



F5 Technical Brief

BIG-IP WAN Optimization Module Performance

Performance guidelines and testing results for the
BIG-IP WAN Optimization module.

by Kollivakkam Raghavan

Senior Manager, Product Development



Contents

Executive Summary	3
--------------------------	----------

Characteristics of Traffic over the WAN	4
------------------------------------------------	----------

WAN Optimization Solution Considerations	6
BIG-IP WAN Optimization Module features	7
BIG-IP WOM Sizing Considerations	8

Summary Results	9
BIG-IP configuration for data center to data center replication use cases	10

Test Results	12
HTTP	13
FTP	14
CIFS	15
Oracle Streams Replication	16
NetApp FlexCache	17
NetApp SnapMirror	18
vMotion	20

Appendix: Hardware and Software Specifications	22
-------------------------------------------------------	-----------



Executive Summary

An increasing amount of data is being transferred between data centers, driven by data replication and backup, storage requirements, and data center consolidation. The trend toward virtualization and cloud computing, while reducing IT costs and providing agility, also adds latency to application delivery. WAN optimization reduces the replication time and overcomes latency issues between data centers by optimizing and accelerating data over the WAN.

F5 BIG IP® WAN Optimization Module™ (WOM) overcomes network and application issues on the WAN to ensure that application performance, data replication, and disaster recovery requirements are met. These services are integrated directly on your F5 BIG-IP device and include superior adaptive compression, intelligent data deduplication, protocol optimization, and traffic control capabilities that dramatically reduce data replication times and enable more efficient use of your existing bandwidth. Encryption of data over the WAN ensures you meet regulatory and security requirements when transferring business-critical information. In addition, by offloading CPU-intensive encryption and compression processes from servers, F5 can save you from adding costly hardware upgrades.

This document has several objectives:

- Characterize the performance of the BIG-IP WAN Optimization Module, to clearly show the value of using a BIG-IP system for data center to data center data replication.
- Provide guidelines to administrators on the configuration settings to optimize WAN performance. This includes the BIG-IP configuration as well as characterization of specific protocols and the optimizations that give the maximum benefit.
- Provide test results with context that can be independently verified.

The greatest strength of the BIG-IP system is the layered software architecture which allows F5 to provide multiple modules on a single device. This flexibility requires proper configuration of the modules for optimal performance. For those new to symmetric WAN optimization, the configuration can be a little confusing. The BIG-IP product team has developed several built-in features to simplify configuration of these WAN optimization services on your BIG-IP system.



Characteristics of Traffic over the WAN

The problem of application performance on the WAN can be traced to the following sources:

1. Latency

Latency is the time taken by a packet to be returned to the sender and approximately equals the Round Trip Time (RTT). In reality, the actual latency is affected by many factors such as the network path taken to the destination (which includes the physical media that is used – fiber versus satellite versus copper), and the number of hops or routers between the source and the destination, with each hop adding some latency to the packet. Latency is typically expressed in milliseconds.

2. Bandwidth

Bandwidth corresponds to the amount of data that can be carried in a physical medium in a given time interval (usually a second). Bandwidth is typically expressed in Mega or Kilo bits per second (Mbps or Kbps).

3. Bandwidth delay product

Bandwidth delay product, also known as BDP, is a product of latency and bandwidth and is typically used to express the amount of data that can be held in a given physical medium. The BDP, whose units are in (M|K) represents the capacity of the physical medium. The goal of all WAN optimization products is to keep the WAN completely full at all times.

4. Chatty protocols

Chatty protocols are inefficient protocols, which clients use to communicate with servers. The protocols cause many applications that run well on the LAN to fail to deliver adequate performance on the WAN. On a LAN where the RTT is negligible, the chatty nature of these protocols is not significant enough to cause performance problems. However, on a WAN with even a moderate RTT, the performance impact becomes significant enough to not only be noticeable, but in many cases renders the application unusable. The chattiness of a protocol is directly proportional to the RTT of the link.

5. Packet Loss

Packet loss occurs when one or more packets fail to reach their destination, typically due to quality of the physical medium. Higher quality networks, which are significantly more expensive, have lower packet loss and perform better. However, on public networks (like the Internet), loss of 1% is typical, and results in retransmissions in reliable protocols like TCP. Retransmissions make the latency problem even worse for TCP based applications.

Armed with these factors, we can now identify several ways to optimize traffic over the WAN. The best way to optimize WAN traffic is to have a symmetric pair of devices on both sides of the WAN, which allows the devices to perform many optimizations that would not otherwise be possible. The WAN optimization approaches to overcome the problems of application performance over the WAN include the following:



1. Improve bandwidth utilization

Techniques like data deduplication exploit the redundancy inherent in the data traversing the WAN by maintaining dictionaries on each device and sending references to the redundant data, thereby significantly reducing the amount of data over the WAN. This reduces the congestion and increases the effective capacity of the WAN. Compression is also used to compress data over the network to yield further bandwidth savings. Deduplication in conjunction with compression is the technique used to maximize bandwidth over the WAN.

2. Reduce perceived latency

This is a protocol specific optimization that is used to provide LAN like response times as perceived by the client program making the request. The technique typically involves a deep understanding of the underlying protocol and using read-ahead and write-behind approaches to stage the data on the device that's closest to the client. Because the WAN optimization devices in this case is a man-in-the-middle, transparently intercepting the client requests, it is not suitable for all protocols.

3. Improve lossy network performance

Using this technique, the WAN optimization device takes advantage of the symmetric deployment to provide TCP optimizations that would not normally be available when communicating with a general purpose server's TCP stack. Since both ends of the WAN TCP connection are "owned" by the WAN optimization device, the TCP stack behavior can be made more aggressive when addressing packet loss conditions in the network.



WAN Optimization Solution Considerations

With the different techniques available for WAN optimization, it is important to understand that simply turning on all available optimizations in every situation does not provide the best performance gains. Several considerations must be taken into account when determining the best set of optimizations to use for any given situation. The available optimization techniques are like a set of tools available to solve a problem; selecting the right set of tools to use for a given situation is critical for maximum performance. The following section describes the various factors involved in making this decision and provides a table for each available optimization type.

