SPECTRA BLACKPEARL

Accelerating archive replication across wide area networks with PORTrockIT

BRIDGEWORKS

EXECUTIVE SUMMARY

BlackPearl is a Spectra Logic platform that aims to make it easy for organisations to preserve data simply, affordably and safely. It provides a RESTful interface that allows users to store files as objects, and uses an onboard cache to accelerate the movement of data on and off long-term storage media, such as disk arrays and tape libraries.

One of the key use cases for BlackPearl is to enable geographically dispersed archiving, where data is preserved at multiple sites to reduce the risk of data loss in the event of a disaster. BlackPearl makes this easy to achieve by simply connecting two BlackPearl devices across a wide area network (WAN) connection.

However, whenever data is transferred across a WAN, it is subject to the effects of latency and packet loss, which can have a dramatic impact on data throughput.

This paper shows how using BlackPearl together with PORTrockIT can transform the performance of data replication across the WAN by counteracting the impact of both latency and packet loss.

*"In our tests, PORTrockIT accelerated BlackPearl's replication performance significantly – in one case, completing the job more than **5 times faster**."*

In every scenario we tested, PORTrockIT was able to increase transfer rates significantly. In one case, it reduced the total transfer time from 247 minutes to just 47 minutes, completing the job more than five times faster.

BlackPearl can help your business create and maintain geographically dispersed archives, while PORTrockIT adds value by mitigating the risk of losing valuable data and making it easier to demonstrate compliance with data retention regulations.

WHY SPEED MATTERS

Over the past decade, as the world has continued to become more digital, most organisations have experienced massive growth in data volumes. According to Statista, global data creation is projected to grow to more than 180 zettabytes by 2025.

At the same time, there has been an increasing focus on questions of data protection, with governments, regulators and customers requiring ever more rigorous standards and service level agreements for long-term data retention and governance.

One of the most effective ways to ensure that your organisation's data is protected for the long term is to adopt a solution such as BlackPearl, which provides a combination of advanced data policy management software and scalable, resilient hardware. BlackPearl can automatically scan your file systems (without impacting performance), identify files that need to be archived, and use its built-in cache to efficiently transfer them to long-term storage media, such as tape.

However, to make this archiving strategy resilient, it is critical to store multiple copies of your archived data and keep them at different locations. The most convenient and automated way to maintain and update these copies is to use a pair of BlackPearl appliances to replicate the data over a WAN connection.

For technical reasons (which we will explore in the next section), moving large amounts of data across a WAN can be a slow process – and as the size of the files to be replicated increases, this can cause major problems.

For example, if replication processes cannot be completed within the allotted time-window and overrun into business hours, they can have a severe impact on the performance of other business-critical systems.

If these overruns become the norm, the business may be forced to reevaluate its archiving strategy. It could reduce the amount of data it replicates, but this would mean protecting less of its data. It could replicate less often, but this could result in data not being available in the event of a disaster. It could extend the replication window, but this could potentially inconvenience users of other systems.

If none of these tradeoffs are acceptable, a final option might be to try to boost performance by investing in even more expensive server and network infrastructure. However, as we shall see, this is unlikely to solve the problem, because it doesn't address the root causes of poor WAN performance.

THE PROBLEMS: LATENCY AND PACKET LOSS

The chief culprit for poor WAN replication performance is latency – the time delay between a source system sending a packet across the network, and the target system receiving that packet.

The main cause of latency is the physical distance that the packet has to travel. Even with high-speed fibre-optic cabling, latency can increase at a rate of up to 5 microseconds per kilometre travelled.

In addition, the time taken to receive, queue and process packets at either end of the connection, and at any intermediate gateways, can add significantly to the total round-trip time for a system to send a message and receive a response. The further the data has to travel, and the more gateways it has to pass through, the greater the latency.

For network traffic sent via the TCP/IP protocol (and almost all replication traffic falls into this category), high latency can cripple transfer rates. TCP/IP works by sending a group of packets, then waiting for an acknowledgement that the packets have been received before it sends 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.

"Extra investment in bandwidth will simply be ***wasted*** *unless latency and packet-loss issues can be addressed."*

When packet loss occurs, the situation gets even worse. If a packet is lost before it is received by the recipient, or the acknowledgement goes astray before it reaches the sender, 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 falls even further, because the sender is sending fewer packets in the same amount of time.

Companies 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

Bridgeworks 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 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 re-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.

This process of creating virtual connections can also be applied across multiple physical WAN links, seamlessly sending the data across them all, even if only one TCP connection is used to supply the data. This can help to maximise the utilisation of your existing WAN links and provide extra redundancy to preserve connectivity in the event that one of the physical connections fails.

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

"PORTrockIT delivers ***faster network transfer performance****, without any need to make any changes to the rest of the network."*

TURNING THEORY INTO PRACTICE

To demonstrate the results that PORTrockIT can deliver for BlackPearl customers, Bridgeworks conducted a set of performance tests at a testing facility in the UK.

The test infrastructures mimicked a real-world BlackPearl architecture, using an "Injector" server to inject or extract files from BlackPearl, and a WANulator to simulate different levels of latency across a wide area network. Performance figures were taken from the BlackPearl console and recorded.

The tests assessed WAN data transfer performance at various simulated levels of latency. The WANulator also simulated a nominal packet loss rate of 0.1%, representing a typical service-level agreement from most WAN providers.

