White Paper

Intel Information Technology

Computer Manufacturing

WAN Performance and Optimization

Optimizing WAN Performance for the Global Enterprise

To improve throughput on Intel wide area network (WAN) circuits, Intel IT conducted a proof of concept (PoC) study to investigate new WAN optimization technologies and products. We achieved significant benefits to WAN throughput in our test network and in a pilot production network using devices that incorporated a variety of optimization techniques.

Tim Verrall, David Wescott, and Lilin Xie, Intel Corporation

May 2006

IT@Intel

Executive Summary

To enhance wide area network (WAN) throughput, Intel IT conducted a proof of concept (PoC) study to investigate new WAN optimization technologies and products. We achieved significant benefits to WAN throughput in our test network and in a pilot production network using devices that incorporated a variety of optimization techniques.

WAN optimization techniques applied in the right place in the network can greatly mitigate performance problems from protocol constraints and offset effects of high line latencies.

Intel depends on high throughput on its WAN circuits to stay efficient and competitive. Throughput is a result of line latency and protocol constraints; although line latency is unchangeable, new WAN optimization technologies can help improve protocol efficiencies using a variety of techniques to increase throughput.

Through an industry survey, laboratory testing, and pilot production deployment, we investigated two products that offer combinations of WAN optimization technologies in a single appliance. These two products delivered impressive results to improve throughput on high latency circuits.

Testing several file types over a variety of simulated high-latency circuits revealed the following benefits:

- Up to 75 percent file compression using basic techniques
- 99 percent file size reduction using advanced compression
- 99 percent file size reduction for all file types after an initial file transfer
- Up to twelve times faster Transmission Control Protocol (TCP) transfers
- Up to five times faster Common Internet File System (CIFS) communications
- Classification of a wide range of traffic for analysis

WAN optimization techniques applied in the right place in the network can greatly enhance performance and offset the effects of high line latencies.

Contents

The Business Challenge

WAN performance is critical to Intel's business around the globe. Employee productivity and efficient company operations depend on high data throughput across thousands of miles of copper and fiber optic cables. Transmission delays result in low throughput, which can affect time sensitive decisions and production schedules, impacting bottom line profitability.

Throughput is the result of line latency—measured as round-trip time (RTT), the time it takes a signal to go from source to destination and back—and protocol efficiency. Increasing bandwidth on our circuits adds capacity, but it doesn't improve line latency, and faster circuits are more expensive, but not necessarily more efficient.

Intel IT manages many WAN circuits through multiple domestic and international telecommunication carriers. To help manage throughput, we keep line latency to under 80 milliseconds (ms) on domestic circuits and under 300 ms on international circuits.

Line Latency Impact

Line latency is a fixed characteristic of a cable's physical properties, the distance from a signal's origin to its destination. Longer distances will always exhibit higher latencies for the same kind of cable. Line latency has a significant impact on data throughput. Industry research has shown that TCP throughput degrades significantly as the line latency increases (Figure 1 on the next page). For a T3 circuit (45 Mbps), the TCP throughput starts out at the available line rate, but only for low latencies. At higher latencies, the throughput begins to decay rapidly. This effect is so dramatic

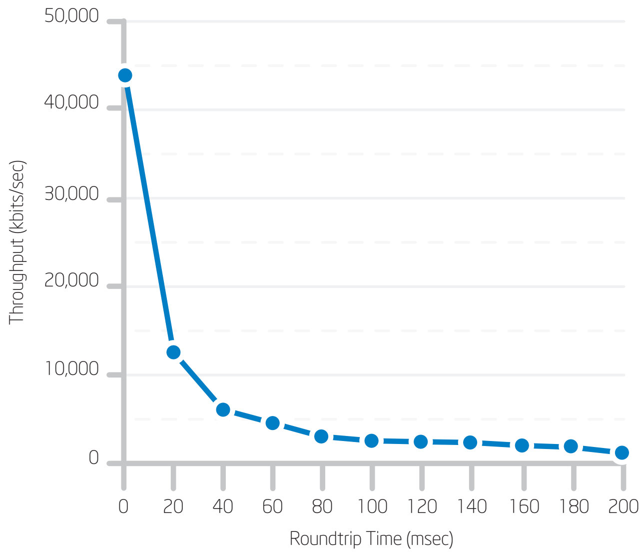

Figure 1. Impact of line latency on throughput.
Single TCP connection's throughput with 32K window size.

