COMMVAULT

Enabling high-speed WAN backups with PORTrockIT

BRIDGEWORKS

EXECUTIVE SUMMARY

Commvault offers one of the most advanced and full-featured data protection solutions on the market, with built-in functionalities such as data compression and source-side deduplication that substantially reduce the time, cost and bandwidth requirements for timely backup and restore operations.

For backups that take place on a local area network (LAN), Commvault's technology is generally all an organisation needs to ensure fast and reliable service. However, when data needs to be moved across a wide area network (WAN), additional bottlenecks can emerge.

As levels of network latency and packet loss increase, the overall speed of data transfer can be affected, which prevents users from getting the full benefit of the Commvault technology.

This paper shows how Bridgeworks' PORTrockIT solution can overcome these two obstacles. By combatting the effects of network latency and packet loss, PORTrockIT can accelerate data transfer rates by a factor of more than 640.

As a result, the combination of PORTrockIT and Commvault can transform backup performance – enabling Commvault to deliver the same powerful advantages across the WAN as it does across the LAN.

*"By combatting the effects of network latency and packet loss, PORTrockIT can accelerate data transfer rates by **a factor of more than 640**."*

WHY PERFORMANCE MATTERS

Effective data protection depends on the ability to make copies of business data in a timely and efficient way. Faced with a wide range of possible disaster scenarios – from the accidental deletion of a file, through the failure of a server, to the complete loss of a whole data centre – organisations seek to optimise their recovery time and recovery point objectives in order to maximise availability and minimise data loss.

The desire to protect larger quantities of data through more frequent backups has led vendors such as Commvault to rethink the way data protection works. As data volumes increase, it becomes impractical to perform full backups regularly – so Commvault has developed compression technologies to reduce the total amount of data that needs to be copied, and deduplication technologies to only copy data that has been created or updated since the previous backup.

These techniques have a significant impact on the speed and cost of backup and restore processes – minimising the total amount of backup data that needs to be stored, and reducing the amount of network bandwidth required to complete the backups within a viable time-window.

*"WAN backups often run into performance issues that **cannot be solved** by optimising bandwidth."*

However, when backups need to take place across a WAN, organisations often run into performance issues that cannot be solved by optimising bandwidth. Even when moving relatively small amounts of compressed, deduplicated data over a very high bandwidth connection, transfer rates can still be extremely slow – hindering the organisation's ability to protect its data.

To understand why this is, we need to take a closer look at how WAN transfer protocols work – and in particular, to understand how they are affected by latency and packet loss.

THE PROBLEMS: LATENCY AND PACKET LOSS

In general, there are two main issues that cause the majority of performance problems when moving data across a WAN.

The first is latency – the time delay between a system sending a packet across the WAN, and the target system receiving that packet. The main causes of latency are:

- The physical distance that the packet has to travel
- The time taken to receive, queue and process packets at either end of the connection
- The time taken to receive, queue and process packets at any intermediate gateways.

The further the data has to travel and the more gateways it has to pass through, the greater the latency is likely to be.

For data transfers that use the TCP/IP protocol, high latency can cripple transfer rates, even over a theoretically high-bandwidth WAN infrastructure. TCP/IP works by sending a group of packets, waiting for an acknowledgement that the packets have been received, and then sending the next group. If the latency of the connection is high, then the sender spends most of its time waiting for acknowledgements, rather than actually sending data. During these periods, the network is effectively idle, with no new data being transferred.

The second issue is packet loss, which occurs when a packet sent from a system on one side of the WAN never arrives at the system that is intended to receive it, or when the acknowledgement from the recipient goes astray before it reaches the sender.

When a packet gets lost, TCP/IP automatically reduces the number of packets it sends in the next group, to compensate for the unreliability of the connection. As a result, network utilisation is greatly reduced, because the sender is sending fewer packets in the same amount of time.

Organisations often try to solve TCP/IP performance issues by investing in more expensive network infrastructure that offers a larger maximum bandwidth. However, this does not fix the problem. As we have seen, latency and packet loss prevent TCP/IP connections from fully utilising the available bandwidth – so any extra investment in bandwidth will simply be wasted unless the latency and packet loss issues can be addressed.

THE SOLUTION: PORTROCKIT

PORTrockIT offers a solution to network latency issues. Instead of sending a group of packets down a single physical connection and waiting for a response, the solution creates a number of parallel virtual connections that send a constant stream of data across the physical connection.

As soon as a virtual connection has sent its packets and starts waiting for an acknowledgement from the recipient, PORTrockIT immediately opens another virtual connection and sends the next set of packets. Further connections are opened until the first connection receives its acknowledgement; this first connection is then used to send another set of packets, and the whole process repeats.

This parallelisation practically eliminates the effects of latency by ensuring that the physical connection is constantly transferring new packets from the sender to the recipient: there is no longer any idle time, and the network's bandwidth can be fully utilised.

The solution also significantly reduces the impact of packet loss. If one of the virtual connections loses a packet, TCP/IP will only reduce the number of packets in the next group sent by that specific virtual connection. All the other virtual connections continue to operate at full speed.

Moreover, PORTrockIT is capable of optimising the flow of data across the WAN in real time, even if network conditions change. The solution incorporates a number of artificial intelligence engines that continuously manage, control and configure multiple aspects of PORTrockIT – enabling the appliance to operate optimally at all times, without any need for input from a network administrator.

In practical terms, PORTrockIT is installed as a pair of appliances, deployed at either end of the WAN. A source server, running a Commvault Virtual Server Agent, simply passes data to the PORTrockIT appliance on the near side of the WAN, which manages the virtual connections to the second PORTrockIT appliance on the far side of the WAN. Once the second PORTrockIT appliance begins receiving packets, it routes them seamlessly to the Commvault MediaAgent, which manages the target storage media.

The effect is simply much faster network transfer performance, without any need to make any changes to the rest of the network architecture.

TURNING THEORY INTO PRACTICE

To demonstrate the results that PORTrockIT can deliver for organisations that want to accelerate WAN data transfer with Commvault, Bridgeworks conducted a set of performance tests, using different Commvault features to simulate four different scenarios:

1. Full backup of VMware virtual machine using Commvault deduplication
2. Full backup of a VMware virtual machine using Commvault compression
3. Auxiliary Copy of a VMware virtual machine
4. DASH Copy of a VMware virtual machine

The test infrastructures mimicked a real-world Commvault architecture, using a WANulator to simulate different levels of WAN latency between source and target servers.

The first set of tests was performed on an unaccelerated architecture (see Figures 1 and 3), where the source and target servers were connected directly to the WANulator. The same tests were then repeated on an architecture that was accelerated by introducing two PORTrockIT appliances, placed on either side of the WANulator, between the source and the target (see Figures 2 and 4).

TEST EQUIPMENT

SOFTWARE:

- Commvault version 11.0

HARDWARE:

- VMware ESX server:
  Windows Server 2012 R2 host,
  IBM System x3250, 8 GB RAM,
  Intel Xeon E31230 3.2 GHz
- Commvault server #1 (used in scenarios 2, 3 and 4):
  Windows Server 2012 R2 host,
  Sun Fire x2250, 16 GB RAM,
  2 x Intel Xeon X5472 3.0 GHz
- Commvault server #2 (used in scenarios 1, 3 and 4):
  Windows Server 2012 R2 host,
  DELL R710, 8 GB RAM,
  2 x Intel Xeon E5506 2.13 GHz
- 2 x Bridgeworks PORTrockIT nodes
- WANulator

Figure 1: Unaccelerated architecture used in scenarios 1 and 2 (full backups of a VMware virtual machine using Commvault dedupe and compression features)

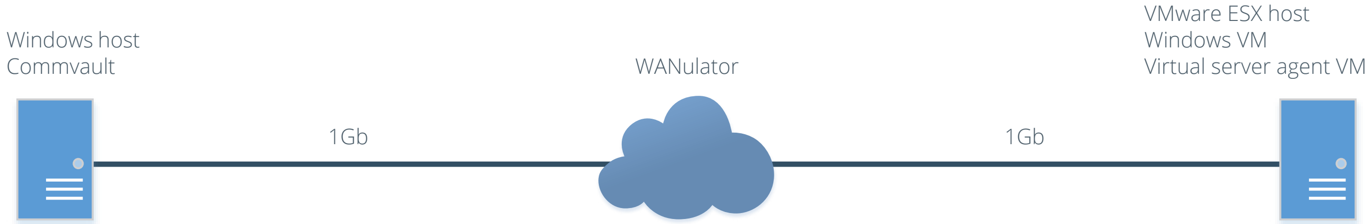

Figure 2: Accelerated architecture with PORTrockIT used in scenarios 1 and 2 (full backups of VMware virtual machines using Commvault dedupe and compression features)