Consideration	TCP Optimization	Deduplication	Compression	Protocol Optimization
Type of data	Not Applicable	Impacted	Impacted	Not Applicable
Type of protocol	Not Applicable	Impacted	Impacted	Not Applicable
Type of network	Impacted	Impacted	Impacted	Impacted
Type of application	Impacted	Impacted	Impacted	Impacted

For each of these considerations, we now describe how the impacted optimizations are affected:

Type of data

Optimization	Description
Deduplication	This technique is very dependent on the nature of the data. If the data traversing the WAN does not have much redundancy (either because of the inherent nature of the data or because redundancy has been externally removed using other mechanisms), this technique will provide limited gains.
Compression	Like deduplication, compression is also affected by the type of data that's being transferred over the WAN. For instance, if the data is comprised of compressed JPEG images, compression techniques on the WAN optimization device will not provide much benefit. On the other hand, if the data consists of XML documents that are being transferred, and assuming that they're not already compressed at the source, compression by the WAN Optimization devices typically provides significant bandwidth savings.

Type of protocol

Optimization	Description
Deduplication	Deduplication works best when the amount of data available for WAN Optimization devices to deduplicate is large. Since some protocols (typically multiple request-response type protocols) send data in smaller chunks and wait for an acknowledgement before sending the next request, deduplication may not provide benefits for these protocols. Database replication applications typically fall in this category. Another factor affecting deduplication is the introduction of latency. Some protocols are extremely sensitive to latency; performance can degrade dramatically when latency increases slightly. Any benefit deduplication might provide is typically lessened and should be more closely tested for these protocols.
Compression	The nature of the protocol affects compression in a similar way to deduplication although to a lesser extent. Latency sensitivity considerations also impact whether compression is appropriate.



Type of network

Optimization	Description
TCP Optimization	TCP optimization can significantly improve the performance of the stack because of the symmetric deployment nature of the WAN optimization devices. The presence of two identical devices on either end of the WAN network allows TCP to be less conservative than general purpose TCP stacks.
Deduplication	The nature of the network affects deduplication indirectly. Before determining whether deduplication is appropriate on the WAN Optimization device, it is important to understand how the WAN is being used. If congestion on the WAN is not an issue, and the WAN is underused to begin with, deduplication will probably not provide any material benefit. This is typically seen in database replication protocols where the maximum throughput of the application is much less than WAN capacity.
Compression	Same as deduplication
Protocol Optimization	When the protocol is such that the WAN is not fully used, protocol optimization (also known as Layer 7 or L7 optimizations) can be used to improve WAN use. For example, the CIFS protocol typically makes requests in 32K chunks. By using read-ahead techniques and by increasing the size of the requests to 64K or 128KB, better WAN use is possible, which leads to improved performance.

Type of application

Optimization	Description
TCP Optimization	This optimization is used only if the underlying communication protocol uses TCP.
Deduplication	Many applications use the request-response protocol to communicate remotely. Chatty applications, like Microsoft Word accessing a document over a WAN, or SQL replication, make multiple requests, with each request needing a response before proceeding to the next request. Depending on the nature of the application, an intermediary may or may not be able to improve this performance. For instance, since database and storage replication applications provide data integrity guarantees, interfering in their operations will almost certainly not be acceptable in any scenario. On the other hand, open a Word document over the WAN can be improved using L7 optimization techniques.
Compression	Same as deduplication
Protocol Optimization	See note for deduplication. Protocols such as CIFS can be optimized using methods such as read-ahead and write-behind to improve WAN use and reduce perceived latency. Other application protocols, like Storage or Database replication, do not lend themselves to these application-specific approaches and cannot be optimized to the same degree.

BIG-IP WAN Optimization Module features

There are several optimization techniques available to solve a particular problem. While some of the features are available for free with the base BIG-IP platform, others are licensed as part of WOM. The free optimization features on the BIG-IP are enabled by provisioning the WOM Lite services, which makes the WAN optimization menu available for configuration. It is important to remember that all the optimizations require a pair of BIG-IP appliances, each appropriately configured, to optimize the traffic over the WAN. The optimizations available are described in the table on the following page.



Module/Feature	TCP Optimization	Deduplication	Compression	Protocol Optimization
WAN Optimization Lite	YES	NO	YES	NO
WAN Optimization Module	YES	YES	YES	YES (CIFS/MAPI)

* For BIG-IP WOM TCP optimizations, a symmetric deployment with two BIG-IP devices is necessary. Even though BIG-IP WOM TCP profiles are available in the base BIG-IP platform, the optimizations will not take effect unless there are BIG-IP devices on both ends of the WAN. In fact, using the WOM TCP profiles without BIG-IP devices on both ends of the WAN could result in detrimental TCP behavior. Also note that non-WOM based TCP optimizations still work in an asymmetric fashion on the BIG-IP system and can be used on single box deployments.

BIG-IP WOM sizing considerations

Sizing a BIG-IP WOM deployment depends on all of the factors mentioned above and it is critical to carefully consider all the factors in order to have a successful installation. The following tables provide guidelines to determine which platform and deduplication method are best suited for particular bandwidth conditions.

For bandwidths up to 45 Mbps - Disk-based deduplication

Desired WAN Throughput	Compression*	WAN Encryption	BIG-IP Platform
Up to 10 Mbps	Symmetric Adaptive Compression	Enabled	1600 and above
Up to 45 Mbps	Symmetric Adaptive Compression	Enabled	6900 and above

Disk-based deduplication is not recommended for bandwidths above 45 Mbps.

For bandwidths above 45 Mbps - Memory-based deduplication

Desired WAN Throughput	Compression*	WAN Encryption	BIG-IP Platform
Up to 155 Mbps	Symmetric Adaptive Compression	Enabled	3600, 3900
Up to 622 Mbps	Symmetric Adaptive Compression	Enabled	6900
Above 622 Mbps	Symmetric Adaptive Compression	Enabled	8900

* **Important:** On many occasions, Proof of Concept (POC) deployments are conducted without the use of a WAN emulator. Because Symmetric Adaptive Compression reacts to data type as well as network conditions, using a 1Gbps link with no latency results in Symmetric Adaptive Compression selecting the NULL compression algorithm (no compression) because the compression adapter correctly detects the network is not a bottleneck in the system and chooses not to add additional latency. It is critical that the compression settings be changed to use a specific algorithm and turn off adaptive compression in this case. The success of the proof of concept can depend on this choice.