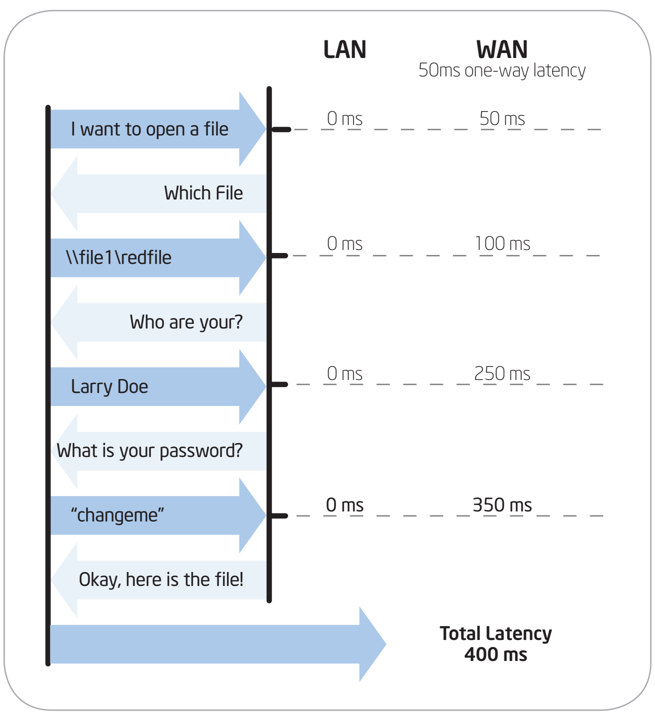

Figure 2. Latency impact from chattiness.

that at 100 ms of latency, TCP is able to use only 10 percent of the available bandwidth.[1]

In the past, we've been able to reduce the impact of line latencies by choosing high quality circuits and finding the shortest routes from sources to destinations. Network protocol constraints, however, create other issues.

Protocol Impact

The protocol governing a network transaction can significantly affect throughput. Examples of protocols include TCP, CIFS, HTTP, File Transfer Protocol (FTP), and others. Protocols define important transmission characteristics:

- Handshaking. The method by which two computers complete communications. Some protocols are very chatty, requiring a lot of communication just to secure a connection and start data transfers. Chattier protocols are less efficient and reduce throughput.
- Data payload window. The amount of unacknowledged data allowed in each communication. Smaller windows require more transmissions—with each transmission's round trip overhead—to transfer a set amount of data. Larger windows are more efficient, with fewer communications and WAN round trip delays.

The CIFS protocol is an example of a chatty protocol. Using CIFS, it can take several communications across the network to request a file, authenticate the user, and deliver the data (Figure 2). Unfortunately, many applications use CIFS, even over high-speed fiber optic WAN connections.

For example, using CIFS, an Intel employee in Central America initiated a query to a database server in Southeast Asia. The query was not presented to the database for almost 20 seconds. Traffic analysis showed that it took 19.5 seconds and 41 communications to submit this command, due to the severe overhead imposed by CIFS. The high network latency, in combination with the large number of communications, reduced the network throughput for this transaction to less than 23 Kbps.

1 Riverbed Technology, Inc. white paper, "It's not about bandwidth – why adding bandwidth or compression doesn't solve the problem of application performance on wide area networks" (2004).

Proof of Concept

Over the last few years, we've tracked many significant advances in the networking and telecommunications industry. We wanted to determine if these advances offer WAN improvements we could take advantage of. We conducted a proof of concept (PoC) study to examine new WAN optimization technologies and test them for application to Intel IT operations.

We executed our PoC in three phases:

- Phase One: Industry Survey. We surveyed the network and telecommunications industry to get an overview of current WAN optimization technologies and products. We then selected two products for testing in our laboratory.
- Phase Two: Laboratory Testing. We benchmarked the two products to evaluate the potential benefits, understand implementation strategies, and provide direct feedback to the product suppliers regarding desirable features and capabilities.
- Phase Three: Production Network Analysis. We selected an Intel network location, installed the products, and evaluated the results.

Phase One: Industry Survey

We researched new WAN technologies, the latest products offered for optimization, and the recommended best practices from industry analysts.