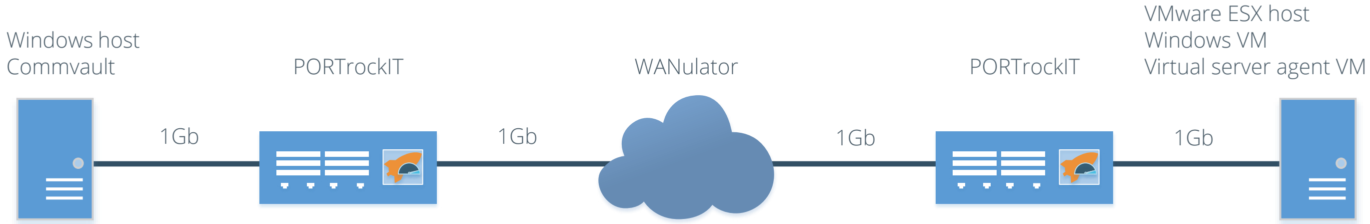

Figure 3: Unaccelerated architecture used in scenarios 3 and 4 (Auxiliary Copy and DASH Copy of a VMware virtual machine)

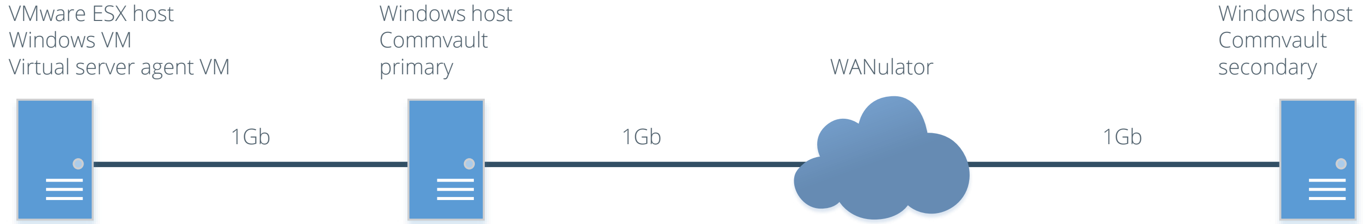

Figure 4: Accelerated architecture with PORTrockIT used in scenarios 3 and 4 (Auxiliary Copy and DASH Copy of a VMware virtual machine)

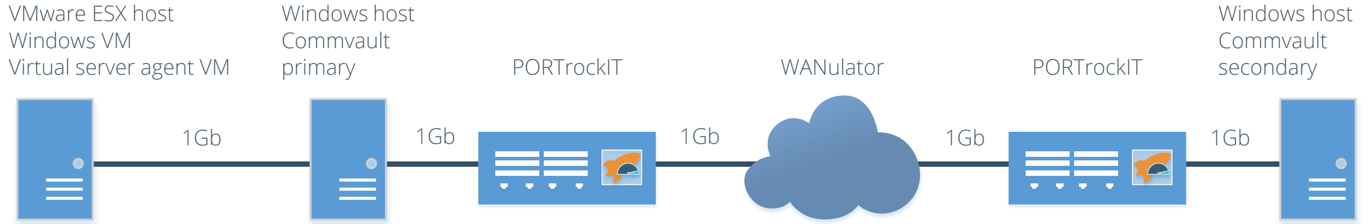

WHAT THE DATA TELLS US

SCENARIO 1: FULL BACKUP OF A VMWARE VIRTUAL MACHINE USING THE COMMVAULT DEDUPE FEATURE

In the first scenario, the Commvault software was used to perform a full backup of a VMware virtual machine.

The source server was an IBM System x3250 server running VMware ESX. The VMware environment hosted both the Windows virtual machine to be backed up, and a second virtual machine running the Commvault Virtual Server Agent. The target server was a Dell R710 running the Commvault MediaAgent.

Multiple tests were run with different values for latency and packet loss. Each test was run first on the unaccelerated architecture (see Figure 1), and then again on the accelerated architecture with PORTrockIT (see Figure 2). After each test, the Commvault deduplication database was cleared, to ensure that a full backup would be performed in the following test.

LATENCY

The first test simulated a scenario with no packet loss, at latencies ranging from 0 ms to 360 ms round trip time (RTT).

From Figure 5, we can see that even small amounts of latency had a negative impact on the transfer rate

Figure 5: Scenario 1 – accelerated and unaccelerated performance for a full backup of a virtual machine using Commvault dedupe with 0% packet loss at various latencies

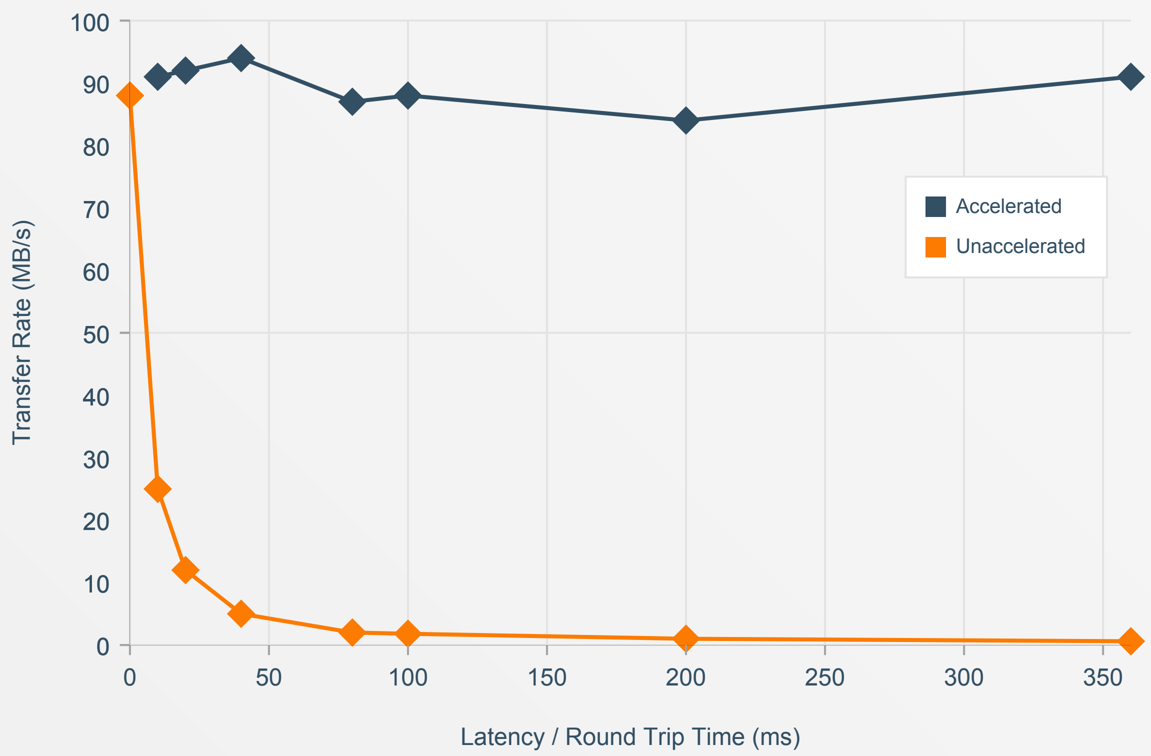

for the backup on the unaccelerated architecture. Moving from 0 ms to 10 ms of latency reduced the transfer rate from 88 MB/s to 25 MB/s – reducing performance by nearly two thirds.

As latency increased further, the transfer rate was reduced to almost nothing. With a round trip time of 360 ms, the transfer rate was just 0.6 MB/s.

By contrast, the accelerated architecture delivered stable performance of between 84 MB/s and 94 MB/s at all latency values. At 10 ms of latency, the accelerated architecture was 3.6 times faster than the unaccelerated architecture, and at 360 ms of latency, it was more than 150 times faster.

COMBINED LATENCY AND PACKET LOSS

Next, a series of tests was conducted to assess the effect of combining packet loss with latency. Three different packet loss scenarios (0.1%, 0.5% and 1%) were tested at various levels of latency. In practice, all real-world WANs are subject to at least some degree of both packet loss and latency, so these scenarios provide an indication of how Commvault might perform with and without PORTrockIT acceleration across a WAN.

