

HP StorageWorks Enterprise File Services (EFS) WAN Accelerator

Frequently Asked Questions



General questions	3
What value will your product provide?	3
Why is improving the performance of applications over WANs important?	3
How exactly does slow application performance over WANs cost money?	3
Why are WANs slower than LANs?	3
What do you mean by "...the nature of network protocols"?	4
Will HP EFS WAN Accelerators help our enterprise to centralize IT infrastructure like storage and file servers?	4
What if we have already embarked on a "re-centralization" project? Can HP EFS WAN Accelerators help us?	4
What if we cannot centralize because we need local infrastructure for other reasons?	4
What are key the benefits of HP EFS WAN Accelerators?	5
Can I buy more bandwidth to increase throughput?	5
Product questions	6
Is your product software or hardware based?	6
Does it support Windows File Shares over WANs? What about Exchange?	6
What about FTP? Will HP EFS WAN Accelerators speed that up also?	6
Will it help my web-based business applications?	6
Where do HP EFS WAN Accelerators get deployed?	6
Do I need to configure the client desktop machines in my remote branch offices?	6
Product functionality	7
What does your EFS WAN Accelerator actually do?	7
How do you keep track of all the files as they get modified over time?	7
Are the files that I open and read at the Remote Branch Office editable (writeable)?	7
How much faster can I expect my applications to be?	7
Do HP EFS WAN Accelerators work with remote backup systems?	8
Don't lots of products have "differencing" algorithms? How is this different?	8
What if your device fails?	8
Do I have to centralize my file server or data storage architecture to take advantage of this?	8



How big is the disk on the HP EFS WAN Accelerator?	8
What is the overhead on your algorithms? Does that slow down my application?	8
Does this system optimize video content over the WAN?	9
Do HP EFS WAN Accelerators encrypt traffic sent over the WAN?	9
How do you handle encrypted traffic?	9
Does this work with HP Storage Server (NAS) and SAN solutions?.....	9
What protocols do you support?	9
Will you support other protocols in the future?	9
Where can I get a more detailed technical explanation of how this works?	9
Competitive information	10
Is this a file caching system?	10
Why is it better than a file cache?	10
HP SDR technique sounds like compression technology. Is it?	10
What about data that has never been sent over the network before, does it go slower through the HP EFS WAN Accelerator because of overhead?.....	11
Availability	12
For more information	13

General questions

What value will your product provide?

HP StorageWorks Enterprise File Services (EFS) WAN Accelerators increase the performance of applications that run on wide area networks (WANs). More specifically, EFS WAN Accelerators radically improve the performance of client-server TCP-based applications and protocols (for example, Messaging Application Programming Interface [MAPI, the Microsoft® Exchange protocol], distributed software development, Common Internet File System [CIFS, Windows® file system], Lotus Notes, Web-based Apps, and so forth) over WAN links by defeating the effects of limited bandwidth and high latency without introducing a new file system or protocol.

In addition, EFS WAN Accelerators provide best-in-class WAN compression, which is important for customers who are trying to delay or avoid a WAN upgrade.

Why is improving the performance of applications over WANs important?

Slow performance over WAN links is one of the root causes of the high cost of IT at distributed enterprises. Improving the throughput over those links will enable significant cost reductions and productivity improvements for most large enterprises.

How exactly does slow application performance over WANs cost money?

Slow performance over WANs costs money in two primary ways:

- It forces enterprises to buy redundant equipment (file servers, mail servers, and storage) for Remote Branch Offices (RBOs) as well as ever-increasing amounts of bandwidth to compensate for the poor performance.
- Slow throughput affects productivity directly when files sizes are large, distances are long, or the bandwidth is limited.

It is not uncommon for cutting-edge enterprises to resort to the modern equivalent of “sneaker net”: burning CDs and sending them by FedEx. Many enterprises endure multi-hour waits for FTPs of critical data like drawings, plans, CAD files, tape-outs, and so forth.

Degraded network performance means that many common business processes and applications that function fine on a LAN work very poorly or not at all over WANs. To get around the problem, most enterprises deploy local file servers, mail servers, storage, backup systems, and so forth in virtually every remote branch office of any size. This costs money up front, and more to maintain over time.