Summary Results

The results presented in this document are the results of extensive testing and tuning of the different configuration settings available to the user. Therefore, one of the goals of this application testing effort has been to provide simple initial configuration for BIG-IP WOM deployments that solve the most common use cases and guidelines for tuning the box to extract maximum performance depending on the use case. Our tests have shown that because of differences in the hardware architectures of BIG-IP platforms the optimum settings are platform dependent.

It is important to understand that the optimum settings are useful for the average case and serve as a good starting point. Depending on the use case, these settings can be further tuned to get optimum performance for that situation. The graphs below show the optimization results for each protocol with and without the presence of the BIG-IP. In each case the effect of using BIG-IP WOM Lite and BIG-IP WOM optimizations is also shown to characterize the impact of each available optimization. The following matrix shows the recommended optimization settings for each protocol.

Note: The Quick Start page on the BIG-IP Configuration utility implements both the correct platform specific TCP settings as well as the recommended optimization with a single click for the applications identified below.

Application/Protocol	TCP Optimization	Compression	Deduplication	L7 Optimization
HTTP	YES	YES	YES	NO
CIFS	YES	YES ¹	YES ¹	YES
Oracle Streams	YES	YES ²	NO	NO
NetApp Flex Cache	YES	YES	YES	NO
NetApp SnapMirror	YES	YES	YES	NO
VMware vMotion	YES	YES	YES	NO
FTP	YES	YES	YES	NO

¹ Deduplication and Compression performance significantly improved in conjunction with CIFS optimization

² Compression for SQL replication does not provide any significant benefit. If the additional latency introduced by compression is a concern then compression can be disabled.



BIG-IP configuration for data center to data center replication use cases

The data center to data center replication use cases are characterized by small numbers of connections. Single connection throughput is therefore an important metric, because the overall aggregate throughput is largely constrained by this metric. By far the most difficult part of the implementation is the TCP profile configuration. This section provides the recommended TCP profile settings for a typical BIG-IP WOM deployment. Configuration of iSession compression and deduplication are much simpler and guidelines for these configurations are provided as well.

BIG-IP TCP configuration

The recommended configuration is one that gives the optimum single connection throughput for a given platform. The following table describes the per-platform setting for the TCP profile used on the WAN facing side of the virtual (the server side profile context on the client BIG-IP WOM and client side profile context on the server BIG-IP WOM).

The BIG-IP web-based Configuration utility automatically implements these guidelines when you use the Quick Start to configure the optimization settings.

Per-Platform TCP profile settings

BIG-IP Platform	Send Buffer	RECV Window
1600	800 KB	800 KB
3600	2,000 KB	2,000 KB
3900	600 KB	600 KB
6900	800 KB	800 KB
8900	800 KB	800 KB

Deduplication configuration considerations

1. For data center to data center replication use cases at high speed links (OC12, 1Gbps, OC192) deduplication in memory mode is recommended.
2. For lower speeds disk mode may be enabled.
3. For latency sensitive applications or where data is not compressible (such as JPEG image replication) deduplication may not provide substantial benefit.
4. Disk based deduplication is only recommended for lower speed links. For instance, a dual drive system can support WAN bandwidths of up to 45 Mbps. Disk mode in single drive systems is suitable for 10 Mbps and lower networks.



Compression configuration considerations

1. For most applications, the default iSession compression settings are sufficient (see Figure 1).
2. For latency sensitive applications or when the data is not compressible, compression should be turned off (only null enabled).
3. If the data is reasonably compressible, LZO can be used when the fastest compression algorithm is desired. Where maximum compression is desired, deflate should be used.

Local Traffic » Profiles : Services : iSession » New iSession Profile...

General Properties

Name	<input type="text"/>
Parent Profile	iSession

Compression Settings Custom

Deduplication	Enabled	<input type="checkbox"/>
Adaptive	Enabled	<input type="checkbox"/>
Deflate	Enabled	<input type="checkbox"/>
Deflate Level	1	<input type="checkbox"/>
LZO	Enabled	<input type="checkbox"/>
Null	Enabled	<input type="checkbox"/>

Outbound iSession to WAN Custom

Reuse Connection	Enabled	<input type="checkbox"/>
Port Transparency	Enabled	<input type="checkbox"/>
Application Data Encryption	Disabled	<input type="checkbox"/>

Inbound iSession from WAN Custom

Target Virtual	match all	<input type="checkbox"/>
----------------	-----------	--------------------------

Cancel Repeat Finished

Figure 1: Default iSession profile on the BIG-IP



Test Results

This section contains the optimization results for each protocol with and without the presence of the BIG-IP WAN Optimization Module. Each set of results contains information on the specific test scenario and network conditions used in the testing.

The F5 optimizations are activated through simple check-boxes in the BIG-IP WOM Quick Start menu and can be configured as specifically as necessary; from all traffic on every subnet, down to one single port on one specific host.