Figures 6, 7 and 8 all show that packet loss exacerbates the problems created by latency. As we saw earlier, at 10 ms of latency with no packet loss, the transfer

Figure 6: Scenario 1 – accelerated and unaccelerated performance for a full backup of a virtual machine using Commvault dedupe with 0.1% packet loss at various latencies

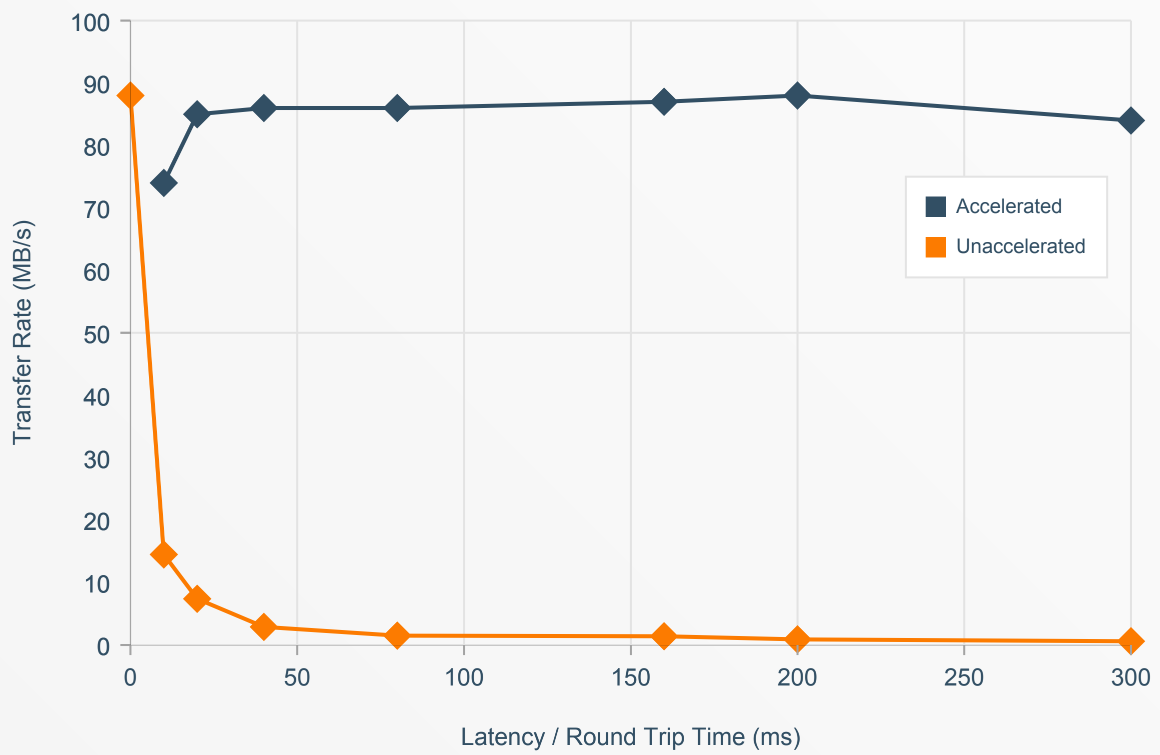

Figure 7: Scenario 1 – accelerated and unaccelerated performance for a full backup of a virtual machine using Commvault dedupe with 0.5% packet loss at various latencies

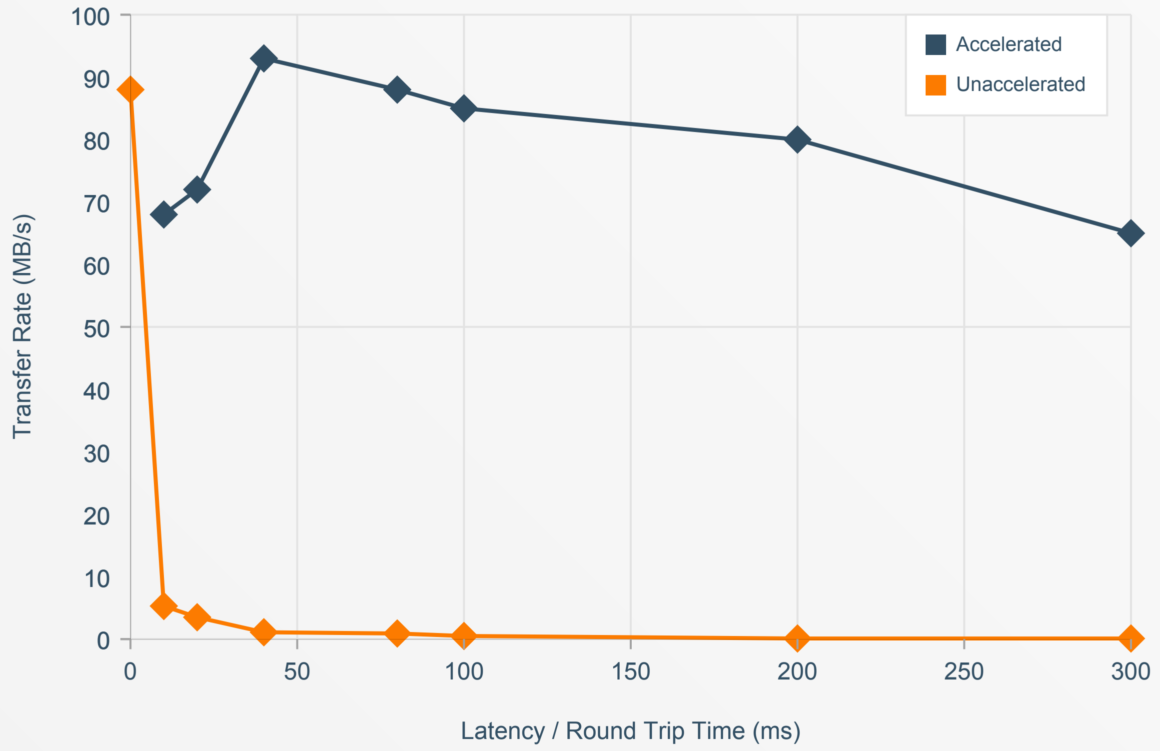

Figure 8: Scenario 1 – accelerated and unaccelerated performance for a full backup of a virtual machine using Commvault dedupe with 1% packet loss at various latencies

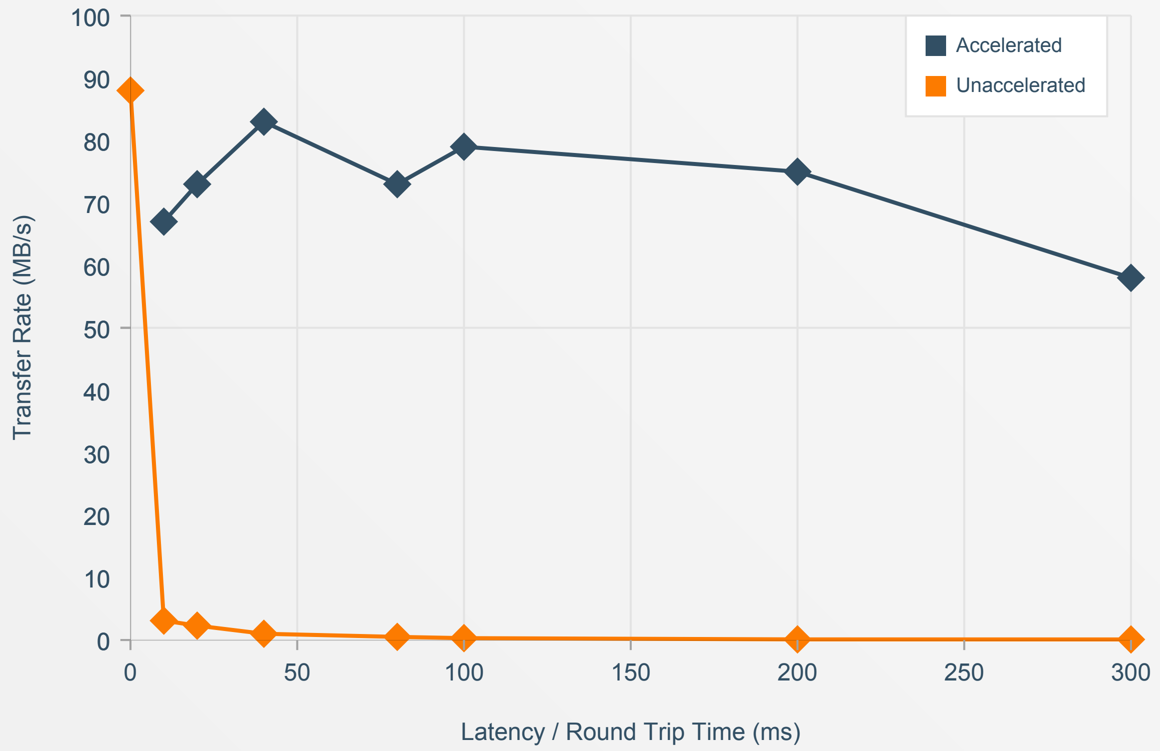

rate of the unaccelerated architecture was 25 MB/s. This rate falls to 14.5 MB/s at 0.1% packet loss; 5.3 MB/s at 0.5% packet loss; and 3.1 MB/s at 1% packet loss.