WAN Optimization Technologies

Over the last few years, WAN optimization technologies have evolved from traffic management, caching, and web load distribution to a much more comprehensive feature set. Newer features include compression, basic caching, bit level pattern caching (also called advanced compression), TCP optimization, application protocol optimization, Quality of Service (QoS), traffic monitoring, and web server front-end management. For us, the following are most significant:

- Compression. Recent algorithm breakthroughs make network compression a more viable solution. These new compression algorithms can improve bandwidth utilization efficiency, thereby reducing bandwidth congestion.
- Caching. By locally storing copies of frequently sent files, caching eliminates repeat WAN traffic.
- Advanced compression. The by-product of new compression algorithms, advanced compression recognizes and stores larger data patterns on a local disk instead of in memory. Because the whole file is generally recognized as a data pattern, advanced compression can achieve the effect of caching over a high-latency circuit.
- Protocol tuning. By using a pair of "spoofing" gateways at the ends of the network links to intercept and forward the traffic, protocol tuning improves error control and flow control to increase network efficiency and help mitigate latency issues. Protocol tuning uses special, highly efficient protocols that encapsulate the network protocol, such as CIFS.
- Quality of Service. Prioritizing WAN traffic can help ensure important data is transmitted first and limit the impact of network queuing and packet drops.

Products by Market Segment

After our survey, we divided WAN products into three market segments. Each segment contained products that are typically deployed in different locations on the network, applying the underlying technologies where they are most effective.

- Comprehensive WAN optimization. These products tend to provide overall systems that include traffic monitoring, basic compression, TCP optimization, application protocol optimization, and some advanced compression. These products are generally deployed at the edge of the network and target all WAN traffic.
- Wide Area File Services. Wide Area File Services (WAFS) products focus on file transfer acceleration using protocol modification, compression, advanced compression, or a combination of these techniques. These products can be implemented at locations in the network other than the WAN edge.
- Acceleration and load balancing. These products focus on HTTP traffic acceleration and load balancing, and data center oriented acceleration technologies. They are usually part of the network core.

Compiling Best Practices

Industry analysts have a number of recommendations concerning the different WAN optimization technologies:

- Compression and advanced compression can benefit all data traffic.
- High bandwidth, high latency links should use protocol tuning to improve link utilization.
- Sophisticated protocol spoofing or WAFS solutions should be considered if there is large scale use of the standard Microsoft Windows* or UNIX* file or e-mail protocols, such as CIFS, Network File System (NFS), and Messaging Application Program Interface (MAPI), over high latency WAN links.
- Web data that is transmitted repeatedly across the WAN is a candidate for both caching and compression. Server offload techniques can also help.
- QoS can help applications that absolutely require low latency to obtain higher priority service over existing WAN links. Route selection can also transmit these applications over low latency links, while other traffic is sent over lower cost, higher latency links.

Evaluation Unit Requirements

For laboratory testing, we selected two evaluation units that could provide most of the capabilities we were interested in. Based on our network conditions and performance challenges, we decided to investigate TCP optimization, basic and advanced compression, and transfer acceleration using protocol modification. Additionally, we felt it was critical to have visibility into the traffic that traveled across our networks, so we made sure our evaluation units could identify, log, and differentiate transactions (such as ports, protocols, and volume).

Phase Two: Lab Evaluation

We benchmarked our evaluation units in a lab environment. We tested each product's capabilities for compression and visibility of data, QoS, effective latency reduction, and the impact of advanced compression.

Test Network

We included both in-line and off-path products in the Phase Two test network. Figure 3 on the next page illustrates one example of the topology. We positioned one product off-path from our simulated WAN to provide basic and advanced compression and TCP optimization as traffic crossed the WAN. We installed the second product off-path to accelerate file transfers. It provided advanced compression and protocol modification before traffic was routed to the WAN.

The products were easy to integrate into the network and manage. We used a WAN delay simulator to dial in different round trip latencies from 0 to 250 ms, and then we measured performance when the product was both active and inactive. We used data transfers that relied on different protocols to see how the products improved the transaction through our test network.

Laboratory Evaluation Results

Both products enabled impressive performance improvement for various applications. Appendix A lists our measurement data. Here is a summary of our results:

- **Data reduction ratio.** The data reduction ratio varied depending on the type of file. We used over ten different file types, with the data reduction ratio ranging from 0 percent for already compressed files to 75 percent for database files. With the advanced compression feature, after the initial file transfer the reduction ratio was always 99 percent for all file types due to caching.
- **Throughput improvement.** For a simulated WAN circuit of 8 Mbps with 250 ms RTT, the overall throughput improvement for the second pass of file transfer was similar to LAN performance levels.
- **Resending files.** After sending a file, the subsequent resend of a slightly modified version saved over 90 percent of data transferred. This is due to the fact that the data store remembers large data patterns at a bit level.
- **TCP acceleration.** The TCP acceleration was about 4 to 12 times faster, depending on circuit conditions.
- **CIFS acceleration.** The CIFS acceleration was about five times faster for a circuit of 6 Mbps with 250 ms RTT.
- **Off-path deployment models.** When a product is inserted in-line, it can "fail to wire" smoothly in the event of a power failure or manual shutdown. Off-path deployment models include "router on a stick," offering flexibility for use in the network.