Use the following list to go to a particular set of results:

- **HTTP** on page 13
- **FTP** on page 14
- **CIFS** on page 15
- **Oracle Streams Replication** on page 16
- **NetApp FlexCache** on page 17
- **NetApp SnapMirror** on page 18
- **vMotion** on page 20

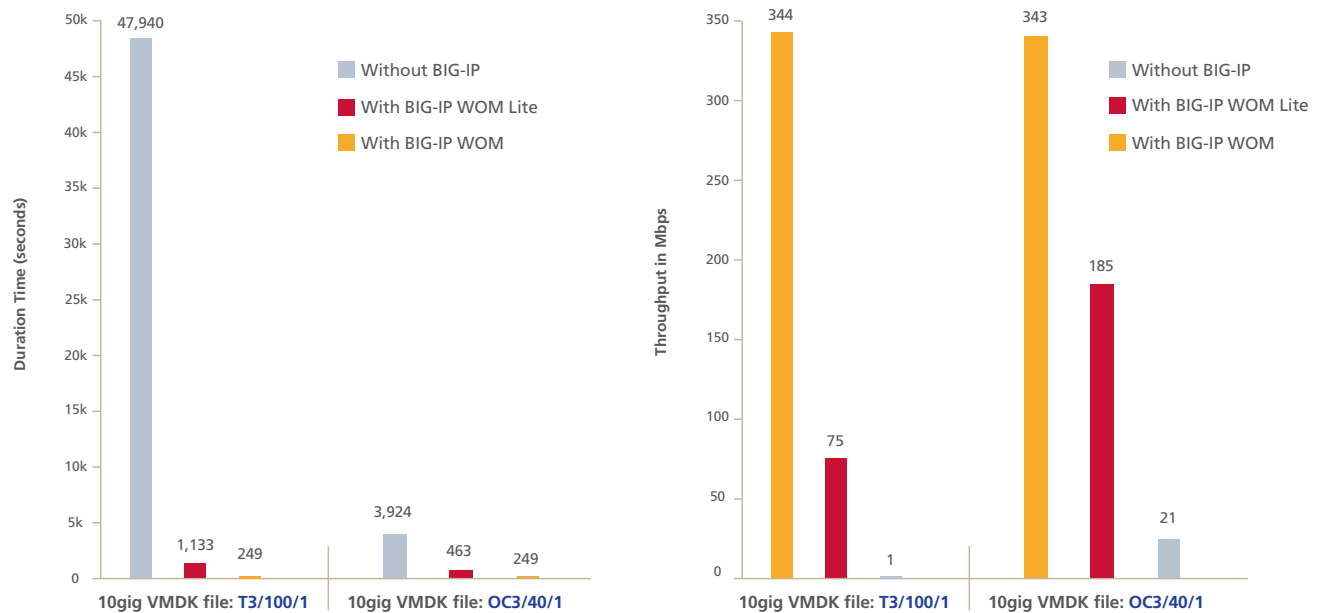
Specific configuration procedures and information can be found in the BIG-IP WOM documentation, available on Ask F5 (<http://support.f5.com/kb/en-us.html>).



HTTP

By enabling HTTP optimization on the BIG-IP WOM, HTTP traffic between local and remote endpoints is accelerated. In the following tests, a large file was transferred through the BIG-IP WOM tunnel in various network conditions. The duration of the transfer time was measured along with the effective throughput. The file being transferred was a 10 gigabyte VMware disk file (VMDK). The network conditions being tested were T3 (45 Mbits per second) with 100 milliseconds of round trip latency and 1% packet loss and OC3 (155 Mbits per second) with 40 milliseconds of latency and 1% packet loss.

For the duration time results (graph on the left), shorter transfer times are better. You can see the BIG-IP WOM accelerated HTTP transfers under these conditions in a range from 4 to 190 times faster than with no acceleration. For effective throughput results (graph on the right), longer bars are better. You can see the BIG-IP WOM increased effective throughput in a range from 75 Megabits per second to 344 Megabits per second for T3 links, and a range of 185 Megabits per second to 343 Megabits per second for OC3 links.



HTTP summary

Performance improvements achieved using the BIG-IP WOM

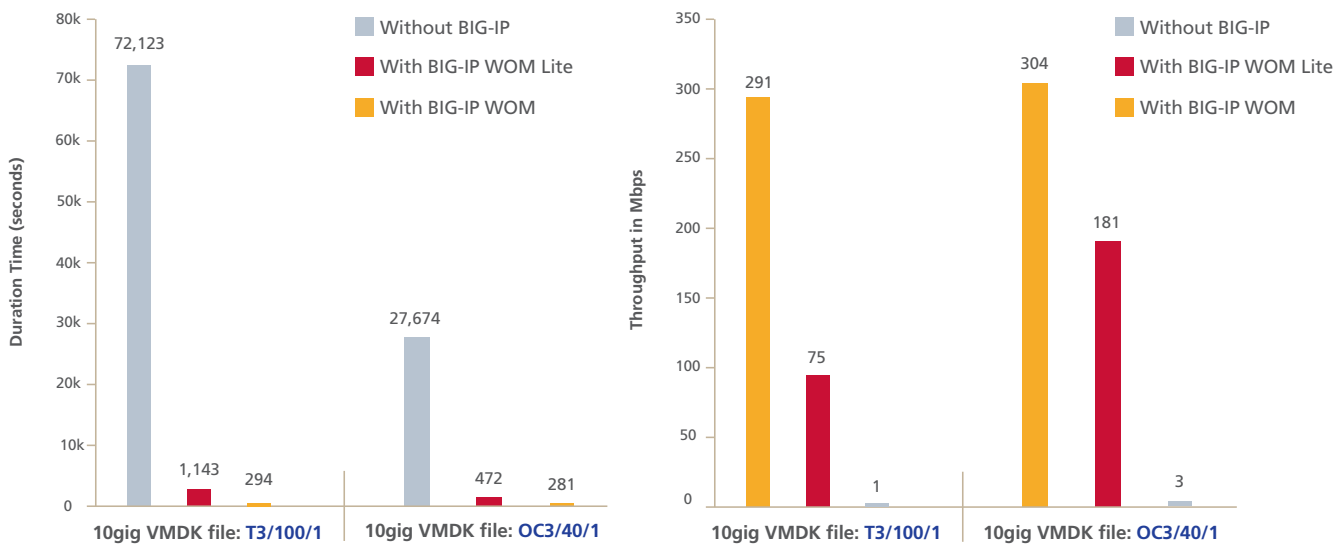
Optimization	T3/100/1	OC3/40/1
TCP	16x	4x
Compression	42x	8x
Deduplication and Compression	190x	16x



FTP

By enabling FTP optimization on the BIG-IP WOM, FTP traffic between local and remote endpoints is accelerated. In the following tests, a large file was transferred through the BIG-IP WOM tunnel in various network conditions. The duration of the transfer time was measured along with the effective throughput. The file being transferred was a 10 gigabyte VMDK file. The network conditions being tested were T3 (45 Mbits per second) with 100 milliseconds of round trip latency and 1% packet loss and OC3 (155 Mbits per second) with 40 milliseconds of latency and 1% packet loss.