With the accelerated architecture, although the combination of latency and packet loss does have some impact on performance, the transfer rates remain high and relatively stable. At 0.1% packet loss, there is almost no degradation in transfer rates as latency increases: speeds range between 74 and 84 MB/s at all times.

At 0.5% and 1% packet loss, performance does gradually degrade as latency increases. However, even in the most difficult conditions (1% packet loss at 360 ms latency), the transfer rate of 58 MB/s is still respectable – and more than 640 times faster than the unaccelerated architecture.

SCENARIO 2: FULL BACKUP OF A VMWARE VIRTUAL MACHINE USING THE COMMVAULT COMPRESSION FEATURE

The second test scenario was similar to the first, but tested the performance of PORTrockIT versus an unaccelerated architecture while using Commvault's data compression features on a full backup of a virtual machine.

The hardware setup was identical to the infrastructure used in the first scenario, with the exception of the Commvault server. In this case, a slightly more powerful Sun Fire x2250 server was used in place of the Dell R710. This was used to ensure the additional computing requirements of the Commvault compression algorithms did not impact the overall performance.

The virtual machine used in the test was a 35GB Windows Server 2012 environment, which Commvault was able to compress by 67%.

LATENCY

Figure 9 shows the impact of latency on the transfer rate of the compressed virtual machine backup.

In this case, the unaccelerated architecture performed slightly better than it did in the dedupe backup test: at 10 ms of latency, the transfer rate was 63 MB/s, compared to 25 MB/s in the dedupe scenario.

Nevertheless, the performance of the unaccelerated architecture still tails off dramatically as latency increases, while the accelerated architecture maintains a stable and high transfer rate of more than 115 MB/s until latency levels rise above 100 ms.

At this point, performance begins to degrade, but even at 360 ms of latency, the transfer rate is 91 MB/s – 45.5 times faster than the unaccelerated architecture.

COMBINED LATENCY AND PACKET LOSS

Again, additional tests were then run at various latencies with 0.1%, 0.5% and 1% packet loss. The results can be seen in Figures 10, 11 and 12.

Broadly speaking, the same pattern can be observed in these tests as in the dedupe scenario: on the unaccelerated architecture, the combination of latency and packet loss has a very significant impact on transfer rate, while the accelerated architecture performs significantly better in all cases.

Note: One curious feature is that the accelerated architecture seems to perform worse at lower levels of latency (below 40 ms) than it does at mid-range levels (40 ms – 100 ms). Bridgeworks believes that this indicates a slight problem with PORTrockIT's artificial intelligence engines, which is now being addressed and will be resolved in future releases of the software.

Despite this issue, the graphs still indicate that PORTrockIT delivers a significant boost to Commvault performance in all cases.

Figure 9: Scenario 2 – accelerated and unaccelerated performance for a full backup of a virtual machine using Commvault compression with 0% packet loss at various latencies

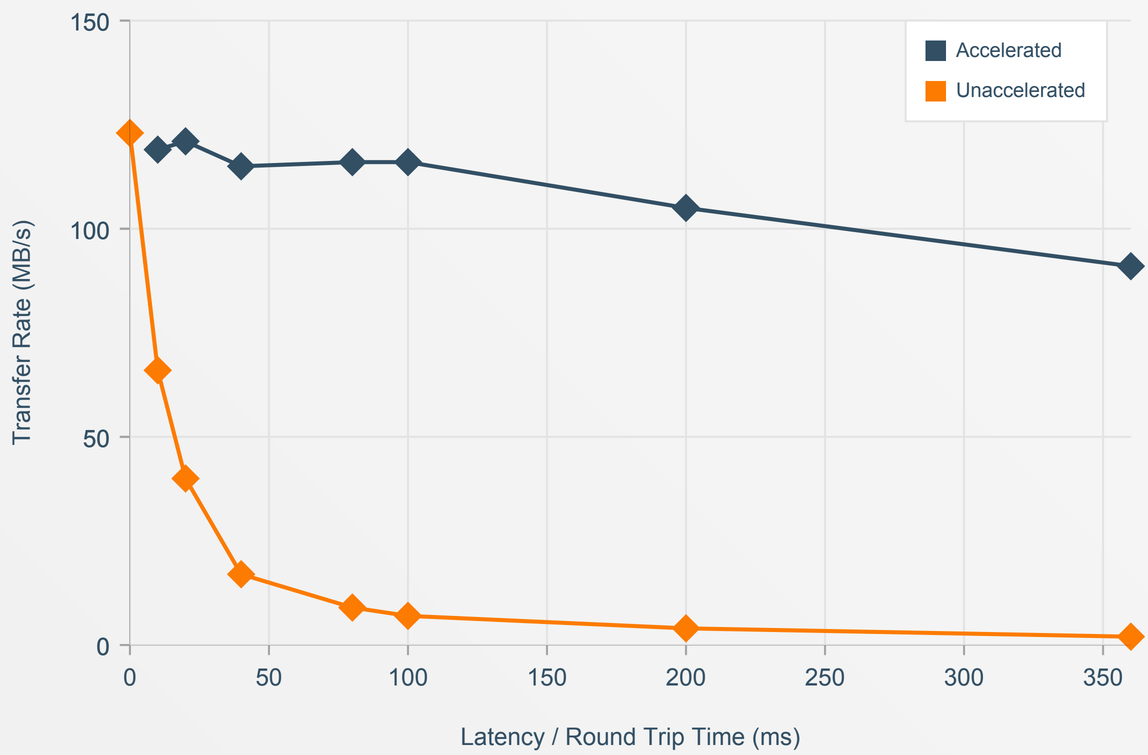

Figure 10: Scenario 2 – accelerated and unaccelerated performance for a full backup of a virtual machine using Commvault compression with 0.1% packet loss at various latencies

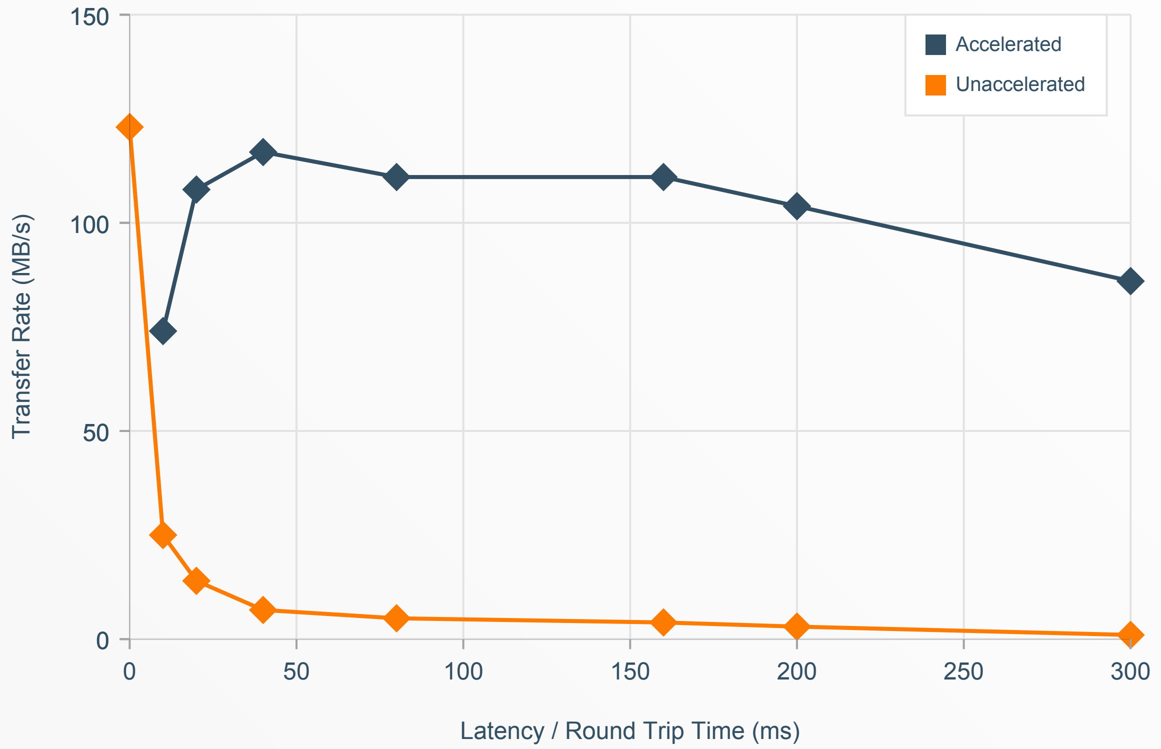

Figure 11: Scenario 2 – accelerated and unaccelerated performance for a full backup of a virtual machine using Commvault compression with 0.5% packet loss at various latencies

