NETAPP SNAPMIRROR

Proving the concept of accelerated replication with PORTrockIT

BRIDGEWORKS

EXECUTIVE SUMMARY

For one major US healthcare provider, the need to protect large volumes of clinical image files by replicating them to a disaster recovery site was causing problems. With significant packet loss on the connection between its two data centres, transfer speeds slowed to a snail's pace – making it impossible to complete the NetApp SnapMirror replication job within a timeframe acceptable to the business.

The company's IT partners, Trace3 and Agilesys, knew of a technology that could potentially solve the problem: PORTrockIT from Bridgeworks. They offered to run a comprehensive proof-of-concept exercise to demonstrate how PORTrockIT could counteract the effects of packet loss, and transform the performance of the replication process.

This paper explores the results of the proof-of-concept, showing how PORTrockIT significantly improved average performance across more than 300 individual test scenarios, including different levels of packet loss, different types of data, and different configurations of the NetApp SnapMirror replication process.

For example, even at 1.5% packet loss (the highest packet loss rate tested), the PORTrockIT solution was able to maintain a transfer rate of between 68 and 98 MB/s, in scenarios where an unaccelerated connection could only deliver speeds between 1 and 7 MB/s.

> *"PORTrockIT achieved transfer rates up to* ***68 times faster*** *than an unaccelerated architecture."*

These results have helped to build a strong business case for the healthcare provider to adopt PORTrockIT – and provide compelling evidence that a similar solution could benefit almost any organisation that uses NetApp SnapMirror for data replication across a wide-area network.

WHY SPEED MATTERS

In the healthcare sector, data volumes are growing rapidly. In particular, the drive to digitise medical records, x-ray and pathology results is generating terabytes of image files, full of sensitive private data that needs to be stored safely and securely, potentially for decades.

In most cases, the safest way to protect this data is to replicate it off-site, to a geographically remote data centre. However, image files can present a challenge for backup and replication solutions, especially those that use data compression to optimise performance, such as NetApp SnapMirror. Since most image formats are already highly compressed, the backup solution cannot slim them down much further, and therefore cannot provide the usual performance gains. As a result, replication jobs can take longer for images than they would for a similar volume of more highly compressible data.

As the amount of data that needs to be protected increases, it can become difficult to complete off-site replications within a reasonable time-window. For example, if an organisation wants to replicate its data every day, it may only have a few hours during the night when few users are online, when the replication job can be executed without affecting performance for business-critical systems. This scenario applies particularly to hospitals and healthcare organisations that operate 24/7, which may have only very brief periods of low activity.

To exacerbate the problem, if a NetApp SnapMirror replication job fails before it is completed, it must be restarted from the beginning. As jobs get larger, both the risk and the consequences of failure become more severe, and the likelihood of overrunning the available window increases.

"PORTrockIT helps to avoid backup issues by ***accelerating*** *the replication process."*

If the data cannot be fully replicated within the required window, the organisation may need to make some tough choices. It could extend the window, but that may impact the performance of other systems. Alternatively, it could either replicate less data, or replicate the same amount of data less frequently. However, this might leave data unprotected, or make it more difficult to recover in the event of a disaster.

A better solution is to find a way to accelerate the replication process – such as PORTrockIT.

THE PROBLEMS: LATENCY AND PACKET LOSS

In general, there are two main issues that cause the majority of performance problems when replicating data across a wide area network (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, and the time taken to receive, queue and process packets at either end of the connection and at any intermediate gateways. The further the data has to travel, and the more gateways it has to pass through, the greater the latency.

For replication jobs 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, 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.

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

> *"Investing in expensive high-bandwidth network infrastructure* ***will not fix*** *TCP/IP problems caused by latency or packet loss."*

Organisations often try to solve TCP/IP performance problems 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 re-used to send another set of packets, and the whole process repeats.

This parallelisation combats 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.

"PORTrockIT uses ***artificial intelligence*** *to accelerate data transfer across multiple parallel connections."*

In practical terms, PORTrockIT is installed as a pair of physical or virtual appliances, deployed at either end of the WAN. The backup server 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 recipient server. 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

A major US healthcare company wanted to replicate large quantities of data – mainly image files – from its main data centre to a secondary location several miles away. The image data was already highly compressed, and due to its sensitive nature, it also needed to be encrypted – limiting SnapMirror's ability to optimise network performance through compression. Moreover, due to the long distance between the sites, the WAN connection suffered from relatively high levels of both latency and packet loss.

