking performance ROI:

### **Benchmark the end-user's experience with business processes**

The point of implementing WAN optimization isn't to get packets from one router to another faster. It's to improve response times for end-users as they perform real-world business tasks. Benchmarks should therefore be based on improvements experienced at the desktop by remote users everywhere across the organization - especially those experiencing the slowest performance with the most important business applications.

### **Benchmark the full range of core business applications**

Companies need the WAN optimization solution that best meets their specific needs. So they should test solutions using key applications and services that are in frequent use or have a high business impact. Since these applications may vary greatly in their characteristics and ability to accelerate, it's important to evaluate how well any WAN optimization solution enhances the performance of each key application and service.

### **Benchmark under a full range of potential network conditions**

Benchmarking under current production conditions alone is insufficient. Those conditions will certainly change as users, locations and applications are added. So WAN optimization solutions should be benchmarked under various types of projected utilization and routing conditions - including worst-case and/or disaster recovery scenarios.

For these reasons, it is essential to benchmark WAN optimization solutions in a controlled testing environment that ensures consistency between tests and provides full visibility into performance to the desktop. This environment should allow engineers to accurately simulate the current network and to create any required "what-if" network scenarios. It should also be able to use production applications at current and projected utilization levels as required to appropriately "stress-test" the competing solutions being evaluated.



With such a benchmarking environment in place, IT organizations can effectively apply testing best practices to determine which WAN optimization solution delivers the greatest business value based on the specific characteristics of its network and its applications.

## A Five-Step Best Practices Benchmarking Process

Given that WAN optimization solutions must be carefully benchmarked to understand their respective performance ROI - and that such comparisons should be conducted in a controlled testing environment - what's the best approach to such comparative solution benchmarking?

The following five-step process is based on the proven best practices of IT organizations that have already taken this benchmarking approach:

### Step 1: Select applications and services for benchmarking

A meaningful benchmark will evaluate the potential impact of WAN optimization based on the current and potential requirements of the business. The first step in any benchmarking process is therefore the selection of the applications and services to be used in the test environment.

Key candidates should be applications currently in use such as:

- Applications that support revenue-critical business processes (such as ERP, CRM systems and transaction processing)
- Applications that are extensively used across the organization (such as email, intranet portals and VoIP)
- Applications that have high visibility at the executive level (such as Business Intelligence, dashboards, financial reporting)
- Applications that have known performance issues across the WAN (such as databases and "fat client" applications)

Applications can also be selected based on their technical characteristics in order to better assess the various strengths and weaknesses of each vendor's WAN optimization solution. In some cases, such applications may not currently be running in the environment - but it is nonetheless prudent to assess how well each vendor's solution handles them. For example, it's a good idea to ensure that the ultimate solution-of-choice will be able to handle real-time services such as VoIP and streaming video.



## Step 2 - Set up the benchmarking test bed

The benchmarking test bed should be capable of both 1) accurately reflecting current conditions across the enterprise environment and 2) emulating any required “what-if” scenarios. The test bed will typically consist of four elements:

- A **WAN emulation appliance** that impairs the flow of packets in a manner that accurately reflects the impact of real-world network factors such as distance, latency, congestion and jitter on end-to-end application performance
- Multiple **desktop PCs** that are configured similarly to those actually being used at various remote locations and are set up to interact with test applications manually or using application-specific “scripts” that simulate the behavior of end-users
- **Application servers** running software and services selected in Step 1, which can either be the actual production servers or dedicated test servers specifically configured to support multiple traffic scenarios
- The **WAN optimization solution** being evaluated