Why are WANs slower than LANs?

The fundamental characteristics of WANs are different than local area networks (LANs) in two important ways:

- WANs have much more limited bandwidth (typically only 1% to 0.05%). For example, a LAN typically has at least 100 Mbps of bandwidth and often 1 Gbps (1,000 Mbps). A typical branch office will have anywhere from 64 kbps up to T1 (1.5 Mbps), and only infrequently higher bandwidth.
- WANs also have much higher latency (typically 100 to 1,000 times higher). Latency is the time it takes a packet to travel from one end of the network to the other, also known as “round trip time.”

This is not only due to the routing inefficiencies inherent in the Internet but also to the speed of light, which fundamentally limits how fast a packet of data can get from one place to another.

These two factors dramatically degrade the throughput of traffic on WANs because of the nature of network protocols.

What do you mean by “...the nature of network protocols”?

There are two layers or protocols that affect applications on WANs: the transport protocol, most commonly TCP, and the application protocol, which rides on top of the transport. Either or both of these can be the bottleneck when applications are run over a WAN, even if you have plenty of bandwidth.

Many commonly used protocols, for example, CIFS (used by Microsoft Windows), NFS, MAPI, HTTP, TCP, FTP, and many others, have a fixed “window size.” This means that on each trip or transaction, there is a maximum amount of data that can be sent. If latency is low like it is on a LAN, each round trip with that limited data happens extremely quickly, so no one notices. However, on a WAN, each round trip takes much longer, so many of the applications that depend on those protocols break down.

In addition to the fixed window size, many of these protocols are extremely “chatty,” which means they take many round trips to accomplish seemingly simple tasks. This chattiness, when combined with the fixed window size, means that the net throughput is severely limited in high-latency environments like WANs.

As latency increases, the throughput decreases at a rate proportional to the inverse of round trip time ($1/RTT$). The decay is hyperbolic (it decreases very quickly after you pass a critical latency threshold).

Will HP EFS WAN Accelerators help our enterprise to centralize IT infrastructure like storage and file servers?

Yes. EFS WAN Accelerators overcome the degradation caused by both high latency and limited bandwidth, which enables companies to centralize key IT infrastructure without the degradation in performance that they would otherwise have to endure.

What if we have already embarked on a “re-centralization” project? Can HP EFS WAN Accelerators help us?

Yes, definitely. For example, many enterprises have taken the first step by centralizing Microsoft Exchange servers in an effort to contain the costs of distributed infrastructure, but have found that the penalty they pay is degraded performance for their increasingly distributed users as they are then subject to the low bandwidth, high-latency environment of a WAN. In many cases, the performance of the centralized deployment is unacceptably slow.

EFS WAN Accelerators will dramatically improve the performance of centralized applications and business processes over the WAN, which will improve productivity and make those architectures perform much better.

What if we cannot centralize because we need local infrastructure for other reasons?

That is fine, too. HP solutions can be used to improve the performance of distributed systems, even if you do not want to or cannot centralize.

What are key the benefits of HP EFS WAN Accelerators?

EFS WAN Accelerators deliver several benefits, depending on the focus of the customer. Here are the most common:

- **Enables universal site consolidation**—Unlike caching appliances that enable file server consolidation, EFS WAN Accelerators enable the consolidation of file servers, Microsoft Exchange servers, Notes servers, backup equipment, and network attached storage (NAS).
- **Reduces the cost of IT**—EFS WAN Accelerators enable centralization of key infrastructure like data storage, mail, and file servers, which is widely known to reduce the costs of building, maintaining, and servicing infrastructure like servers, storage, and other key pieces of technology.
- **Reduces the need for WAN bandwidth**—By recognizing all common data sent across the WAN, storing it, and reusing it instead of resending the data, EFS WAN Accelerators can dramatically reduce the amount of data traversing the WAN, which means customers can postpone bandwidth upgrades, and possibly downgrade existing links.
- **Increases WAN throughput and productivity**—Many large enterprises have a hard time sharing large files across long distances, particularly those in engineering, design, construction, software development, and other data-intensive sectors. EFS WAN Accelerators can dramatically increase the throughput of WAN links through the combination of its latency and bandwidth optimizations (see “What does your EFS WAN Accelerator actually do?”), which translates into time saved by users.
- **Eliminates the WAN as an obstacle to distributed systems design and operations**—Today, interconnecting distant offices is a difficult and expensive problem if you want high performance. EFS WAN Accelerators eliminate many of the issues that make it a difficult challenge, which means you can focus on designing and building the best architecture, and not have to settle for creatively setting up systems to get around WAN issues.
- **Improves the performance of most distributed and collaborative applications**—Because of its application-independent bandwidth optimization, EFS WAN Accelerators will automatically eliminate the common data sent by many applications, including Web-based client-server business apps, mail, document management systems, file sharing, and so forth.
- **Better backup of distributed systems**—Enterprises that currently do back up on local tape backup systems at remote branch offices can now move the backup process to the datacenter. Because of the huge reduction in traffic sent over WAN links when EFS WAN Accelerators are deployed, the concern about bringing down the WAN with too much traffic becomes moot. This leads to lower costs and much better control over the backup process and the data itself.
- **Better control of IT processes**—If enterprises use EFS WAN Accelerators to enable a centralization move of servers and storage, they not only reduce their costs but improve their ability to control IT processes and to implement more consistent policies.
- **Lower security vulnerability**—Centralizing servers and storage also means that there are fewer points of entry into the enterprise’s network, which means less vulnerability to attack.

Can I buy more bandwidth to increase throughput?

You can, but it may not help your throughput. Most people think throughput is roughly equivalent to bandwidth, but that is not always the case on WANs.

Throughput is governed not only by the available bandwidth, but by other factors, too, in particular by latency and packet loss rate. Even if you compress files down to zero bits, you cannot eliminate latency because the speed of light is fixed.

Many protocols are serial in nature, which means that each data transmission must wait for the data to arrive before sending the next batch.

Furthermore, the amount of data sent on each transaction is often limited by the protocol. This means that the maximum throughput is proportional to W/RTT , where “W” is the maximum window size, and “RTT” is the round trip time. For a more complete explanation, refer to the “It is not about bandwidth” white paper.

You may have experienced an example of this problem yourself if, for example, you doubled your bandwidth but did not see 2x the throughput.

Product questions

Is your product software or hardware based?

HP EFS WAN Accelerators utilize software-based technology delivered and packaged as appliances based on HP ProLiant hardware. The notion of an appliance is a special purpose computer that needs no configuration or software installation; it is ready to go when the customer receives it.

Does it support Windows File Shares over WANs? What about Exchange?

Yes, EFS WAN Accelerators optimize all TCP traffic. In addition, additional optimizations are available to support Windows (CIFS) and Exchange (MAPI) in the first release. They will support other key protocols in future releases.

What about FTP? Will HP EFS WAN Accelerators speed that up also?

EFS WAN Accelerators support FTP as well. Using the same mechanisms to recognize commonality across files and to predict transactions will dramatically help the throughput of file transfers by only sending bits that have never been sent before and eliminating unnecessary round trips over the WAN.

Will it help my web-based business applications?

Yes, especially if there is repeatable or predictable data requested by the application.

Where do HP EFS WAN Accelerators get deployed?

EFS WAN Accelerators are deployed in remote offices and in the datacenter across WAN links. They need to be deployed on both ends of each link that is to be optimized. In the datacenter, EFS WAN Accelerators can be clustered to support many remote branch offices.

In the remote office, the appliance is deployed in the data path by way of a dual port Ethernet NIC. In the datacenter, the product is deployed more like a server, out of the data path.

Do I need to configure the client desktop machines in my remote branch offices?

No, there is no configuration required. EFS WAN Accelerators are completely transparent to end users.

Product functionality

What does your EFS WAN Accelerator actually do?

EFS WAN Accelerators intercept TCP requests sent across a WAN link and perform two optimizations that can operate independently or in a synergistic way.