For the duration time results below (graph on the left), shorter transfer times are better. You can see the BIG-IP WOM accelerated FTP transfers under these conditions in a range from 24 to 240 times faster than with no acceleration. For effective throughput results (graph on the right), longer bars are better. You can see the BIG-IP WOM increased effective throughput in a range from 75 to 291 Megabits per second for T3 links and a range of 181 to 304 Megabits per second for OC3 links.



FTP summary

Performance improvements achieved using the BIG-IP WOM

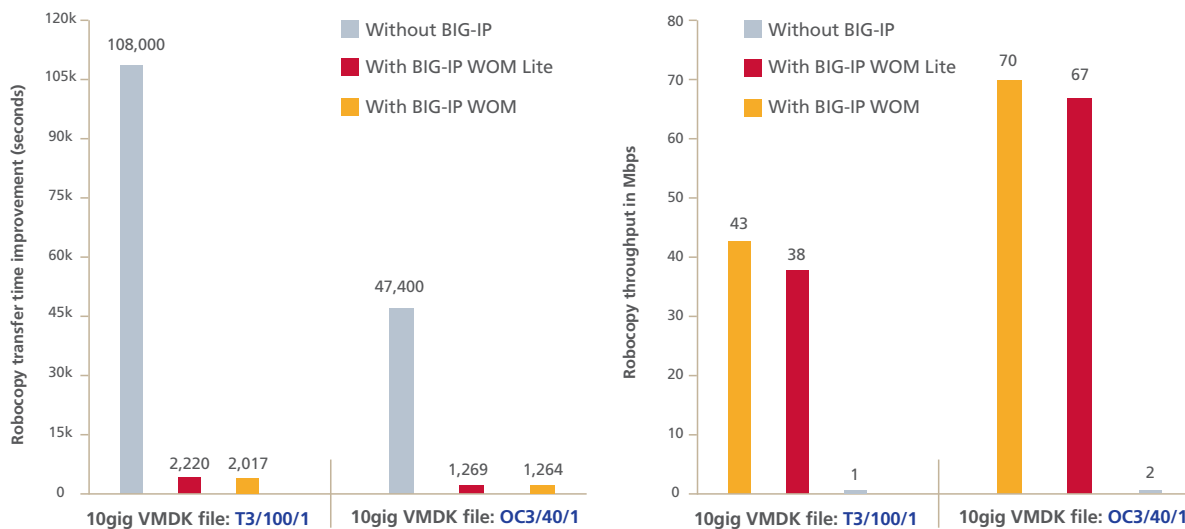
Optimization	T3/100/1	OC3/40/1
TCP	25x	24x
Compression	62x	60x
Deduplication and Compression	240x	100x



CIFS

By enabling CIFS optimization on the BIG-IP WOM, Microsoft Office and Windows File Sharing (CIFS) traffic between local and remote endpoints is accelerated. In the following tests, a large file was transferred through the BIG-IP WOM tunnel in various network conditions. The duration of the transfer time was measured along with the effective throughput. The file being transferred was a 10 gigabyte VMDK file. The network conditions being tested were T3 (45 Mbits per second) with 100 milliseconds of round trip latency and 1% packet loss and OC3 (155 Mbits per second) with 40 milliseconds of latency and 1% packet loss.

For the duration time results (graph on the left), shorter transfer times are better. You can see the BIG-IP WOM accelerated CIFS transfers under these conditions in a range from 1.18 to 55.19 times faster than with no acceleration. For effective throughput results (graph on the right), longer bars are better. You can see the BIG-IP WOM increased effective throughput in a range from 38 to 43 Megabits per second for T3 links and a range of 67 to 70 Megabits per second for OC3 links.



CIFS summary

Performance improvements achieved using the BIG-IP WOM

Optimization	T3/100/1	OC3/40/1
TCP	1.18x	1.56x
Compression	48.65x	37.35x
Deduplication and Compression (2 nd Pass)	55.19x	38.85x

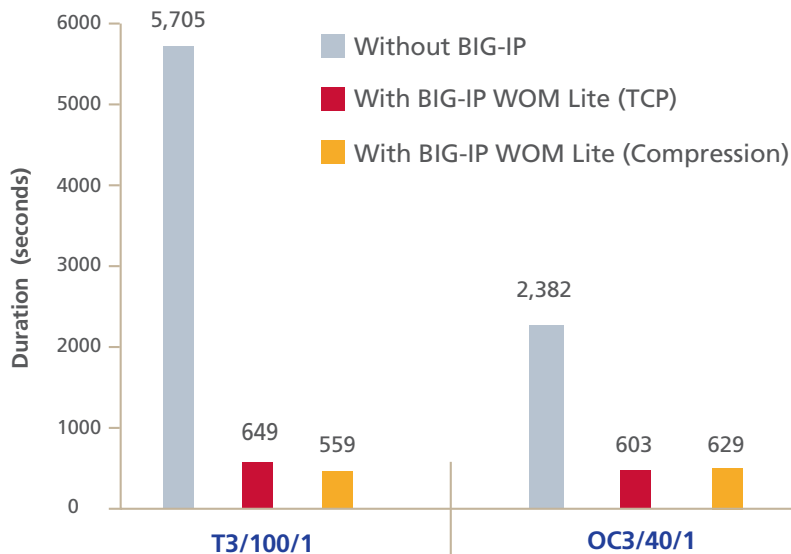


Oracle Streams Replication

By enabling Oracle® Streams optimization on the BIG-IP WOM, Oracle Streams traffic between local and remote endpoints is accelerated. In the following tests, database replication was configured between a primary and secondary database and measured through the BIG-IP WOM tunnel in various network conditions.