Test 1 and **Test 2** were run first with PORTrockIT set to bypass mode, to simulate an unaccelerated network architecture. The tests were then re-run with PORTrockIT in active mode, to simulate an accelerated network architecture. The figures from the two test runs were then compared to assess the difference in performance between the unaccelerated and accelerated architectures.

Test 3 used a slightly different setup, with a pair of Bridgeworks WANrockIT appliances in place of the PORTrockIT appliances. This is because WANrockIT is specifically designed to enable the transfer of SAN block-level protocols across a WAN, which is necessary for testing the specific use case of transferring data between a local BlackPearl and a remote tape library that are linked by a fibre channel connection.

Across all three tests, the most important metrics to consider are:

- **Average transfer rate**, measured in megabytes per second. This allows us to assess the speed at which files can be transferred across the WAN.
- **Total time to completion**, measured in minutes. This allows us to judge how long the entire file transfer job takes.

TEST 1: CACHE-TO-CACHE

Test 1 assessed the performance of transferring a 1 TB payload of files from the local BlackPearl cache to the remote BlackPearl cache and back again at 20, 40, 60, 80 and 100 ms of latency with 0.1% packet loss.

First, a number of files were created on the "Injector" server, and transferred to the local BlackPearl. These files were then replicated across the WAN to the remote BlackPearl.

Next, the local BlackPearl cache was cleared, and the Injector requested the files from the local BlackPearl. This caused the local BlackPearl to request the files from the remote BlackPearl over the WAN.

Figure 1: Accelerated and unaccelerated performance at 20ms latency with 0.1% packet loss during Test 1.

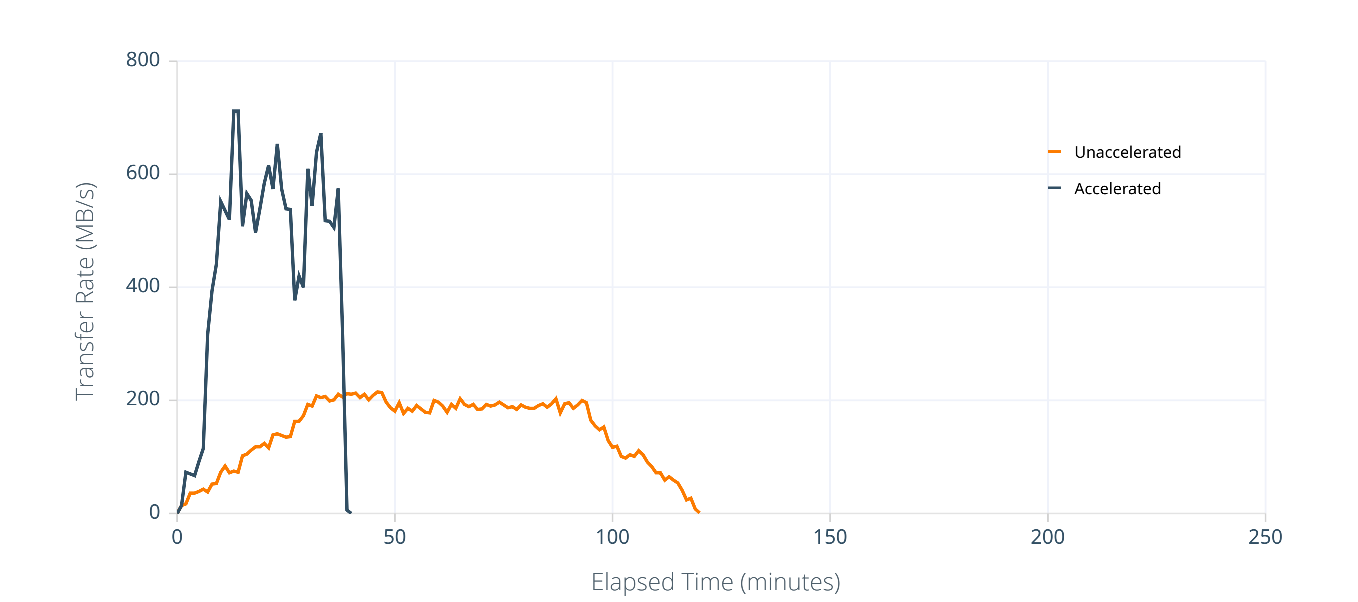

Figure 2: Accelerated and unaccelerated performance at 100ms latency with 0.1% packet loss during Test 1.