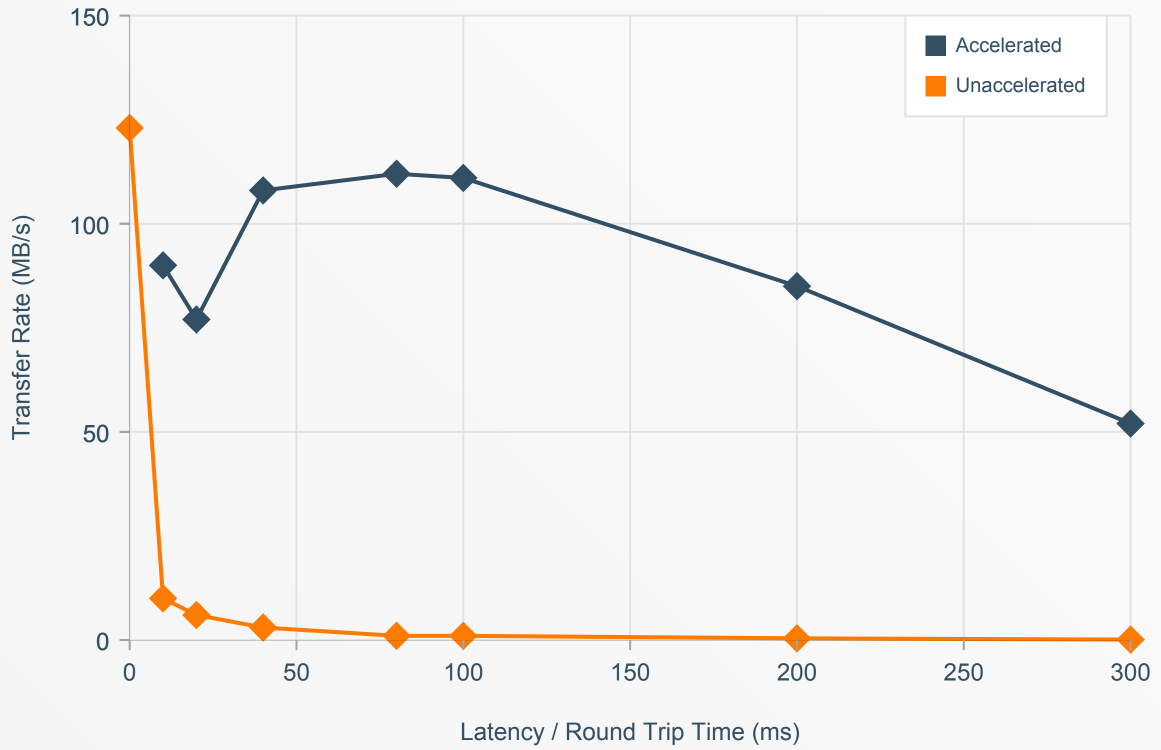

Figure 12: Scenario 2 – accelerated and unaccelerated performance for a full backup of a virtual machine using Commvault compression with 1% packet loss at various latencies

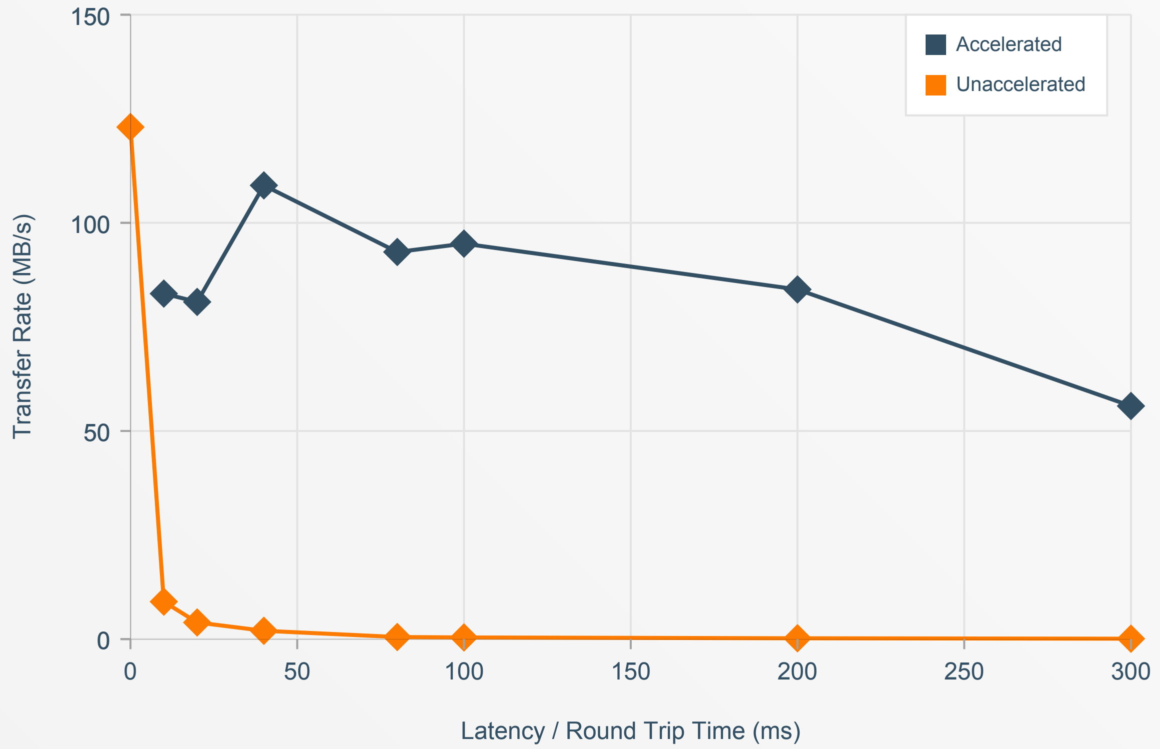

SCENARIO 3: AUXILIARY COPY OF A VMWARE VIRTUAL MACHINE

The third scenario involved testing Commvault's Auxiliary Copy feature, which is used to make an additional copy of an existing backup. The unaccelerated and accelerated architectures are shown in Figures 3 and 4.

Before the test began, a virtual machine was backed up from the VMware ESX server (the IBM System x3250) to the primary Commvault server (the Dell R710). The test then involved using Auxiliary Copy to replicate the backed up data to the secondary Commvault server (the Sun Fire x2250), across the WAN.

Once again, multiple tests were run with different values for latency and packet loss. Each test was run first on the unaccelerated architecture, and then again on the accelerated architecture with PORTrockIT.

LATENCY AND PACKET LOSS

Figures 13 to 16 show the impact of latency on transfer rates at different levels of packet loss (0%, 0.1%, 0.5% and 1%). To make the best use of the limited time available for testing, a smaller range of latency values were tested: 10 ms, 100 ms, and 200 ms.

Figure 13: Scenario 3 – accelerated and unaccelerated performance for an Auxiliary Copy of a virtual machine with 0% packet loss at various latencies

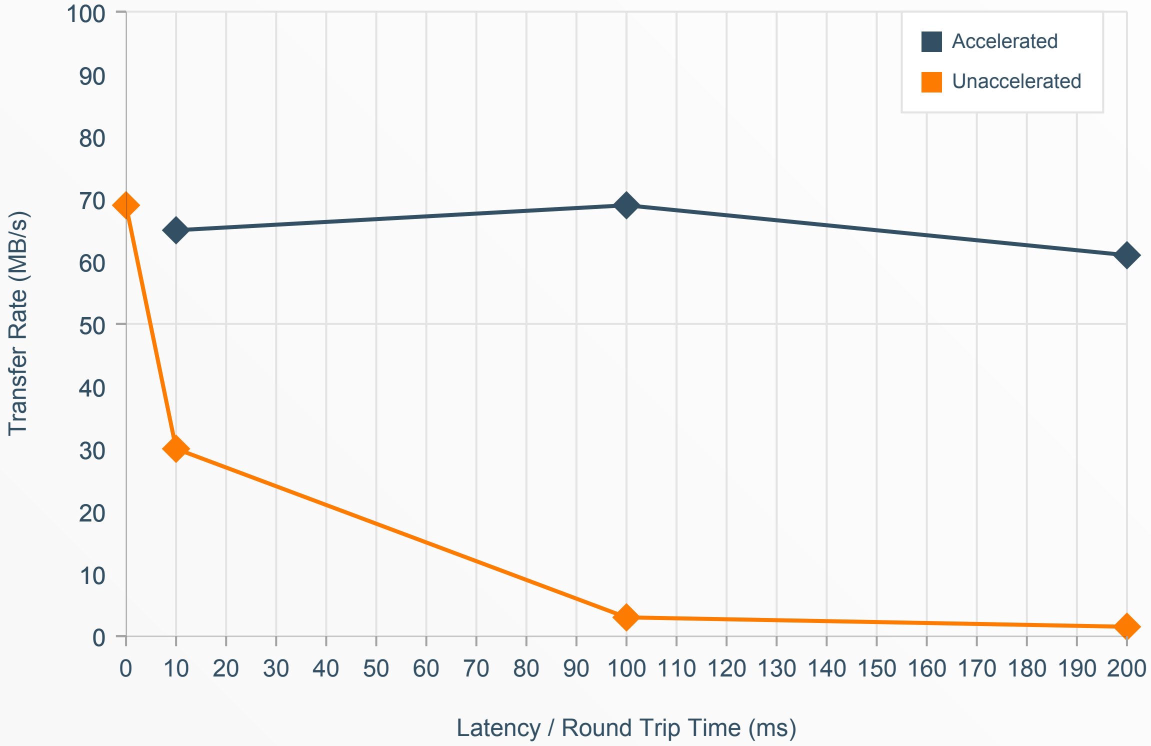

Figure 14: Scenario 3 – accelerated and unaccelerated performance for an Auxiliary Copy of a virtual machine with 0.1% packet loss at various latencies