- **Scalable data referencing (SDR)**—The first optimization EFS WAN Accelerators perform is to send as little data as possible, yet still fulfill the client request. It does this by maintaining a “Data Store,” which is a novel method of storing very granular pieces of data and references to that data so they reside on both ends of a WAN link in a way that they can be re-used to rebuild files very quickly.
- **Transaction prediction**—The second thing EFS WAN Accelerators do is to try to avoid taking round trips whenever possible. The HP EFS WAN Accelerator observes the behavior of known protocols or applications (like Windows File System or CIFS, for example), and takes advantage of the highly repetitive and predictable behavior exhibited by them. In this way, EFS WAN Accelerators can predict what will happen next, and send data before a client asks for it.

By combining these two techniques, requests can be fulfilled by a server before they are even requested by a client. Furthermore, data that ordinarily is sent across the network is often already present on the client side (in the Data Store) because of prior activity.

How do you keep track of all the files as they get modified over time?

We do not keep track because we do not need to. Our architecture only keeps track of the “references” for our own purposes. The system has no knowledge of the files, filenames, directories, file systems, and so forth. SDR keeps track of data and references and then stores those references in such a way that they can be retrieved very quickly if needed to fulfill a request.

Are the files that I open and read at the Remote Branch Office editable (writeable)?

Yes. The system is fully symmetric. All changes are sent upstream and downstream transparently. You do not have to do anything different. That means that in a branch office, you can open a file that is stored in the datacenter and in many cases the bits comprising the files would be stored locally on the HP Data Store. As you edit and save the file, all changes are sent in real time back to the datacenter, ensuring that others who need the file always have the latest version.

How much faster can I expect my applications to be?

That depends a great deal on the workload of your users. The more in-common data there is in your company’s work, and the more iterative and repeatable the workload is, the more EFS WAN Accelerators can optimize your interactions.

Likewise, the higher the latency and the lower the bandwidth of your WAN connections, the greater potential there is for HP technology to make a dramatic difference. Expected benefits should be in the range of 5 to 100 times faster, depending on the content, workload, number of users, and so forth.

Do HP EFS WAN Accelerators work with remote backup systems?

Yes, EFS WAN Accelerators are very effective for remote backup because by their very nature they are quite iterative. The actual amount of data that needs to be sent using EFS WAN Accelerators will, in many cases, be a tiny fraction of the actual data being backed up.

With EFS WAN Accelerators you can even just move your servers from the remote office to the datacenter where they can be backed up using your centralized equipment, and where the process can be overseen by qualified IT professionals.

Don't lots of products have "differencing" algorithms? How is this different?

The HP approach to finding the common elements that make up data is very different.

- First of all, a very fine-grained data is used. This gives much better granularity, which means EFS WAN Accelerators can incorporate data changes without sending as much data over the WAN.
- Second, SDR algorithms are extremely good at handling changes in the underlying files and adapting in a way that minimizes the amount of data that changes.
- Third, the referencing scheme is independent of the applications, unlike differencing technologies that are typically specific to one application (like CAD, backup, and so forth).

What if your device fails?

The product can be deployed in a high-availability configuration (redundant boxes) so that a failover would occur to the backup box.

Each EFS WAN Accelerator has a dual port Ethernet NIC that has an electromechanical relay that cuts over to wire in the event of power failure or a software crash. In this case, if you do not have a redundant EFS WAN Accelerator deployed, you are back to your current state of an un-optimized WAN link.

Do I have to centralize my file server or data storage architecture to take advantage of this?

No, it is completely up to you. If you want to centralize Exchange or file servers, storage, and so forth, you can. If you do not want to or cannot because of network or other operational issues, HP EFS WAN Accelerators can be used as a performance and productivity improvement tool.

How big is the disk on the HP EFS WAN Accelerator?

There are five different models, designed for small, medium, and large offices, as well as for datacenters. The raw storage on the appliances ranges from 250 GB to 1.8 TB.

What is the overhead on your algorithms? Does that slow down my application?

There is some overhead to perform the SDR, but it is miniscule compared with the optimization achieved with our standard compression, and even smaller when compared to the benefits of the SDR algorithms.

Does this system optimize video content over the WAN?