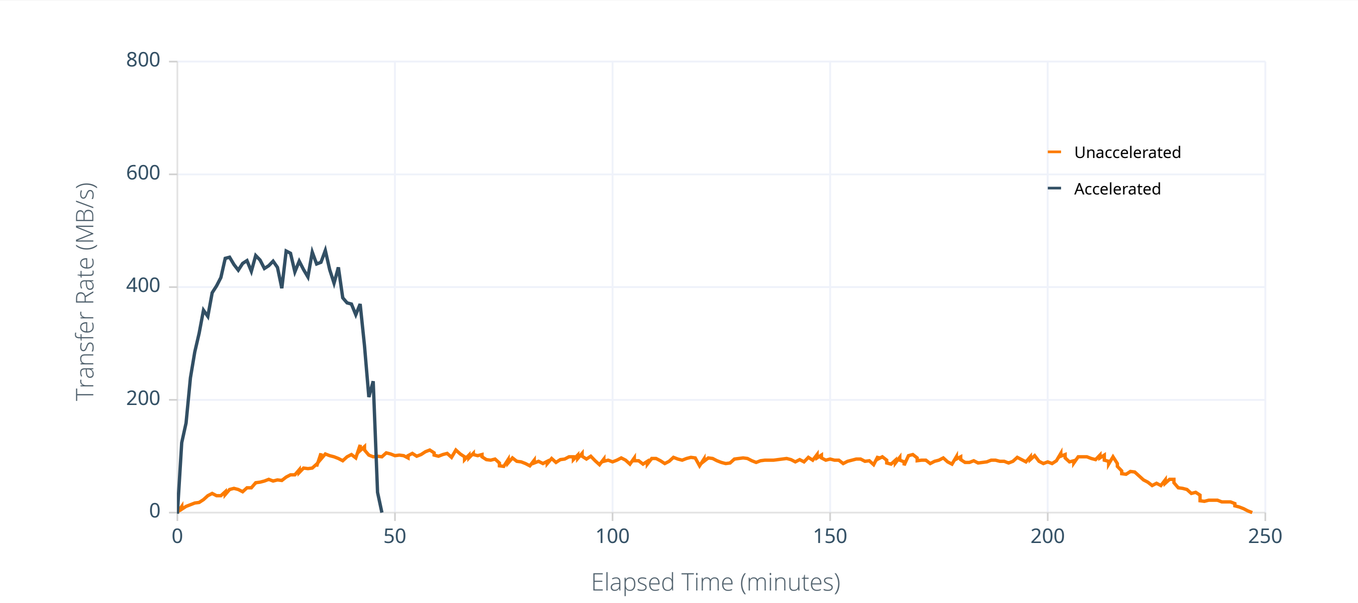

Figure 3: Network topology used in Test 1.

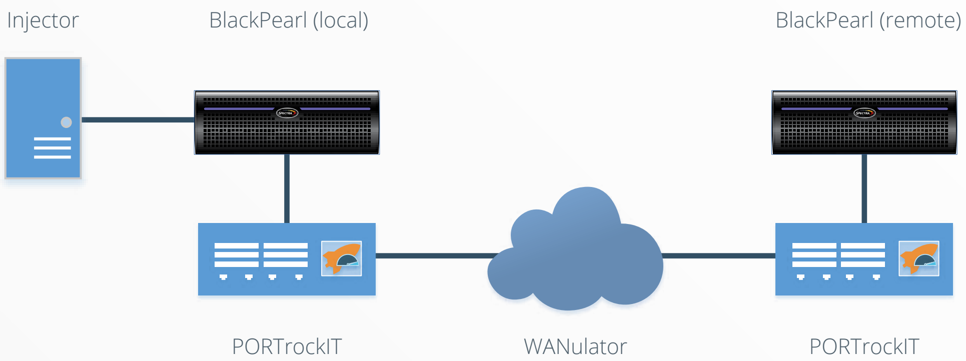

Figure 4: Average transfer rates for accelerated and unaccelerated architectures at various latencies with 0.1% packet loss during Test 1.

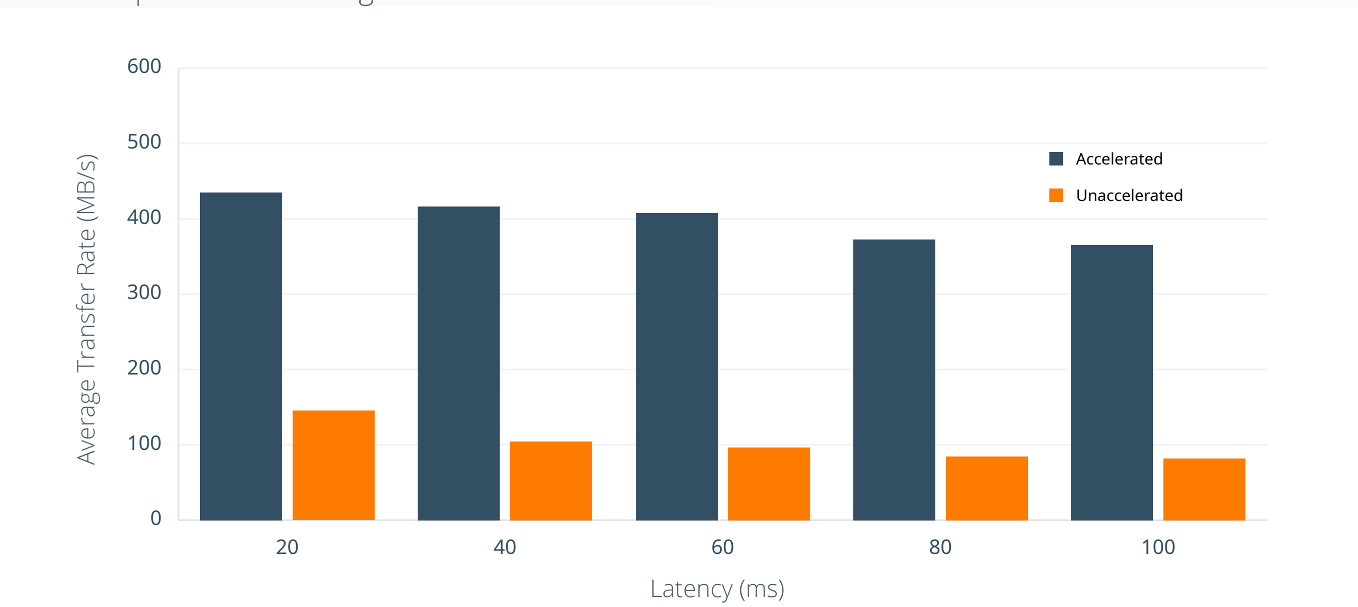

Figure 5: Total time to completion for accelerated and unaccelerated architectures at various latencies with 0.1% packet loss during Test 1.