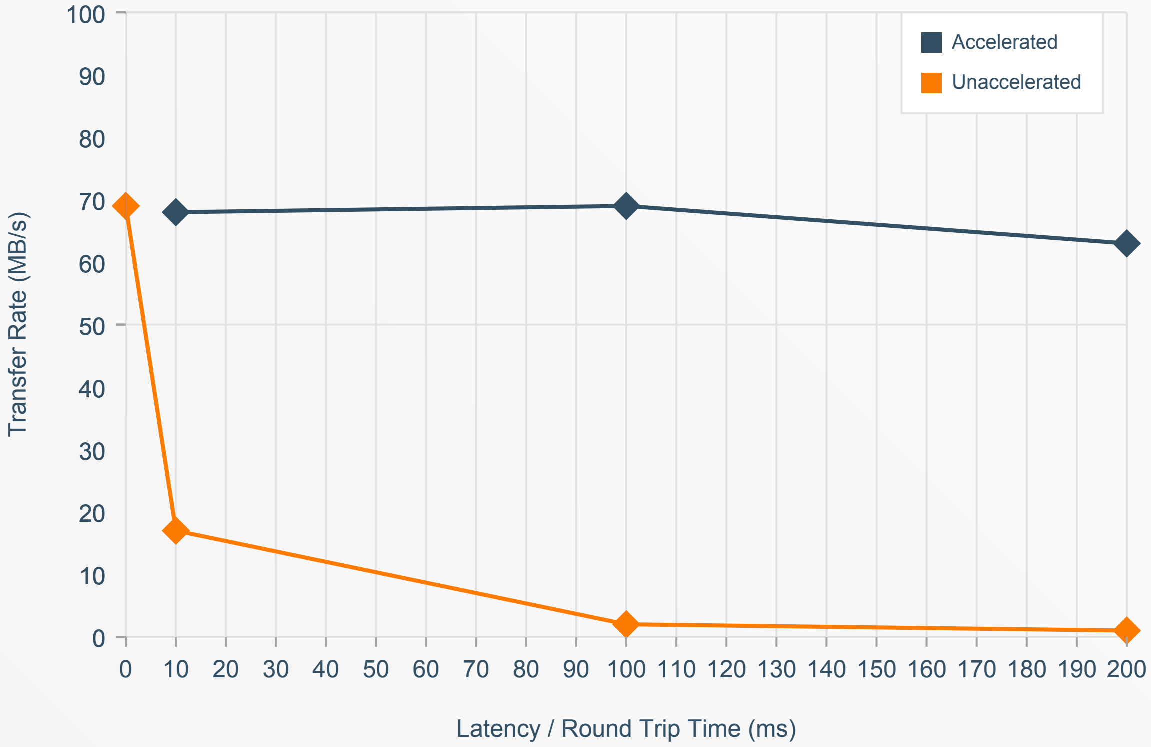

Figure 15: Scenario 3 – accelerated and unaccelerated performance for an Auxiliary Copy of a virtual machine with 0.5% packet loss at various latencies

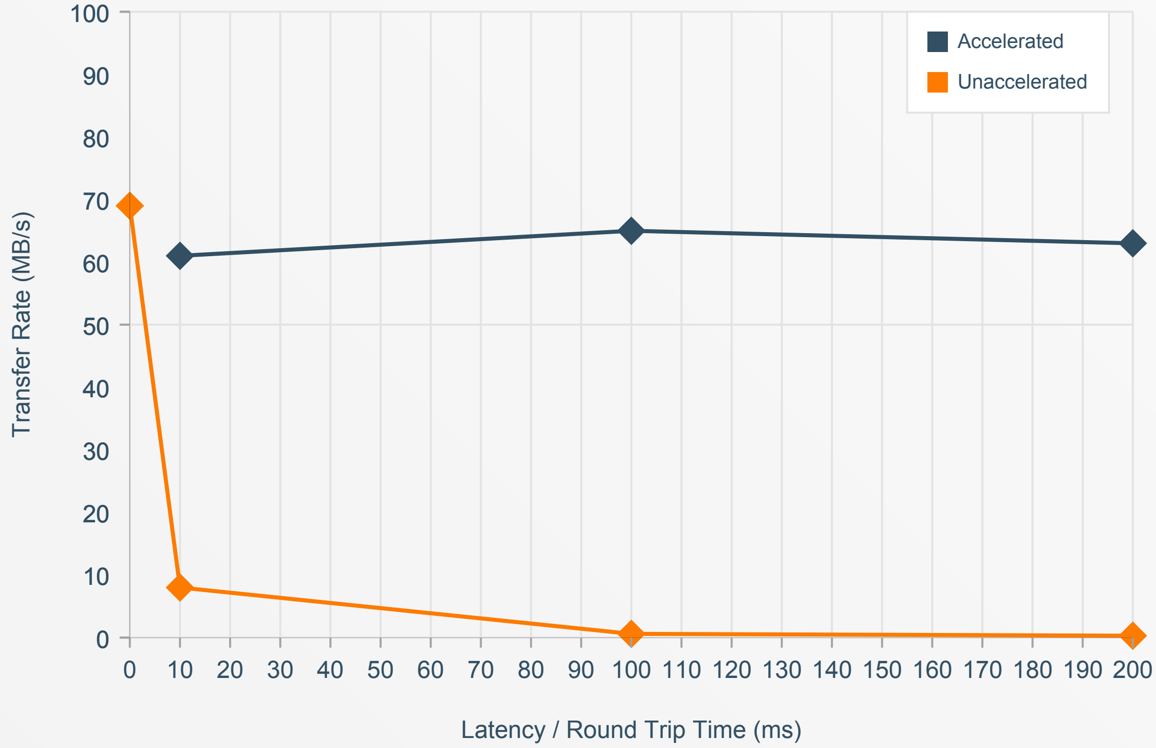

Figure 16: Scenario 3 – accelerated and unaccelerated performance for an Auxiliary Copy of a virtual machine with 1% packet loss at various latencies