For Oracle Streams replication, the following graph does not show throughput because database replication is typically not bandwidth constrained. Additionally, compression provides no additional benefit for higher speed links and was disabled. The network conditions being tested were T3 (45 Mbits per second) with 100 milliseconds of round trip latency and 1% packet loss and OC3 (155 Mbits per second) with 40 milliseconds of latency and 1% packet loss.

For the duration time results, shorter transfer times are better. You can see that BIG-IP WOM accelerated Oracle Streams transfers under these conditions in a range from 4 to 10 times faster than with no acceleration in place.



Oracle Streams Replication summary

Performance improvements achieved using the BIG-IP WOM

Optimization	T3/100/1	OC3/40/1
TCP	9x	4x
Compression	10x	4x
Deduplication and Compression	N/A	N/A

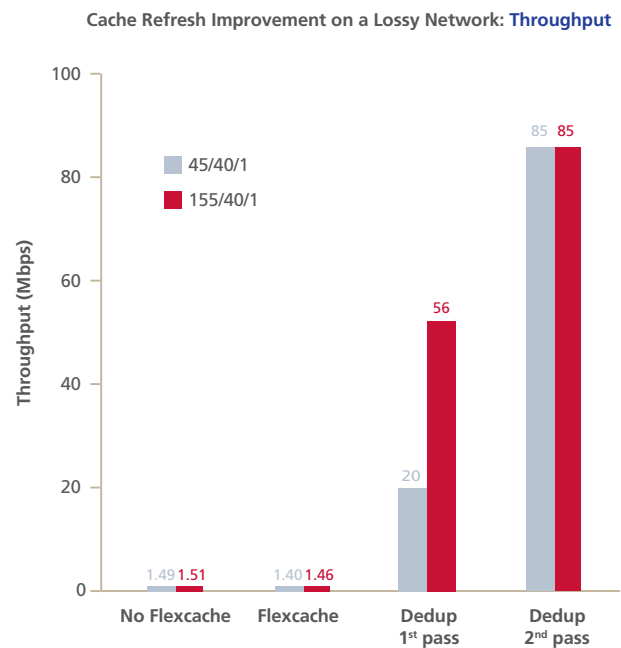
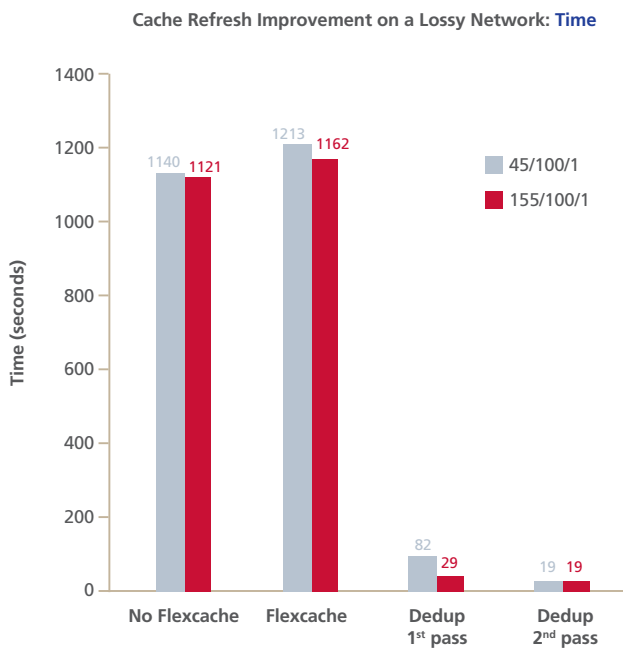


NetApp FlexCache

By enabling NetApp® FlexCache™ optimization on the BIG-IP WOM, NetApp FlexCache traffic between local and remote endpoints is enabled and accelerated.

The use case for NetApp FlexCache is to improve the cold cache refresh times over the WAN. The following tests were performed assuming a 100% cache refresh scenario. The duration of the transfer time was measured along with the effective throughput. The network conditions being tested were T3 (45 Mbits per second) with 100 milliseconds of round trip latency and 1% packet loss and OC3 (155 Mbits per second) with 40 milliseconds of latency and 1% packet loss.

For the duration time results (graph on the left), you can see the BIG-IP WOM accelerated NetApp FlexCache transfers under these conditions in a range from 15 to 60 times faster than with no acceleration in place. For effective throughput results (graph on the right), you can see the BIG-IP WOM increased effective throughput in a range from 1.40 to 85 Megabits per second for T3 links and a range of 1.46 to 85 Megabits per second for OC3 links.



NetApp FlexCache summary

Performance improvements achieved using the BIG-IP WOM

Optimization	T3/100/1	OC3/40/1
Deduplication First Pass	15x	32x
Deduplication subsequent passes	60x	51x

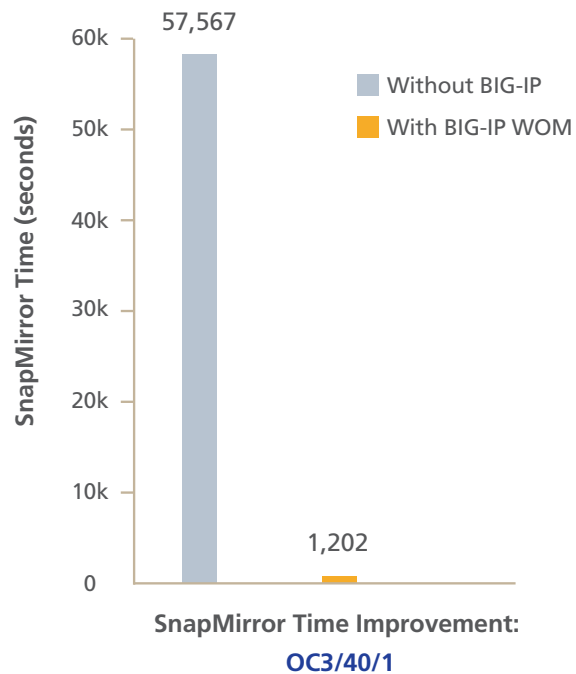
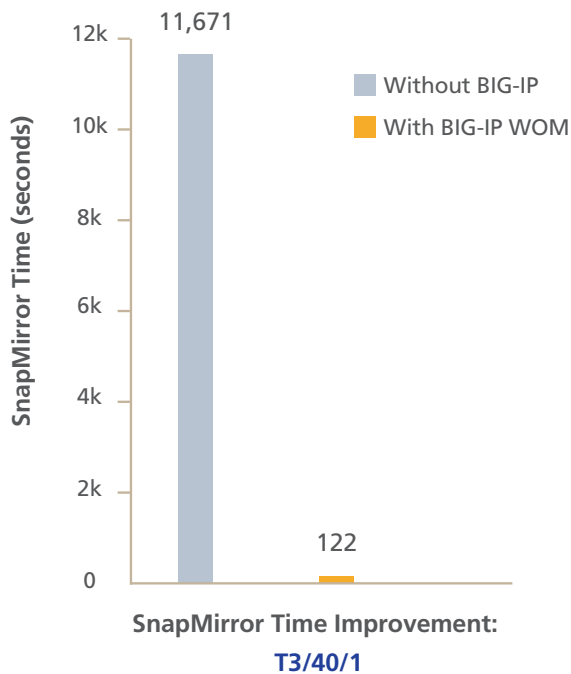