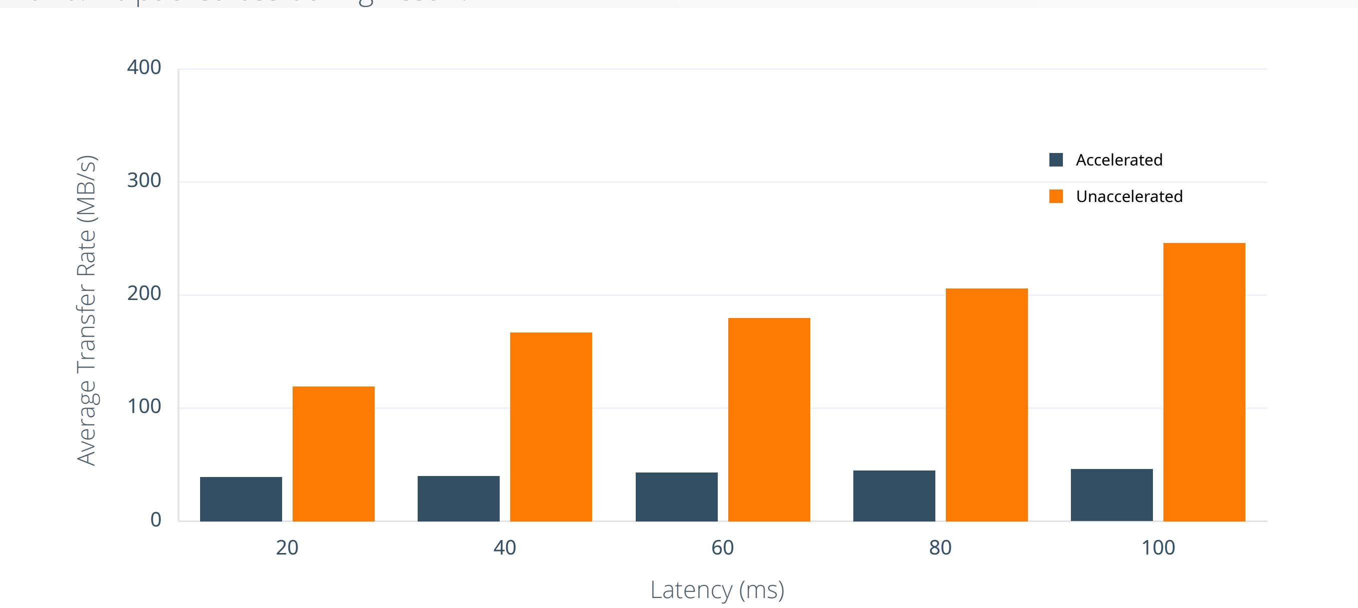

TEST 2: CACHE-TO-CACHE WITH REMOTE TAPE LIBRARY

Test 2 extended Test 1 by adding a tape library connected to the BlackPearl appliance at the remote end of the WAN. This simulated a scenario where BlackPearl is used as a distributed archive for long-term data retention.

The Injector filled up the local and remote BlackPearl caches, and the remote cache was then pushed out to a tape library connected to the remote BlackPearl.

Next, both caches were cleared, and the Injector requested the files from the local BlackPearl. This forced the remote BlackPearl to read from the tape library and pass the data across the WAN to the local BlackPearl and back to the Injector.

Once again a 1 TB payload of files was transferred at 20, 40, 60, 80 and 100 ms of latency with 0.1% packet loss.

Figure 6: Accelerated and unaccelerated performance at 20ms latency with 0.1% packet loss during Test 2.