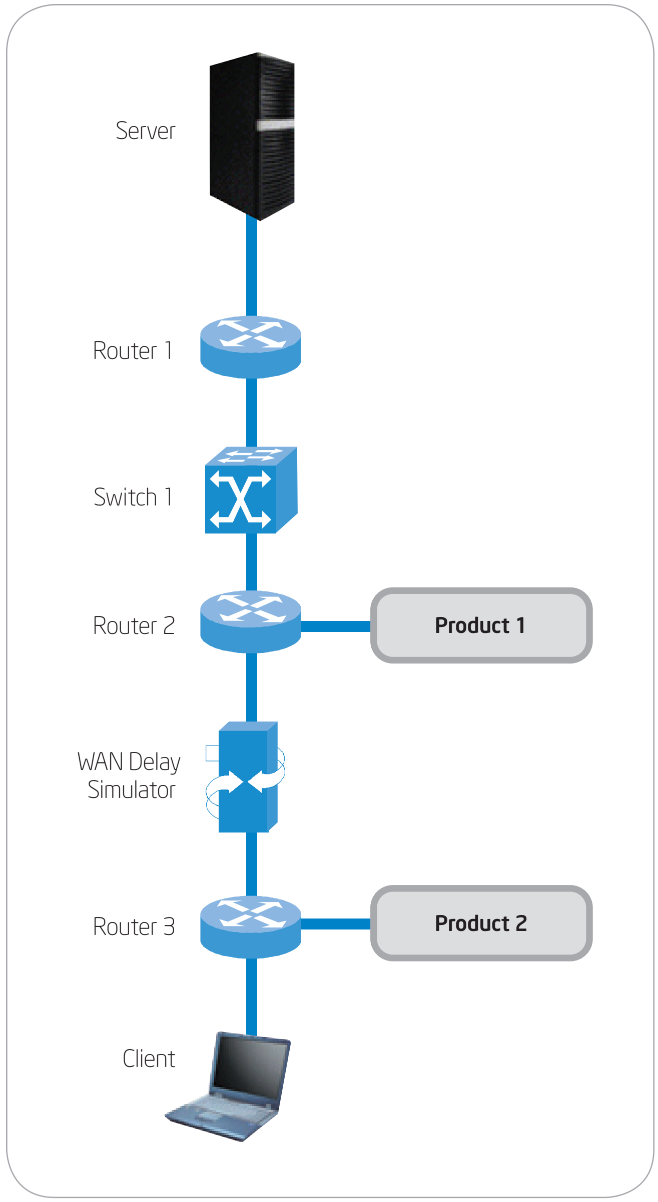

Figure 3. Laboratory evaluation test network.