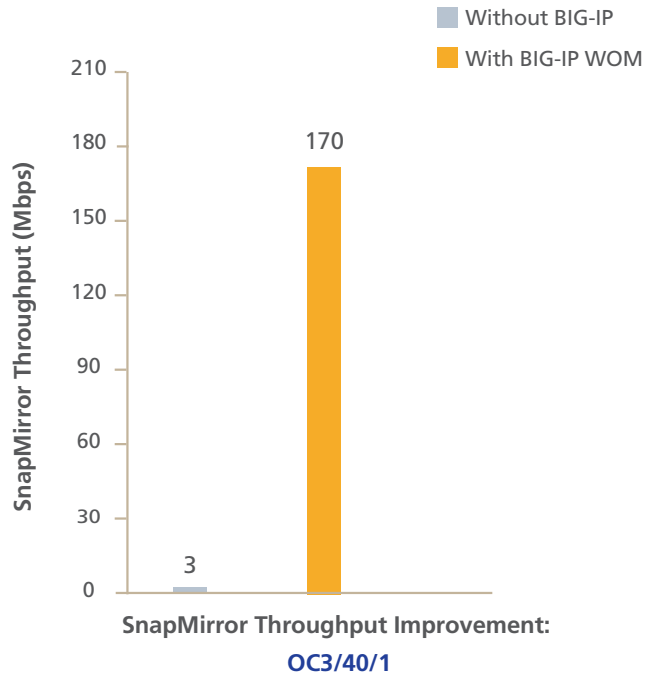
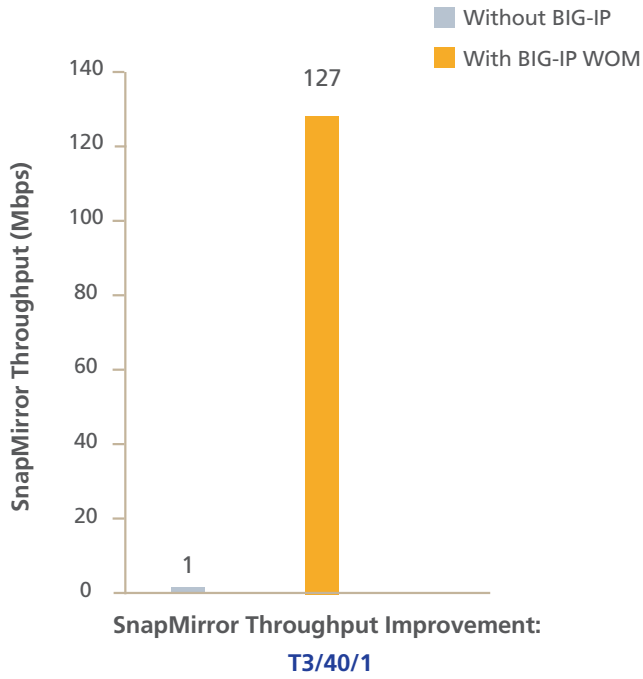
NetApp SnapMirror

By enabling NetApp SnapMirror® optimization on the BIG-IP WOM NetApp SnapMirror traffic between local and remote endpoints is accelerated.

The use case for NetApp SnapMirror is to improve throughput. The following tests were done with NetApp deduplication enabled. The following graphs show that using the BIG-IP with WOM improves the time taken to create a SnapMirror over a WAN. The duration of the SnapMirror time was measured along with the effective throughput. The network conditions being tested were T3 (45 Mbits per second) with 40 milliseconds of round trip latency and 1% packet loss and OC3 (155 Mbits/second) with 40 milliseconds of latency and 1% packet loss.

For the duration time results, shorter transfer times are better. You can see the BIG-IP WOM accelerated NetApp SnapMirror transfers under these conditions in a range from 20 to 95 times faster than with no acceleration (the first graph represents first pass results while the seconds graph represents subsequent passes). For effective throughput results (graph the following page), longer bars are better. You can see the BIG-IP WOM increased effective throughput to 127 Megabits per second for T3 links and 170 Megabits per second for OC3 links (the first graph represents first pass results while the second graph represents subsequent passes).