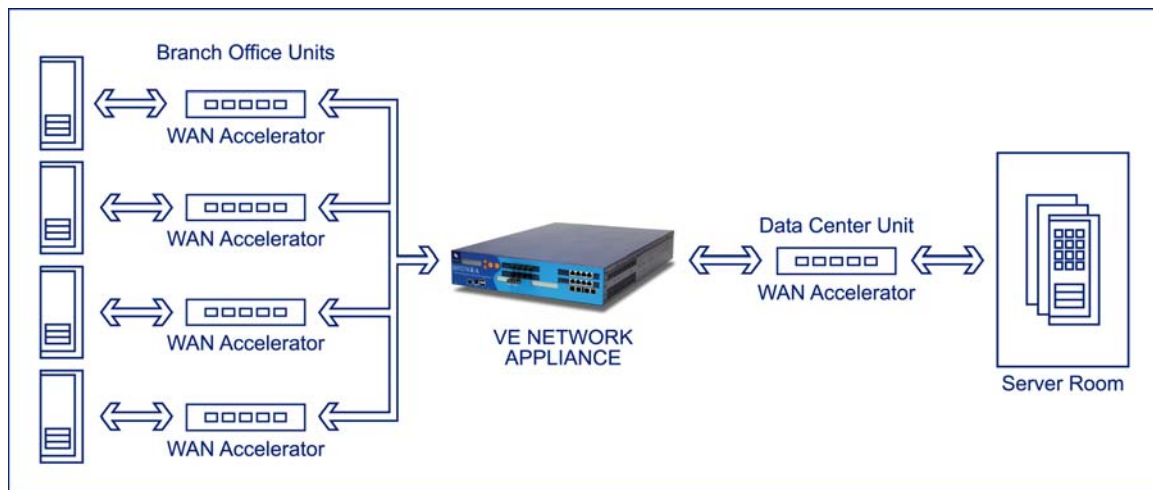


Figure 1: The WAN optimization benchmarking environment

The test bed must obviously also be able capture the results of the tests and present them in a form that is understandable to the decision-making team.

### Step 3 - Define and implement the network scenarios

Once the test bed is in place, the testing team can configure it to appropriately benchmark the WAN optimization solutions being evaluated. There are basically three sets of factors that must be considered when creating these configurations:

#### 1) Network topology

The topology of the network test bed will typically be a data center and multiple branch offices. Although the model can potentially include dozens of remote sites, an effective benchmark can use a configuration with a relatively small number of sites that represent a cross-section of an organization's actual remote offices. Figure 2 shows a typical model with ten remote sites located around the world and a data center located in North America.

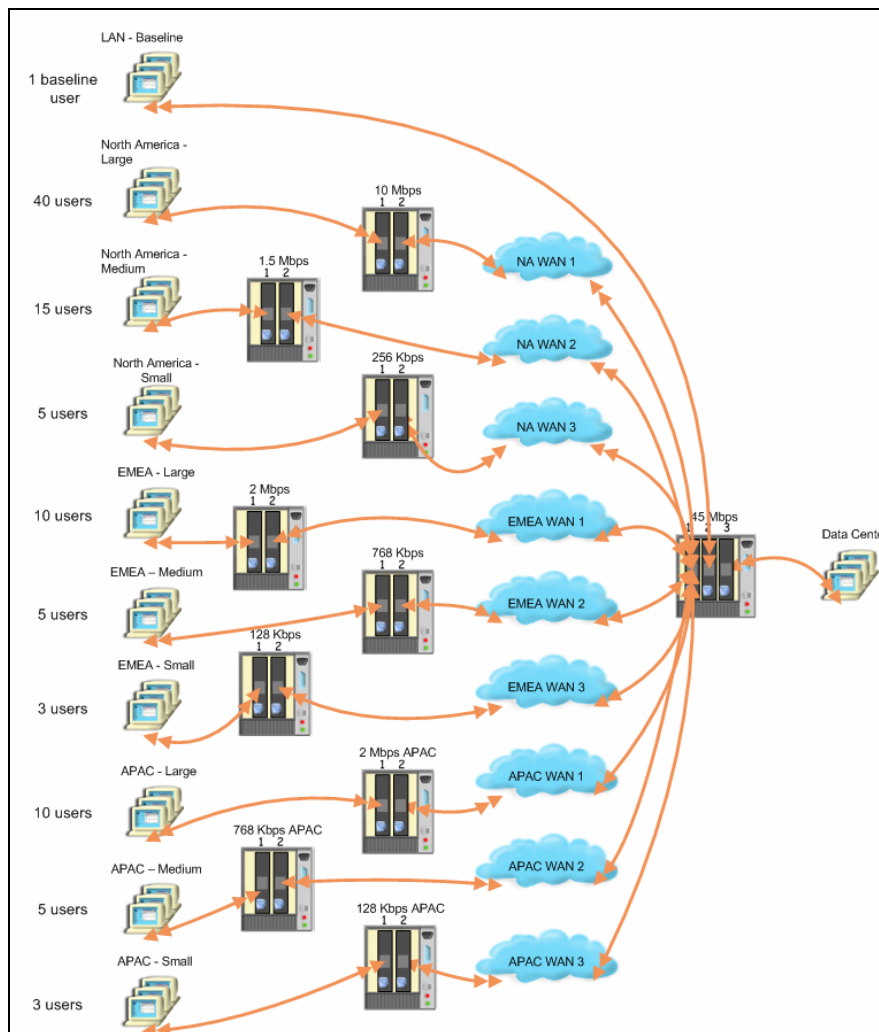


Figure 2: A typical model with ten remote sites located around the world and a data center located in North America, as drawn in Microsoft Visio®