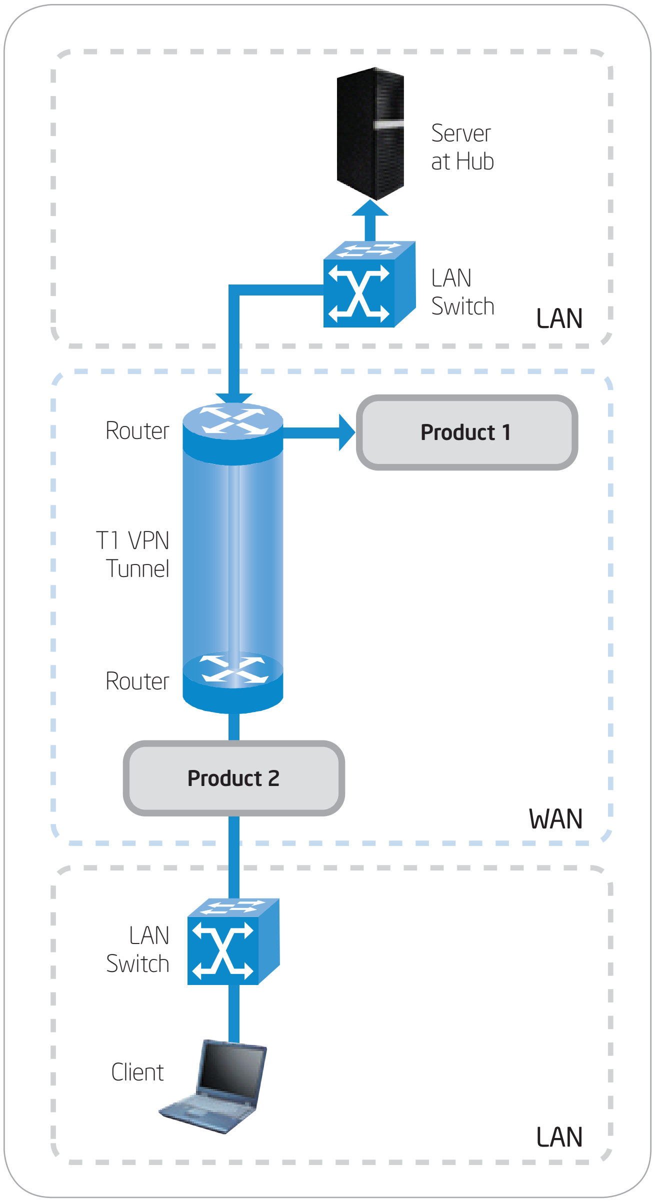

Figure 4. Test site deployment.

With our impressive benchmark results, we felt these products could offer important improvements to our WAN services. Our next step was to test them in a real world production environment.

Phase Three: Production Network Testing

To help determine a site for our production test network, we completed a traffic analysis for some representative locations. We discovered the following:

- Almost every site saw Web and CIFS traffic, both of which use considerable bandwidth.
- Engineering computing sites commonly saw file mirroring traffic.
- Many other well-known and unknown applications consume large amounts of bandwidth.

Table 1 on the next page summarizes the types of traffic we found at the possible test sites. TCP port numbers (in parentheses) gave us visibility into the traffic and helped us classify it prior to testing. We ultimately selected an Intel site with about 30 users and one WAN router.

Test Setup

At the test site, the WAN router connects to a remote Intel hub using a T1 virtual private network (VPN) tunnel. We placed the test appliances at each end of the T1 to optimize the traffic across the link. Figure 4 shows the deployment topology. We placed one product in-line in the WAN circuit. We implemented the other product at the remote hub as an off-path device and configured it so that the hub router routed only traffic to and from the test site to the product.

Test Methodology

We monitored traffic for five weeks and collected data to measure real performance:

- The appliance's monitoring page, to classify traffic
- User-timed transactions
- WAN circuit statistics gathered from network analysis software

Results

The test site products had excellent traffic discovery and monitoring capability, classifying over 75 percent of the traffic and giving us visibility into the traffic that flowed through the network. Figure 5 charts the traffic types. Most traffic consisted of the following:

- Client-to-server mail traffic between the remote test site and the mail server at the hub site
- Internet (Web-proxy)
- Intranet (HTTP)
- CIFS
- Internal resources
- Enterprise resource planning (ERP) application traffic

In addition, we discovered the following performance improvements:

- An average 60 percent data reduction verified by two independent data sources—the product's monitoring page and circuit statistics
- User-invoked large FTP file transfers that were 13 times faster
- HTTP traffic with up to a 14 percent time savings

Acceleration did not increase for more than 96 percent of CIFS traffic. When we investigated, we discovered that the product did not support the type of CIFS transaction used in our production network. The supplier confirmed that the product's next release would support these types of CIFS transactions.

Table 1. Commonly seen WAN traffic.
(TCP port numbers in parentheses)

Office	Business Applications
Intranet, Internet (80, 8080, 443, 911)	Rysnc (873)
CIFS (139, 445)	Business continuity software (10566)
Mail (dynamic port, 25)	NFS (2049)
DNS replication (UDP port 53)	SSH (22)
Intel internal network activities (507, 1406, 1405, 1850)	Shell (514)
Database server (3180, 1433)	FTP (20)
Network-based meeting tool (1503)	
Tool: VNC (5900-5906)	
Tool: Terminal Services (3389)	
X11 (6000-6006)	