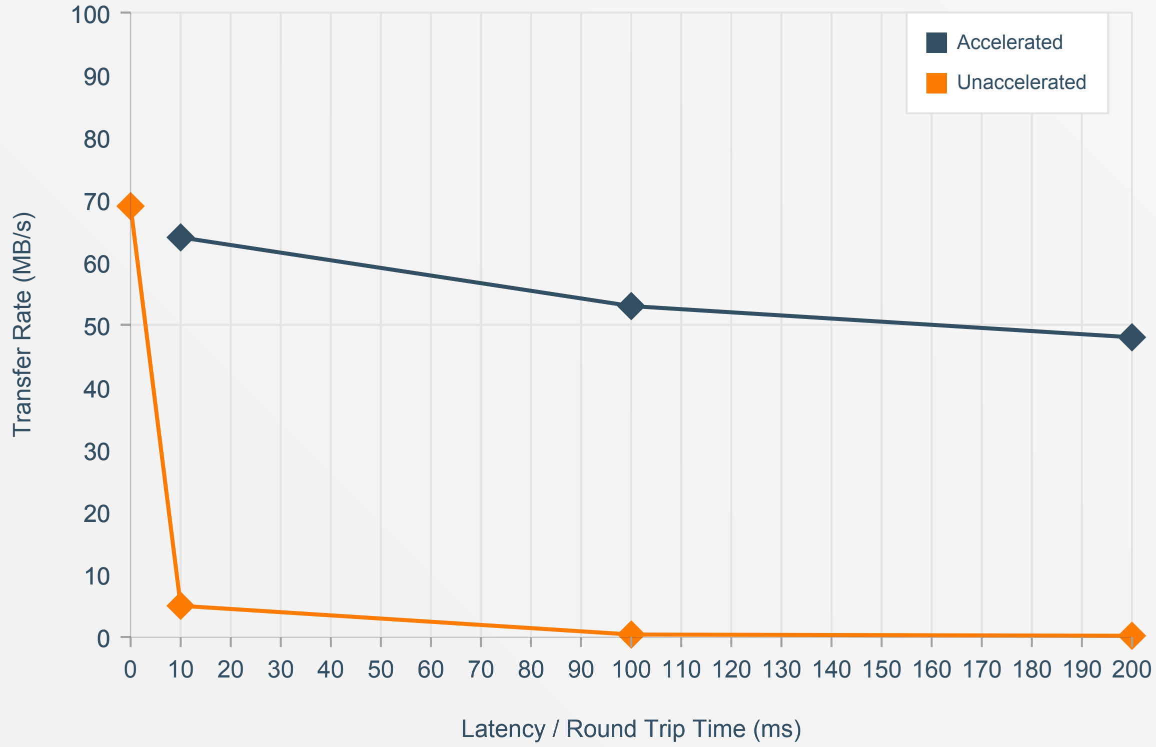

In all four graphs, we can see that once again, the unaccelerated architecture struggles as latency increases, and that higher rates of packet loss exacerbate the performance issues.

We can also see that the accelerated architecture deals much more effectively with both latency and packet loss, providing stable transfer rates in all conditions, with only a slight degradation at the highest packet loss rates. Even at 200 ms of latency and 1% packet loss, the accelerated architecture maintains a transfer rate of 48 MB/s, compared to 0.2 MB/s for the unaccelerated architecture – making it 240 times faster.

SCENARIO 4: DASH COPY OF A VMWARE VIRTUAL MACHINE

The final test scenario looked at the performance of Commvault's Auxiliary Copy feature with DASH Copy enabled. DASH Copy is an advanced client-side deduplication feature offered by Commvault, which aims to make more efficient use of network bandwidth and storage resources by only transmitting unique blocks of data across the network – reducing the data volume and duration of Auxiliary Copy jobs by up to 90%.

This scenario used exactly the same setup as the Auxiliary Copy scenario (see Figures 3 and 4). The only difference was that the DASH Copy option was enabled in the Commvault console. After each test, the Commvault deduplication database was cleared, to ensure that a full backup would be performed in the following test.

"At 200 ms of latency and 1% packet loss, PORTrockIT is ***240 times faster****."*

LATENCY AND PACKET LOSS

Figures 17 to 20 show the results of DASH Copy tests conducted at various latencies with packet loss rates of 0%, 0.1%, 0.5% and 1% respectively.

As in the previous scenarios, the unaccelerated architecture fares poorly: there is significant performance degradation as latency increases, and higher levels of packet loss exacerbate the problem.

Unlike the previous scenarios, latency and packet loss also seem to have a similar – though less severe – effect on the accelerated architecture. This is because the DASH Copy feature is running its deduplication process on the client side in real time during the test, and it is suspected that it cannot supply enough deduplicated data to the PORTrockIT appliance. As a result, PORTrockIT is starved of data to send, and cannot keep the connection fully saturated with data throughout the duration of the test.

Figure 17: Scenario 4 – accelerated and unaccelerated performance for a DASH Copy of a virtual machine with 0% packet loss at various latencies

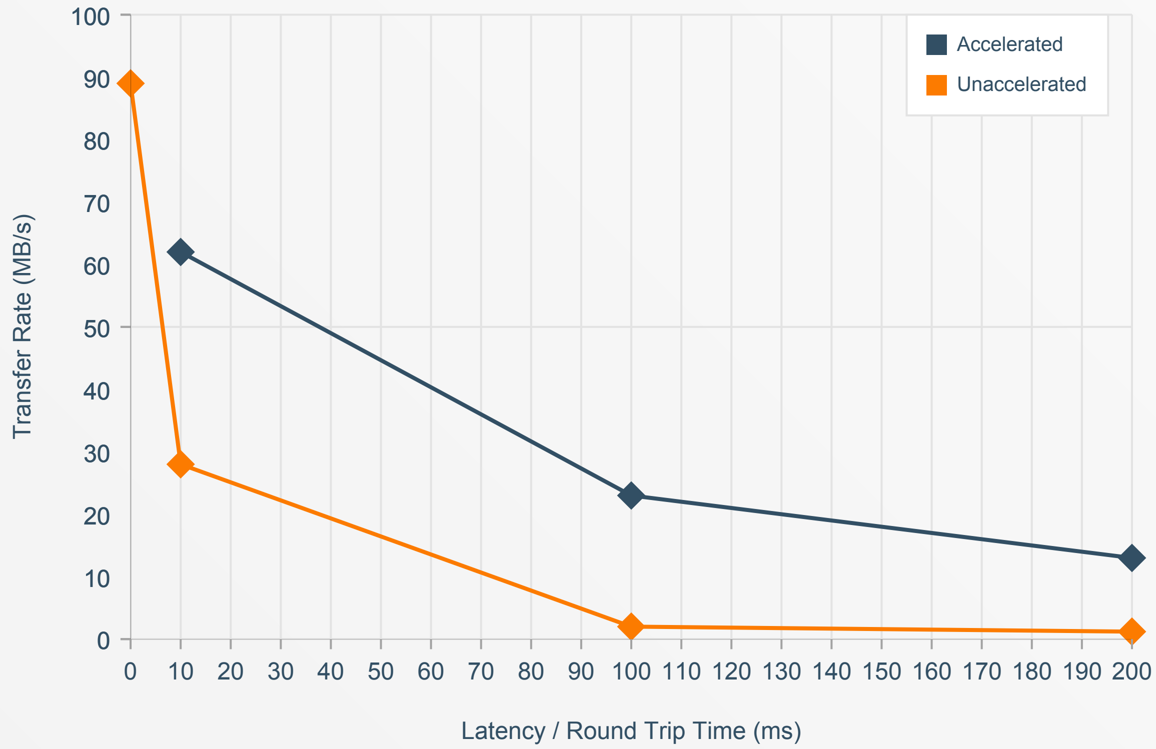

Figure 18: Scenario 4 – accelerated and unaccelerated performance for a DASH Copy of a virtual machine with 0.1% packet loss at various latencies

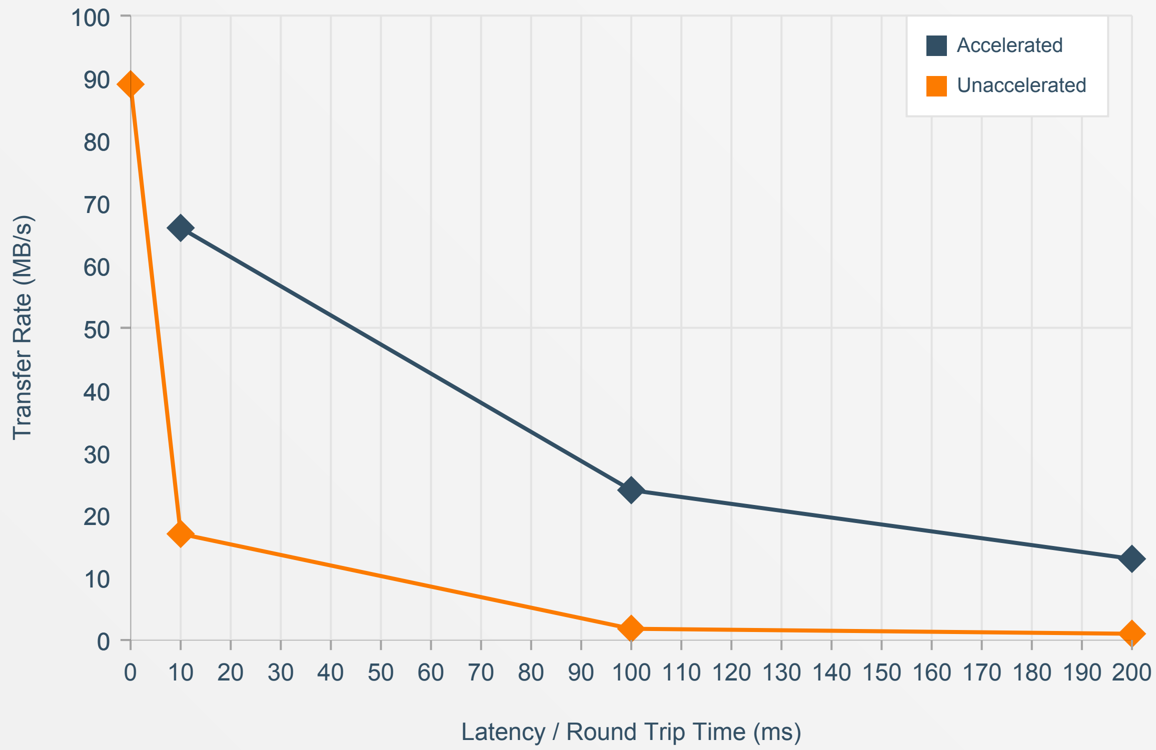

Figure 19: Scenario 4 – accelerated and unaccelerated performance for a DASH Copy of a virtual machine with 0.5% packet loss at various latencies