EFS WAN Accelerators will handle on-demand video just like any other file that may be requested more than once. Live content, however, will not be optimized by this system.

Do HP EFS WAN Accelerators encrypt traffic sent over the WAN?

No, but EFS WAN Accelerators can be integrated into encrypted environments (see “How do you handle encrypted traffic?”).

The traffic sent between two EFS WAN Accelerators would not be in a very usable state as it is sent in the form of references and pieces of data, and in many cases since only tiny fractions of entire files are sent, it would be difficult to make use of the data. However, HP recommends that customers implement their own encryption solutions.

How do you handle encrypted traffic?

EFS WAN Accelerators must be deployed “outside” of encrypted traffic. The product needs to be able to read the file to detect repeating patterns in the data and to predict transactions, so if the data is already encrypted before it gets to the EFS WAN Accelerator, the optimizations will not be possible.

All the customer has to do is deploy EFS WAN Accelerator in the right place in the network. HP professional services are available to help customers deploy the product correctly.

Does this work with HP Storage Server (NAS) and SAN solutions?

EFS WAN Accelerators are only relevant for distributed systems operating over WANs because that is where customers have the low bandwidth and high latency. EFS WAN Accelerators can be used to interconnect Storage Server (NAS) “islands” that need to be accessed by distant users, but it is not really relevant for “back office” storage technologies that are interconnected by very high bandwidth, low-latency technology.

What protocols do you support?

The first version of the EFS WAN Accelerator supports all TCP-based traffic for bandwidth optimization. For latency optimization (Transaction Prediction) EFS WAN Accelerators support CIFS, MAPI, TCP, FTP, HTTP, and any web-based application (like enterprise ERP applications). Future releases (which may be downloaded to existing deployments through an over-the-wire upgrade facility) will include additional functionality.

Will you support other protocols in the future?

Yes, over time, HP will add new capability to the HP EFS WAN Accelerator software. All new features will be downloadable and customers are entitled to one year of free software upgrades from HP Support.

Where can I get a more detailed technical explanation of how this works?

For more information, visit <http://www.hp.com/go/efs> to access datasheets, QuickSpecs, and solution white papers. For additional HP StorageWorks EFS Accelerator white papers, visit <http://h18006.www1.hp.com/storage/efswhitepapers.html>.

Competitive information

Is this a file caching system?

No, EFS WAN Accelerators are not file caching devices. File caches store copies of files for fast access by the second requester, whereas EFS WAN Accelerators store data and “references,” not the files themselves. Every file access request is sent to the originating server, which ensures there is absolutely no chance that an out-of-date copy is ever used.

EFS WAN Accelerators have no knowledge of the files from which the references were created. The references can be thought of as the common set of building blocks that can be used to rebuild any file that has been sent across the network. References can belong to more than one file.

Why is it better than a file cache?

- **Broader support for more data types**—The primary benefit is that EFS WAN Accelerators optimize all TCP traffic, not just files like a file cache. That means that all your company’s email, web traffic, applications, and files are optimized for WAN performance, and you do not have to buy a separate box for mail, web traffic, files, and applications.
- **Easier configuration**—EFS WAN Accelerators are also much simpler and easier to deploy and configure. There is no proxy configuration required as there is in a caching product. EFS WAN Accelerators find each other automatically using HP auto-discovery algorithms.

HP SDR technique sounds like compression technology. Is it?

The system uses some techniques that are similar to those used by some compression systems. It may be thought of as a hybrid of compression and other technologies, but it has several important differences:

- **Protocol latency optimization at the application level**—While some companies claim to handle latency, they do it at the TCP level, not at the application level. Often it is the application level that is the rate limiting step, not the TPC level and those types of solutions do not address the effect of latency on application protocols. For many real-world scenarios, latency bounds performance, regardless of the bandwidth available. The performance of protocols like CIFS, for example, decays very quickly as latency increases: On a cross country DS3 with 100-ms latency, for example, a Windows “copy” will see throughput in the range of 600 kbps, even though there are nominally 45 Mbps of bandwidth available.
- **Enables site consolidation**—Compression cannot be the basis of a file server, mail server, or NAS consolidation project. If your company is trying to centralize infrastructure without giving up performance, compression alone will not deliver the benefit you are looking for because of the way latency limits throughput.
- **Disk-based architecture**—EFS WAN Accelerators store the data on disk, not in memory like most compression or WAN optimization solutions. This means that it is possible to have a much larger Data Store than memory-based compression systems. Because they are on disk, which is relatively cheap, the references created and stored are stored almost indefinitely (actually until the disk gets full and they are finally purged based on their popularity and age).
- **Very long time windows**—The long life of references means that recognizing common data extends over very long time periods, potentially days, weeks, months, or even years, depending on the workload. Compression systems operate on much smaller time scales and data sets. This means that common elements of files sent weeks or months earlier will still be recognized, and they will not be sent again across the WAN.

- **Higher effective compression**—Unlike packet-level compression technology or “WAN optimization” products, these “compression ratios” are not bounded. For example, a typical WAN optimization technology is limited to about 3:1 or 4:1 compression ratio. If a user copied a 200-MB file over a WAN, the EFS WAN Accelerator could potentially send zero packets if the same or a similar version of the file was sent weeks before. The best a packet-level system would do would be to send about 50 MB instead of 200 MB—it would have no knowledge of the same file (or even a closely related one) that was sent weeks earlier.
- **Very large file sizes**—EFS WAN Accelerators have no practical limits on the size of files it can handle. Both file caches and compression systems have limits on the size of the files they can optimize.

What about data that has never been sent over the network before, does it go slower through the HP EFS WAN Accelerator because of overhead?

New data will actually be transmitted more quickly because EFS WAN Accelerator implements a standard “L-Z” compression algorithm on data that has never been seen before. As is typical with this method, the data is reduced by 3:1 or so. This reduction in data overwhelms the small overhead associated with SDR. After the new data has been sent through the system and processed using SDR, it never has to be sent again.

Availability

HP StorageWorks EFS WAN Accelerators are available in five configurations plus an EFS WAN Accelerator Manager to meet the needs of any site. The following table summarizes the specifications.

Model	DL320-510	DL320-1010	DL320-2010	DL380-3010	DL380-5010	DL320-M25
Part number	391683-B21	391684-21	391685-21	391686-B21	391687-B21	391688-B21
Profile	1U	1U	1U	2U	2U	1U
WAN capacity	512 kbps	2 Mbps	4 Mbps	10 Mbps	45 Mbps	–
Raw disk store	250 GB	250 GB	500 GB	1.2 TB	1.8 TB	80 GB
Optimized TCP connections	200	625	1300	2400	4000–5000	–
User capacity (consolidated)*	Up to 25	Up to 80	Up to 200	Up to 1,200	Up to 1,600	–
User capacity (non consolidated)*	Up to 150	Up to 500	Up to 1,100	Up to 3,500	Up to 6,000	–
Fan-out***	–	–	Up to 20:1	Up to 50:1	Up to 150:1	Up to 500:1
RAID	No	No	Yes	Yes	Yes	No
Dual power supplies	No	No		Yes	Yes	No
Hot-swappable disks	No	No	No	Yes	Yes	No
GigE fail-through	Yes	Yes	Yes	Yes	Yes	–
Auto-discovery	Yes	Yes	Yes	Yes	Yes	Yes
Management interface	WEB UI, CLI, SNMP	WEB UI, CLI				

* This is a guideline only based on remote site servers consolidation environments. Actual number of supported users will vary based on use case.

** This is a guideline only based on environments where servers remain in remote site. Actual number of supported users will vary based on use case.

*** This is a guideline only and covers remote site EFS WAN Accelerators to datacenter WAN Accelerators or WAN Accelerator Managers. Actual fan-out will vary based on use case.

For more information

For more information on HP StorageWorks EFS WAN Accelerators, visit:

<http://www.hp.com/go/efs>

For additional HP StorageWorks EFS Accelerator white papers, visit:

<http://h18006.www1.hp.com/storage/efswhitepapers.html>

© 2005 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation.

5983-2464EN, 05/2005



i n v e n t