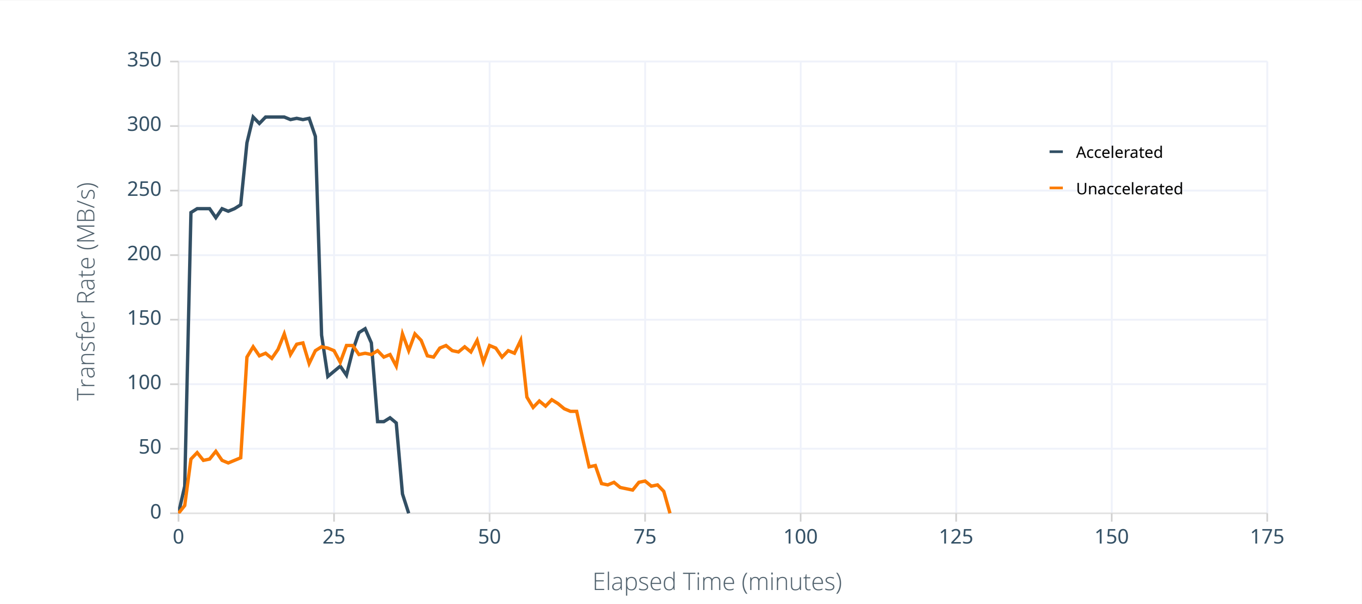

Figure 7: Accelerated and unaccelerated performance at 100ms latency with 0.1% packet loss during Test 2.

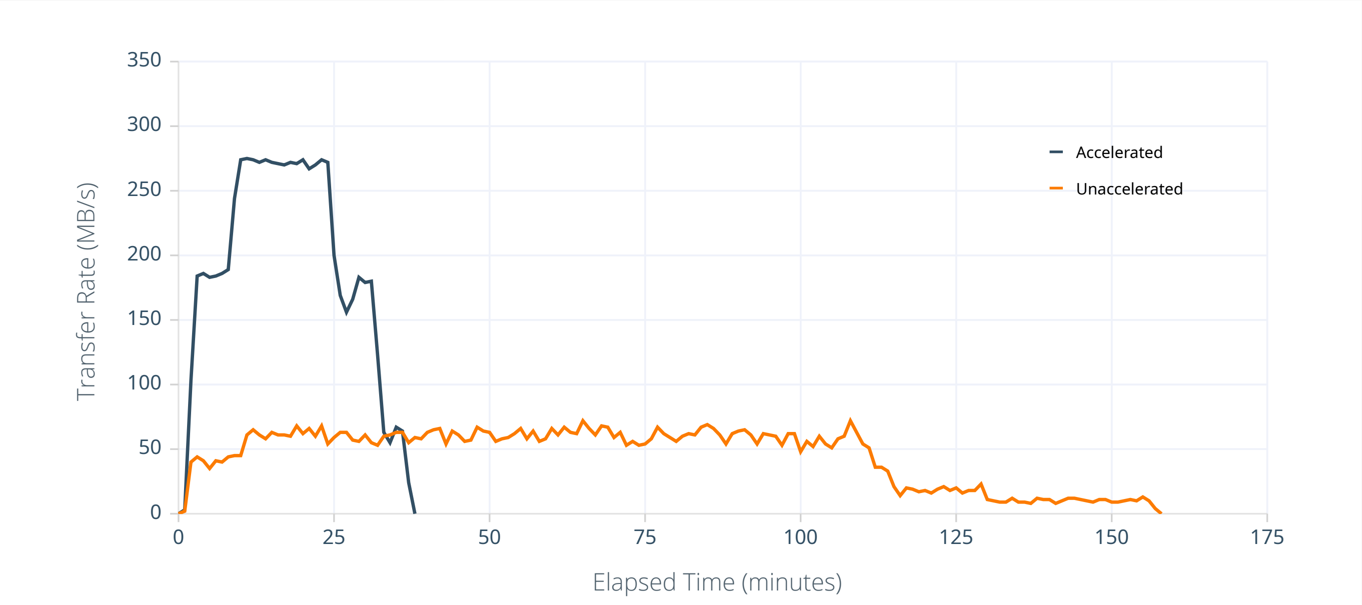

Figure 8: Network topology used in Test 2.

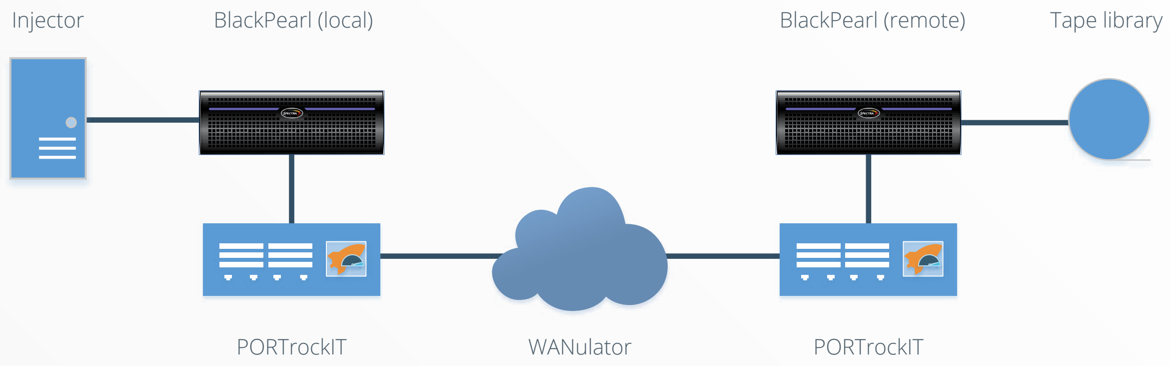

Figure 9: Average transfer rates for accelerated and unaccelerated architectures at various latencies with 0.1% packet loss during Test 2.