NetApp SnapMirror summary

Performance improvements achieved using the BIG-IP WOM

Optimization	T3/100/1	OC3/40/1
Adaptive Compression with Deduplication 1 st Pass	20x	42x
Adaptive Compression with Deduplication subsequent passes	95x	48x



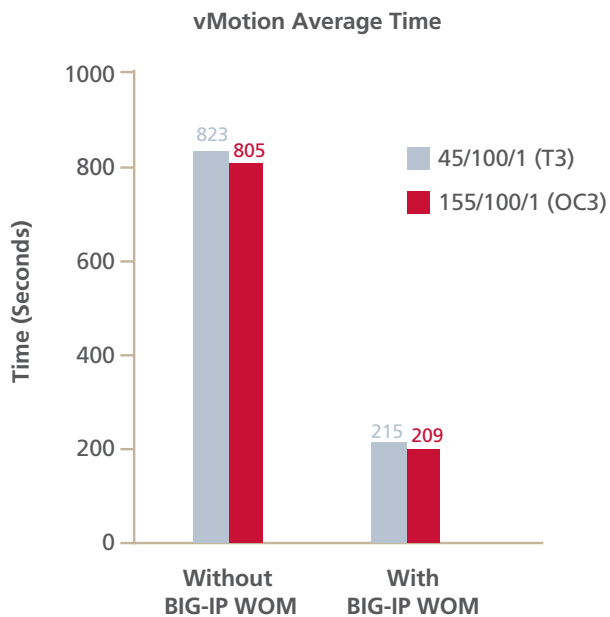
vMotion

By enabling VMware® Live vMotion™ optimization on the BIG-IP WOM, vMotion traffic between local and remote endpoints is enabled and accelerated.

The use case for VMware vMotion is to enable long distance vMotion and improve transfer time. These tests were done with memory-based vMotion. The following results show how using BIG-IP with WOM improves the time taken to transfer a Virtual Machine's memory without dropping any active connections over a WAN. With VMware vMotion, the migration of multiple machines of the same operating system allows deduplication to be more effective than the first pass. For example, the migration of four Windows 2008 Servers results in many deduplication hits and increased migration times for machines.

The duration of the transfer time was measured along with the effective throughput. The machine being transferred was configured with two gigabytes of RAM which was fully active during the test (using a load generator to create memory swapping and the most real world conditions). The network conditions being tested were T3 (45 Mbits per second) with 100 milliseconds of round trip latency and 1% packet loss and OC3 (155 Mbits per second) with 100 milliseconds of latency and 1% packet loss.

For the duration time, shorter transfer times are better. You can see the BIG-IP WOM accelerated VMware Live vMotion transfers under these conditions in a range from 3.8 to 3.9 times faster than with no acceleration in place. Without the BIG-IP WOM on both links however, the vMotion transfer failed in greater than 50 percent of all tests. With the BIG-IP WOM, the vMotion success rate was 100%.





vMotion summary

Performance improvements achieved using the BIG-IP WOM

Optimization	T3/100/1	OC3/40/1
First Pass with Deduplication, TCP optimizations and compression	3.8x	3.9x



Appendix: Hardware and Software Specifications

Hardware (HTTP)

Product	Platform	Functionality	Quantity
BIG-IP	3900	WAN Optimization Module	2
Cisco® Switch	CAT6k	Routing/Switching	1
Dell® Servers	2950	ESX Server	2
Linux Client and Server VMs running on VMware ESX™			2
LANForge®	WJ-400	WAN Simulator	1

Software (HTTP)

Platform	Version	Description
BIG-IP	10.2	RTM Build
ESX	VMware ESXi 4.0.0 build-164009	Host OS
Microsoft® Windows®	MS Windows Server 2003 R2, SP2, Enterprise x64 Edition	Guest OS
CAT6K	12.2(18)SXD7b	Release Software (fc1)
LANForge	5.0.2	WAN Simulator

Hardware (FTP)

Product	Platform	Functionality	Quantity
BIG-IP	3900	WAN Optimization Module	2
Cisco Switch	CAT6k	Routing/Switching	1
Dell Servers	2950	ESX Server	2
Linux Client and Server VMs running on ESX			2
LANForge	WJ-400	WAN Simulator	1



Software (FTP)

Platform	Version	Description
BIG-IP	10.1	RTM Build
ESX	VMware ESXi 4.0.0 build-164009	Host OS
Linux	Linux version 2.6.18-8.el5PAE (gcc version 4.1.1 20070105 (Red Hat 4.1.1-52))	Red Hat
CAT6K	12.2(18)SXD7b	Release Software (fc1)
LANForge	5.0.2	WAN Simulator

Hardware (Robocopy)

Product	Platform	Functionality	Quantity
BIG-IP	3600	WAN Optimization Module	2
Cisco Switch	CAT6k	Routing/Switching	1
Dell Servers	2950	ESX Server	2
Windows 2003 servers and XP client and server VMs running on ESX			2
LANForge	WJ-400	WAN Simulator	1

Software (Robocopy)

Platform	Version	Description
BIG-IP	10.2	RTM Build
ESX	VMware ESXi 4.0.0 build-164009	Host OS
Microsoft Windows	MS Windows Server 2003 R2, SP2, Enterprise x64 Edition	Guest OS
CAT6K	12.2(18)SXD7b	Release Software (fc1)
LANForge	5.0.2	WAN Simulator



Hardware (Oracle SQL Replication)

Product	Platform	Functionality	Quantity
BIG-IP	3600	WAN Optimization Module	2
Cisco Switch	CAT6k	Routing/Switching	1
Dell Servers	2950	ESX Server	2
LANForge	WJ-400	WAN Simulator	1

Software (Oracle SQL Replication)

Platform	Version	Description
BIG-IP	10.2	RTM Build
ESX	VMware ESXi 4.0.0 build-164009	Host OS
Microsoft Windows	MS Windows Server 2003 R2, SP2, Enterprise x64 Edition	Guest OS
SQL Server	Oracle 11g	SQL Version
CAT6K	12.2(18)SXD7b	Release Software (fc1)
ORACLE 11g	12.2(18)SXD7b	Release Software (fc1)
SQL Version	5.0.2	WAN Simulator

Hardware (vMotion)

Product	Platform	Functionality	Quantity
BIG-IP	6900	WAN Optimization Module	2
Cisco Switch	CAT6k	Routing/Switching	1
Dell Servers	2950	ESX Server	2
Linux Client and Server VMs running on ESX			2
LANForge	WJ-400	WAN Simulator	1

Technical Brief

BIG-IP WAN Optimization Module Performance

Software (vMotion)

Platform	Version	Description
BIG-IP	10.2	RTM Build
ESX	VMware ESXi 4.0.0 build-164009	Host OS
Microsoft Windows	MS Windows Server 2003 R2, SP2, Enterprise x86 Edition	Guest OS
LANForge	5.0.2	WAN Simulator

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 www.f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com