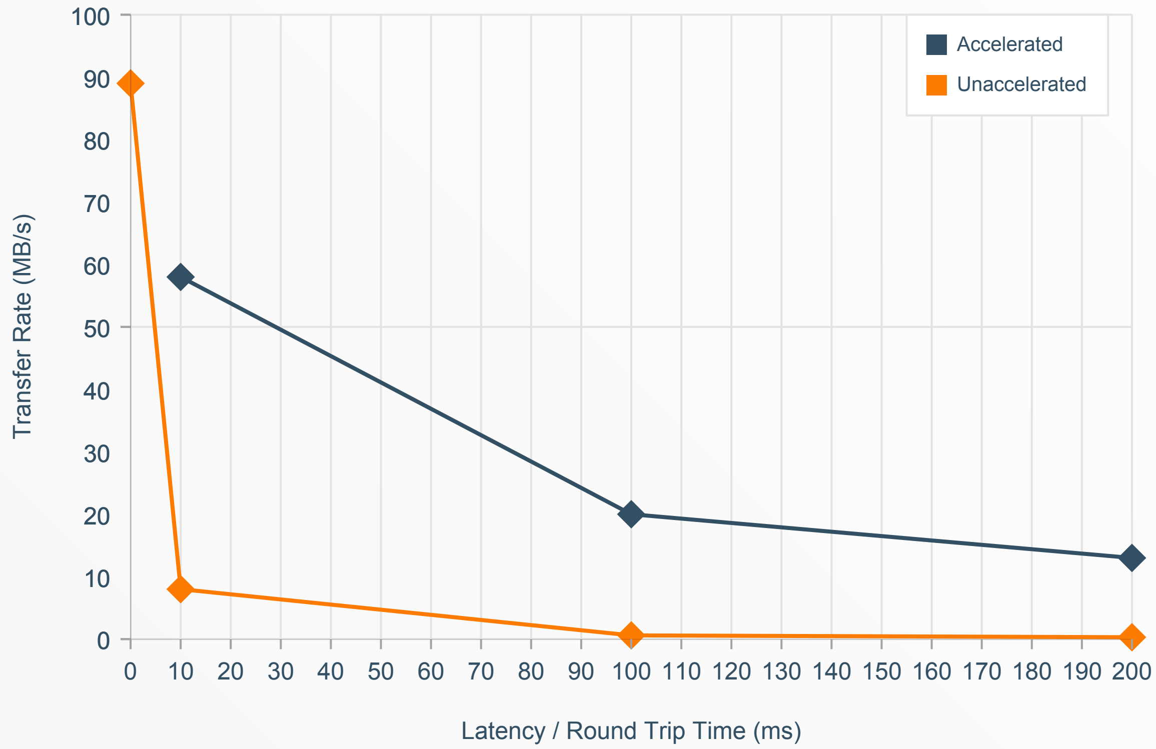

Figure 20: Scenario 4 – accelerated and unaccelerated performance for a DASH Copy of a virtual machine with 1% packet loss at various latencies

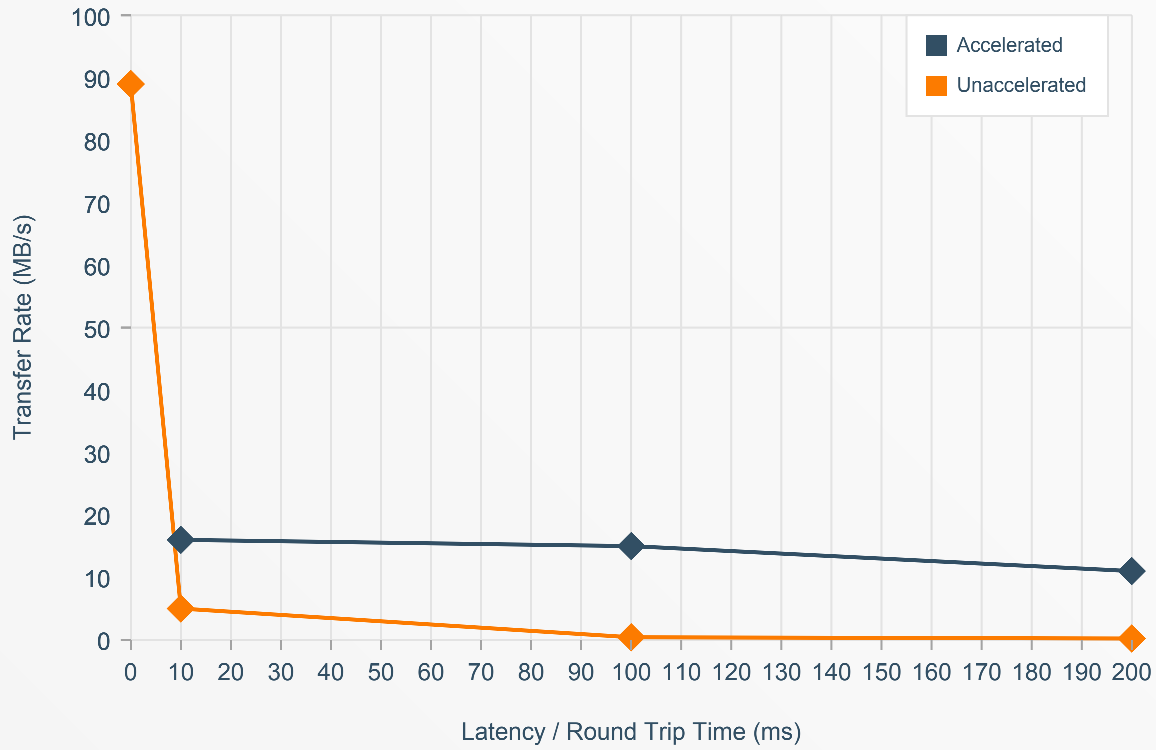

Nevertheless, the accelerated architecture still delivers a significant performance increase in all latency and packet loss conditions. In all the scenarios tested, transfer rates with PORTrockIT were at least twice as fast as with the unaccelerated architecture; and in the most difficult conditions, PORTrockIT was 55 times as fast.

> *"In all scenarios tested, PORTrockIT was* ***at least twice as fast*** *as an unaccelerated architecture; and in the most difficult conditions, it was* ***55 times as fast****."*

REALISING THE BUSINESS BENEFITS

Commvault offers some of the most advanced data protection solutions available on the market today. Advanced features such as data compression and client-side deduplication already deliver significant benefits for companies that want to optimise backup performance without upgrading their network infrastructure or storage capacity.

PORTrockIT provides a simple, reliable and cost-effective way to leverage these Commvault features across a WAN while minimising the effects of latency and packet loss on overall performance. By simply installing PORTrockIT appliances at the edges of an existing WAN, organisations can gain the full value of their investment in Commvault – keeping their data protected more effectively than ever before.

“With Commvault and PORTrockIT, your backup infrastructure works ***smarter, not harder****.”*

The combination of Commvault and PORTrockIT is an example of technologies working smarter, not harder. The volume of data that organisations need to manage and protect is increasing year-on-year, at a rate that would have seemed incredible ten years ago. As a result, simply throwing money at faster backup servers, more bandwidth and larger storage systems is no longer a practical option.

Organisations need to adopt backup technologies that help them get greater value from existing infrastructure by reducing the amount of data that needs to be copied and stored, and moving it across the network faster and with less impact on other systems. With help from Bridgeworks and Commvault, these technologies are now within every organisation’s grasp.

ABOUT THE AUTHOR

David Trossell has been part of the IT industry for over 30 years, working for infrastructure specialists such as Rediffusion, Norsk Data and Spectra Logic before joining Bridgeworks in 2000 as CEO/CTO. He is a recognised visionary in the storage technology industry, and has been instrumental in setting the company’s strategic direction and developing its innovative range of solutions. David is the primary inventor behind Bridgeworks’ intellectual property, and has authored or co-authored 16 international patents.

TAKE THE NEXT STEPS

To learn more about PORTrockIT and other smart networking solutions from Bridgeworks, please visit www.4bridgeworks.com, or call us on +44 (0) 1590 615 444.

BRIDGEWORKS