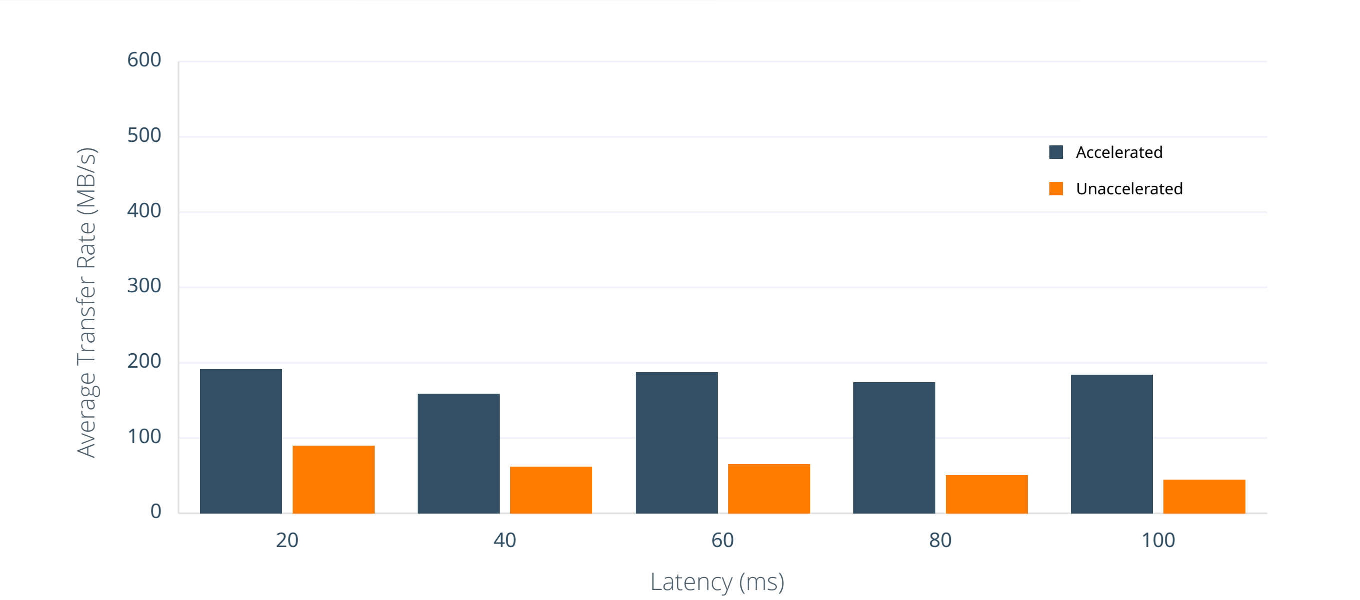

Figure 10: Total time to completion for accelerated and unaccelerated architectures at various latencies with 0.1% packet loss during Test 2.

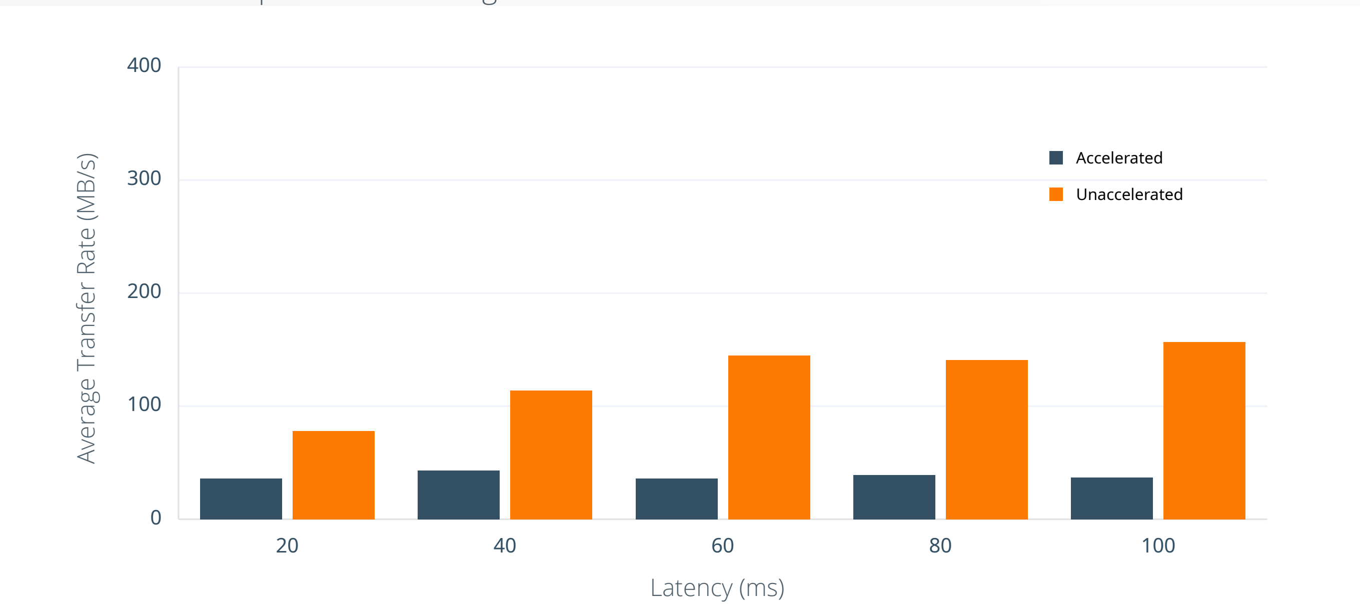

TEST 3: REMOTE TAPE LIBRARY ACROSS THE WAN

In **Test 3**, one BlackPearl appliance was placed on the local side of the WANulator and connected via a fibre channel connection to a remote tape library, using WANrockIT as the accelerator. This would test the speed of reading directly from tape across the WAN.