## 2) Input of current network parameters

To assess the impact of the WAN optimization solution under current network conditions, the test bed must be configured to accurately reflect those conditions. This can be done using technology that actually “records” those conditions in real time from the production network for “playback” on the WAN emulation appliance, or by manually defining parameters for each remote office based on available information.

Here are some typical parameters for the ten offices in our sample configuration:

Region	Office Size	Bandwidth	One-way WAN latency	Number of users
North America	Large	10 Mbps	25 msec	40
	Medium	1.5 Mbps		15
	Small	256 Kbps		5
EMEA	Large	2 Mbps	50 msec	10
	Medium	768 Kbps		5
	Small	128 Kbps		3
APAC	Large	2 Mbps	70 msec	10
	Medium	768 Kbps		5
	Small	128 Kbps		3

Figure 3: Sample configuration metrics

## 3) Variety of “what-if” scenarios

As noted earlier, benchmark testing should be performed under a variety of scenarios to ensure that the WAN optimization solutions can fulfill the future requirements of the business. In particular, testers will want to:

- Include a scenario that combines **high latency and packet loss impairments**, since latency mitigation algorithms may become less effective when packet loss is introduced
- Include a **disaster recovery scenario** that introduces intermittent loss of network connectivity, since some WAN optimization solutions rely on synchronization between devices in remote offices and those in the data center
- Include various levels of network saturation scenarios, including a **fully utilized network scenario** – since QoS solutions prove their worth only under such conditions. It’s also important to test QoS solutions with a mix of real-time applications and data applications to make sure that the prioritization one doesn’t adversely impact the other.



### Step 4 - Perform the benchmark tests

With all the preparation complete, benchmarking tests can be run under the network scenarios defined in Step 3. Business processes should be run multiple times in each test scenario to replicate the real-world effect of users repeatedly accessing the same data.

Below is an example of what the test result might look like with two different WAN optimization solutions:

Site	LAN Baseline (Sec)	Current WAN Baseline (Sec)	Average Response Time Solution A (Sec)	Average Response Time Solution B (Sec)
North America (large office)	15	23	19	21
North America (medium office)	15	25	22	24
North America (small office)	15	29	30	32
EMEA (large office)	15	35	30	32
EMEA (medium office)	15	38	33	35
EMEA (small office)	15	42	35	38
APAC (large office)	15	48	40	47
APAC (medium office)	15	54	42	49
APAC (small office)	15	58	45	51
Total		352	296	329
Average response time per site		$352/10 = 35.2$	$296/10 = 29.6$	$329/10 = 32.9$
Average gain per site vs. WAN baseline			$35.2 - 29.6 = 5.6$	$35.2 - 32.9 = 2.3$

Figure 4: Comparison of response times between two WAN optimization solutions for an inventory update process

Testers can also experience the impact of each solution first-hand by performing manual tests from the desktops that are included in the test bed while the automated test is running.



### Step 5 - Calculate and compare performance ROI

Based on the benchmarks metrics captured across all testing scenarios, the testing team can produce the performance ROI figures that decision-makers need to select the optimal WAN optimization solution for the business.

Consider the following results for four different solutions tested for three different business processes:

Business Process	Average Performance Gain with Vendor A (Sec)	Average Performance Gain with Vendor B (Sec)	Average Performance Gain with Vendor C (Sec)	Average Performance Gain with Vendor D (Sec)
Update inventory record	5.34	2.45	3.18	4.33
Access file on intranet	4.22	7.7	4.5	5.2
Find a contact in the CRM system	4.5	3.3	5.2	6.2
<b>Total</b>	<b>14.06</b>	<b>13.45</b>	<b>12.88</b>	<b>15.73</b>

Figure 5: Benchmark comparison for four WAN optimization solutions

To calculate performance ROI, we also need to quantify the “I” - which is the total cost of ownership for each solution given the number of remote sites where the solution must be installed. This number should include hardware costs, setup costs, management costs, and any additional professional services required.