The company's IT partners, Trace3 and Agilesys, proposed a solution based on NetApp SnapMirror and Bridgeworks PORTrockIT appliances. To demonstrate to their client that the solution would work in practice, they offered to run a comprehensive proof-of-concept exercise at Trace3's testing facility in Phoenix, Arizona.

The proof-of-concept infrastructure mimicked a real-world SnapMirror architecture, using a WANulator to simulate different levels of latency and packet loss between the source (the NetApp storage unit that would send the data), and the target (the NetApp unit that would receive the replicated data).

All the tests were performed first on an unaccelerated architecture, where the source and target were connected directly to the WANulator (see figure 1). The same tests were then repeated on an architecture that was accelerated by introducing two physical PORTrockIT appliances, which sat either side of the WANulator, between the source and the target (see figure 2).

A battery of more than 300 tests was designed to simulate different SnapMirror replication configurations (for example, sync, async and semi-sync settings in ONTAP cluster mode), and different size and types of data (for example, replications jobs containing

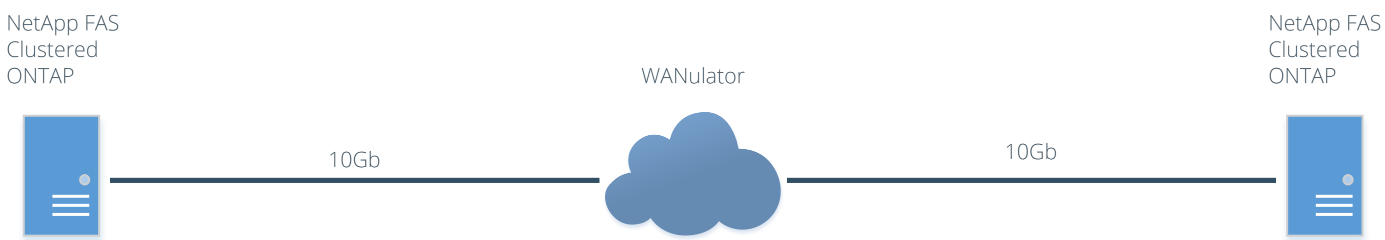

Figure 1: Unaccelerated architecture

multiple 10+ KB files versus multiple 10+ GB files, and jobs containing highly compressible files versus compressed/ encrypted files).

The performance of each of these configurations was then tested under different levels of latency and packet loss, and the results from all of the configurations were used to calculate NetApp SnapMirror's average performance in the various latency and packet loss scenarios.

After conducting each performance test, a data integrity check was performed to confirm that data was not lost or corrupted during testing, and that the files written to and read from the NetApp storage were identical to the original source files.

TEST EQUIPMENT

SOFTWARE:

- WANulator ISO image
- Test Data Generation Tool

HARDWARE:

- Windows Server 2008 R2
- 2 x NetApp FAS8060 units
- 2 x PORTrockIT 210 nodes
- WANulator host

NETWORK:

- Two standard 10 Gb/s LAN connection management ports
- One 10 Gb/s 10 GbE interface card (with 1 SFP+ LC 50 micron multimode fiber connector)

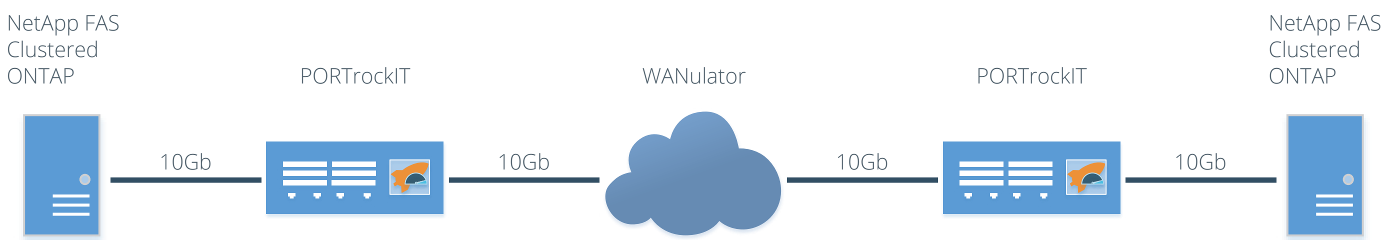

Figure 2: Accelerated architecture with PORTrockIT

WHAT THE DATA TELLS US

LATENCY

The first set of tests simulated a scenario with no packet loss, at latencies ranging from 0 ms to 360 ms round trip time (RTT). Data was replicated from the source to the target using NetApp SnapMirror, first via the unaccelerated architecture, and then again via the accelerated architecture with PORTrockIT.