The Injector filled the local BlackPearl cache, which was then sent across the WAN to be stored on the tape library. Next, the local cache was cleared, and the Injector requested the data from the BlackPearl. This forced the BlackPearl to pull the data back across the WAN from the remote tape drive.

The payload for this test was 400 GB of files, transferred using the accelerated architecture at latencies of 80 ms and 100 ms with 0.1% packet loss. This represents typical latencies between data centres located on each coast of the USA, or between continents.

Figure 11: Accelerated performance at 80ms latency with 0.1% packet loss during Test 3.

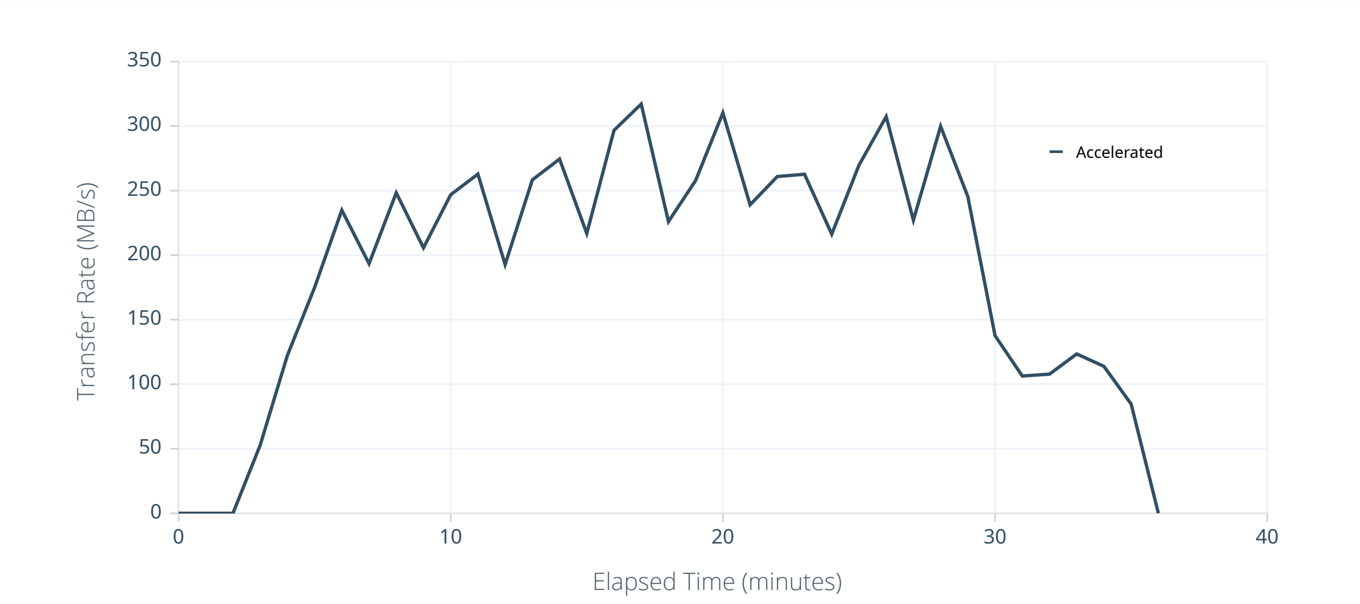

Figure 12 Accelerated performance at 100ms latency with 0.1% packet loss during Test 3.