For 20 remote sites, these costs might be as follows:

TCO - Vendor A	TCO - Vendor B	TCO - Vendor C	TCO - Vendor D
\$120,000	\$110,000	\$90,000	\$120,000

Figure 6: Sample TCO for four WAN optimization solutions

The total performance gain for each solution can be calculated by multiplying the average gain per business process by the total number of sites where the WAN optimization devices will be deployed (20, in this example).



By dividing the performance gain by TCO, we can calculate the performance ROI figure for each solution:

	Vendor A	Vendor B	Vendor C	Vendor D
TCO	\$120,000	\$110,000	\$90,000	\$120,000
Performance Gain	281.2	269	257.6	314.6
Performance ROI (sec/\$1000)	$281.1/120 = 2.34$	$269/110 = 2.45$	$257.6/90 = 2.86$	$314.6/120 = 2.62$

Figure 7: A sample "performance ROI" report

Of course, performance ROI may not be the sole factor in the acquisition of a WAN optimization solution. The financial stability of the vendor, the quality of support and service, and other contract terms may also come into play. Decision-makers may also choose to use the result of such benchmarking to wrest price and contract concessions from a favored vendor. But performance ROI is clearly a powerful metric for sorting through the competing claims of WAN optimization vendors - and for making the purchasing decision that best suits the requirements of the business.

## Conclusion

In today's networked enterprise, employee productivity can be highly contingent on the performance of critical applications over the corporate WAN. For this reason, WAN optimization technology can deliver significant business value. However, to wisely select the WAN optimization solution that will deliver the greatest return-on-investment to the business, IT decision-makers need a way to comparatively evaluate competing vendors' offerings. The five-step methodology described in this paper and the performance ROI metric that it produces provides an accurate basis for such a value comparison. By applying this methodology and metric to the acquisition of WAN optimization technology, IT organizations can optimally align their expenditures with the specific needs of the business.



## About Shunra

Shunra is the pioneer and market leader in predicting how business applications and network services will perform for remote end-users - before deployment. The Shunra VE solutions simulate any production network environment in the pre-deployment lab. It delivers a powerful, flexible and easy way to test the performance of applications or network equipment under a wide variety of network impairments - as if they were running in a real production environment. Through this process users understand the impact that the network and applications have on each other's performance and on remote end-users' experiences, and uncover and resolve production-related problems before rollout.

The benefits of using Shunra VE to evaluate WAN optimization solutions include:

**Simplification of the evaluation process:** Shunra VE makes it easy to accurately and repeatedly recreate production network conditions in the lab.

**Optimized integrity of the evaluation process:** By ensuring the accuracy and consistency of testing conditions, Shunra VE protects the integrity of vendor selection

**Deep analysis:** Shunra helps pinpoint even the most subtle bottlenecks in any business processes, so decision-makers can appropriately address both latency-related and bandwidth-related constraints.

**Faster time-to-benefit deployment time:** By fully validating the use of the WAN acceleration solution-of-choice, Shunra eliminates the need for extensive field trials or gradual rollouts - allowing the business to more quickly realize the full benefit of its technology investment.

Over 1,700 leading companies worldwide use Shunra's award-winning solutions - including 3M, Boeing, Cisco, Dow Chemical, EMC, FedEx, General Motors, JPMorgan Chase, Kelly Services, Merrill Lynch, Motorola, Nestle, Pitney Bowes, and Vodafone. Shunra's headquarters are located in New York City, and the company has a major R&D center in Israel, and offices in the UK and Sweden. Shunra is also supported through a global network of channel partners. For more information, please visit [www.shunra.com](http://www.shunra.com).

## Shunra Professional Services

Shunra's professional services organization provides expert engineering support for focused engagements ranging from the planning of network changes to the independent evaluation of third-party technology solutions in the context of business-specific requirements. Combining proven best practices with Shunra's industry-leading network performance emulation technology, Shunra Professional Services provides high-value turnkey services that enable clients to make fully informed decisions based accurate and relevant benchmarks of application and service performance to remote end-users. For more information on Shunra Professional Services, email [experts@shunra.com](mailto:experts@shunra.com)