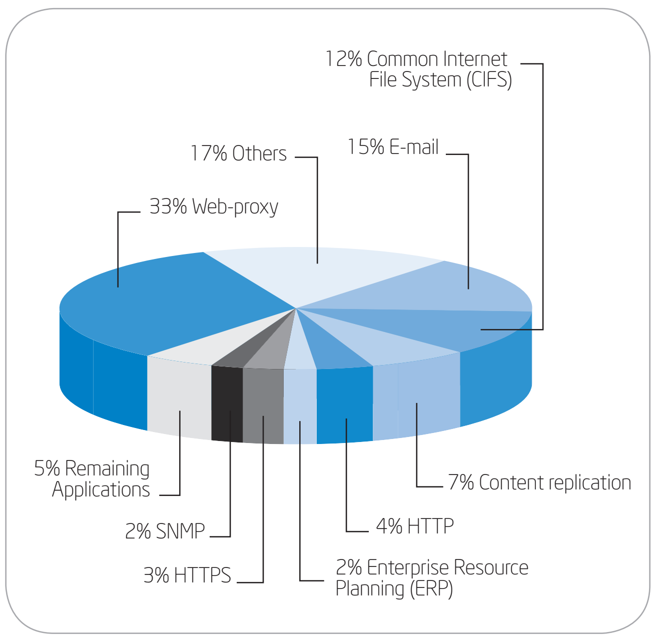

Figure 5. Production test traffic profile. Percent of Traffic To and From WAN by Application

Conclusion

WAN optimization has greatly matured over the last several years and now includes technologies that can be applied with impressive results for high latency WAN circuits. We saw significant benefits in terms of accelerated transfers and smaller file sizes for the major protocols that cross our WAN circuits, including a 99 percent reduction in file sizes for second-pass advanced compression, up to 12 times faster TCP transfers, and five times faster CIFS transfers in our test network. We will continue to investigate WAN optimization for Intel deployments.

Appendix A. Benchmark Results

Tables 2 through 5 list our measured data from laboratory testing.

Table 2. Compression results of FTP file copy at 8 Mbps, 250 ms RTT

Type	Size	1st Pass Reduction Ratio (%)	2nd Pass Reduction Ratio (%)
.ttc	15MB	49	99
.db	9MB	74	99
.exe	70MB	4.5	99
.iso	686MB	2.2	99
.mis	17MB	3.8	99
.ppt	3MB	14.6	99
.zip	46MB	1.6	99
.iso	12MB	1.9	99
.ezqb	34MB	3.1	99
.dat	904MB	72.2	99
Split	1.9MB	42	99
.doc	1.2MB	48	99
.pdf	4.6MB	14..5	99

Table 3. TCP acceleration and sequencing caching at 8 Mbps, 250 ms RTT

Copy Action by FTP		Throughput seconds (kBytes/sec)			
File Type	File Size	LAN	Baseline	Appliance 1st Pass	Appliance 2nd Pass
.exe	70 MB	57.5 (1261)	328 (220)	52.2 (1388)	17 (4256)
.iso	686 MB	88 (7982)	3025 (232)	762 (922)	99 (7077)
.zip	46 MB	8.6 (5539)	213 (223)	84 (565)	9.1 (5219)
.ios	12 MB	1.7 (7064)	50 (240)	12.8 (937)	1.87 (6423)

Table 4. Advanced compression at T1, 60 ms RTT

	File Type	File Size	Data Transferred across WAN (bytes)
Baseline:	.ppt	18 MB	18M
Pass 1:	.ppt	18 MB	16M
Pass 2:	.ppt	18 MB	294K
Pass 3:	.ppt	Slightly modified 18 MB	335K

Table 5. CIFS acceleration at 6 Mbps, 250 ms RTT

Drag & Drop Activity:		Transfer Time (secs)		Acceleration
File Type	File Size	Baseline	Acceleration	(x faster)
.db	9 MB	87	20	4.35x
.bin	12 MB	240	20	12.00x
.ppt	18 MB	200	50	4.00x
.msi	17 MB	163	51	3.20x

Authors

Tim Verrall is a senior staff network engineer with Intel Information Technology.
David Wescott is a senior network specialist with Intel Information Technology.
Lilin Xie is a senior network engineer with Intel Information Technology.

Acronyms

CIFS	Common Internet File System
ERP	enterprise resource planning
FTP	File Transfer Protocol
MAPI	Messaging Application Program Interface
Mbps	megabits per second
ms	millisecond
NFS	Network File System
PoC	proof of concept
QoS	Quality of Service
RTT	round trip time
TCP	Transmission Control Protocol
VPN	virtual private network
WAFS	Wide Area File Services
WAN	wide area network

www.intel.com/IT

Printed in USA ♻ Please Recycle
0506/ARM/RDA/PDF Order Number: 313396-001US