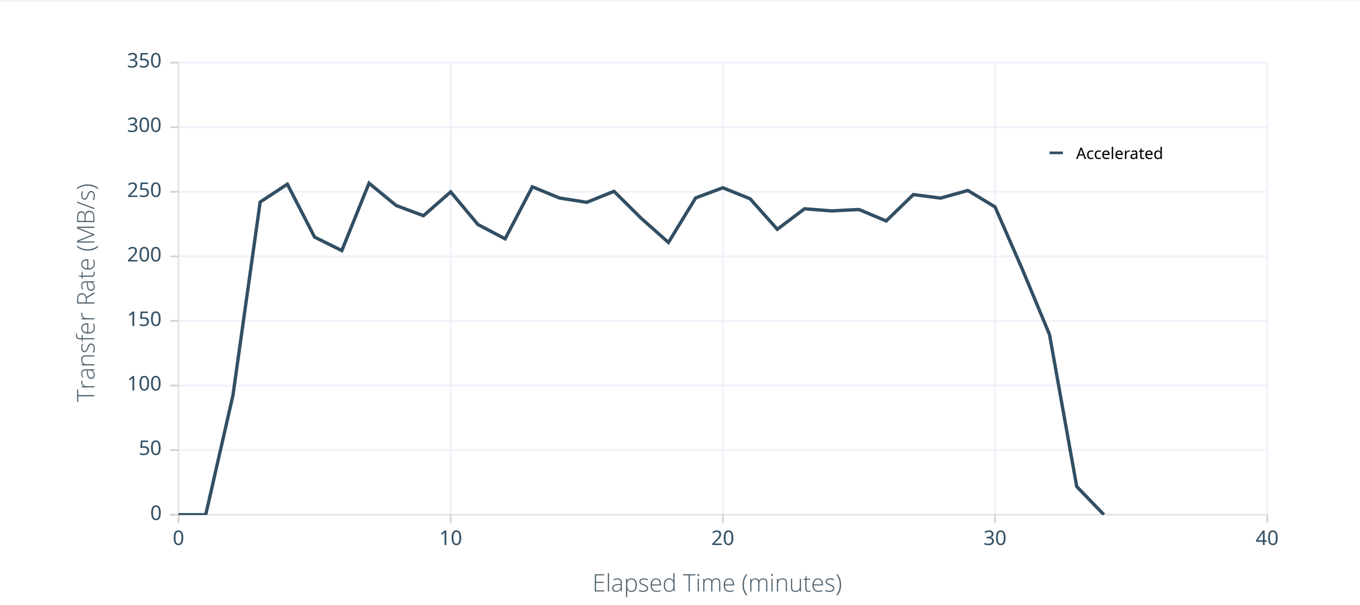

Figure 13: Network topology used in Test 3.

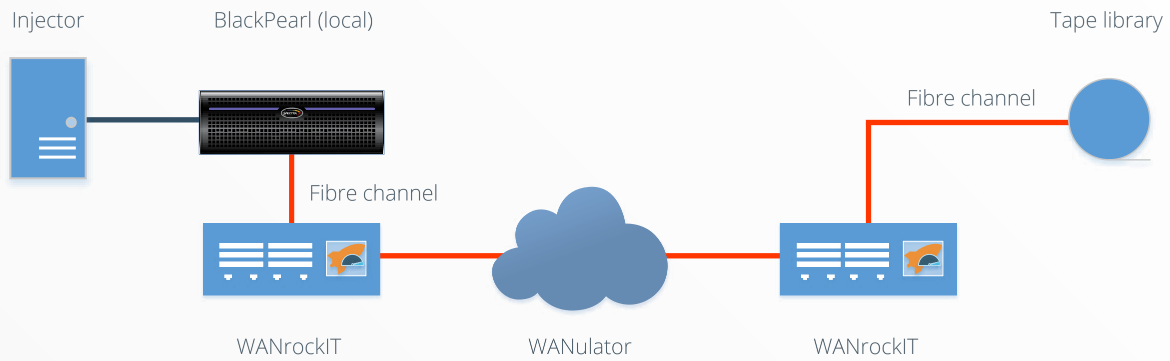

WHAT THE DATA TELLS US

WITHOUT PORTROCKIT

For the unaccelerated network architecture, the results of **Test 1** and **Test 2** show a negative correlation between latency and performance. As latency increases, average transfer rates (ATR) decline and we see a longer total time to completion (TTC).

For example, in **Test 1**, the TTC was 119 minutes when there was 20 ms of latency, but rose to 246 minutes at 100 ms of latency – more than twice as long.

The results of **Test 2** were similar: the TTC was 78 minutes at 20 ms of latency, but 157 minutes at 100 ms of latency – again, more than double the time to complete the transfer.

In a real-world scenario, this variability in TTC would be a major concern, as any unexpected increase in latency could cause transfers to time-out or overrun their scheduled window – putting valuable business data at risk.

WITH PORTROCKIT

When we look at the architecture accelerated by PORTrockIT, the results are very different. At all levels of latency, the ATR improves dramatically when PORTrockIT is activated. For example, in **Test 1**, the worst-case ATR of the accelerated architecture (365 MB/s at 100 ms of latency) is still more than 2.5 times as fast as the best-case ATR of the unaccelerated architecture (145 MB/s at 20 ms of latency).

We can also see that PORTrockIT keeps TTC stable, regardless of latency. In **Test 1**, the shortest TTC was 39 minutes at 20 ms of latency, and the longest was 46 minutes at 100 ms of latency. In **Test 2**, the fastest TTC was 36 minutes and the slowest 43 minutes.

In both cases, that's a difference of just 7 minutes between the slowest and fastest transfers – which suggests that real-world users can expect their transfers to complete on time, even if latency increases.

REALISING THE BUSINESS BENEFITS

Our initial investigation into the transfer of data between BlackPearl devices highlights the extensive work that Spectra Logic has undertaken to enhance the performance of the transfer process.

However, in the results of **Tests 1 and 2**, we can observe the combined effects of latency and packet loss on BlackPearl's ability to transfer data across a WAN at high speeds. The root cause of the performance degradation is uncertain, although it is possible that latency and packet loss cause cache-to-cache synchronisation to stall, which in turn creates backlogs.

Our tests demonstrate that when using PORTrockIT's artificial intelligence capabilities to actively mitigate the effects of latency and packet loss, it is possible to make significant gains in WAN transfer performance.

With PORTrockIT, customers are able to maximise throughput across the WAN when synchronising and/or migrating the BlackPearl cache, or recovering data from remote tape devices. This not only results in faster, more reliable transfers – it also gives customers greater flexibility in where to deploy their BlackPearl devices without sacrificing performance.

Finally, one of the key features of BlackPearl is its ability to offload data from its cache to object storage-based tape for long-term storage. Remote tape storage is a key area of focus at the moment, as it adds an extra layer of protection against data security attacks and helps to ensure that data can always be restored in a disaster recovery scenario.

As **Test 3** demonstrates, by using Bridgeworks WANrockIT, a WAN acceleration product for SAN protocols, Spectra Logic customers can not only geographically separate their BlackPearl instances, but also use a third geographical site for air-gapped tape storage facilities without suffering the performance penalties associated with latency and packet loss.

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. David 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 our website or call us today:

- www.4bridgeworks.com
- **UK:** +44 (0) 1590 615 444
- **USA:** +1 888 238 1805