Looking at Figure 3, the results show that there was relatively little difference between the two architectures – both delivered very high performance at low levels of latency, with a sharp dip in transfer rate as latency neared 100 ms.

At these higher levels of latency, the PORTrockIT architecture did offer a significant performance improvement – a 44% increase at 100 ms, and a 45% increase at 150 ms.

The results show that even without PORTrockIT acceleration, SnapMirror provides an efficient means of replicating data across a WAN when packet loss is not an issue. High levels of latency do impact performance, but transfer rates of around 100 MB/s can still be considered acceptable in most replication scenarios.

However, next, let's look at what happens when packet loss is introduced into the equation.

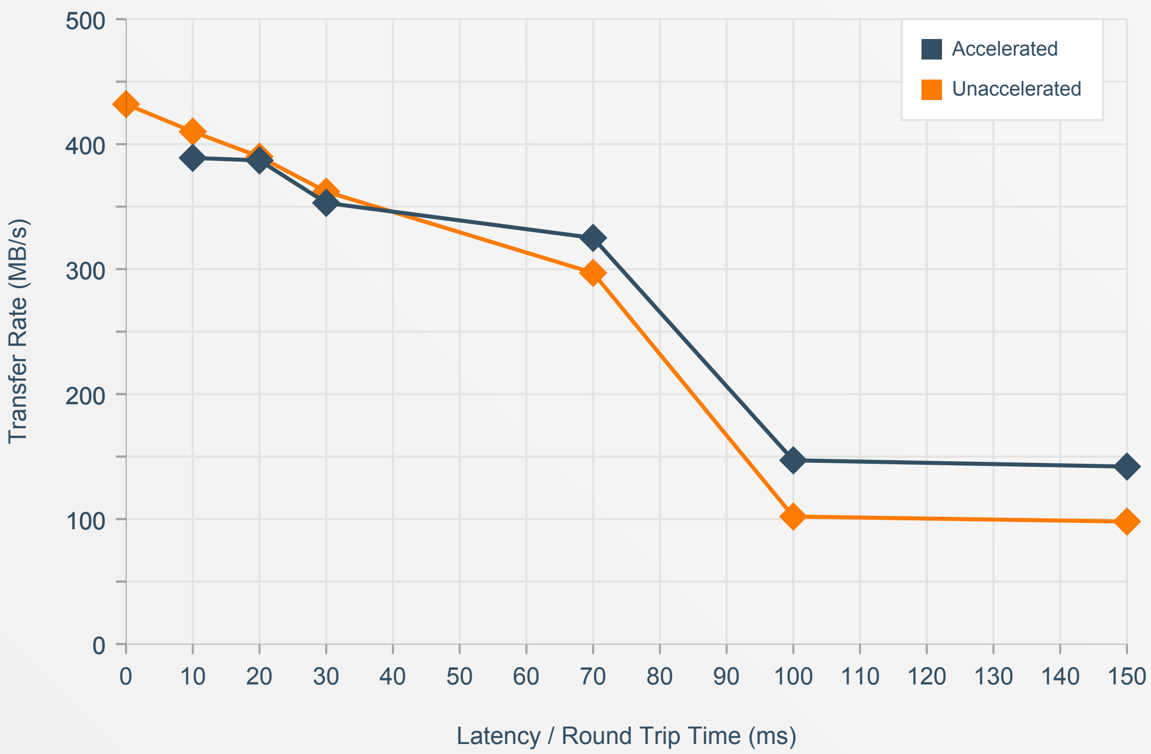

Figure 3: Accelerated and unaccelerated performance at various latencies with 0% packet loss

PACKET LOSS

Three other scenarios with different levels of packet loss (0.5%, 1% and 1.5%) were also tested at various levels of latency across multiple different SnapMirror configurations. Figures 4, 5 and 6 all show that the unaccelerated architecture saw severe performance degradation from the introduction of packet loss. In all three scenarios, the accelerated architecture performed considerably better.

Even in the most extreme example (150 ms of latency with 1.5% packet loss) the accelerated architecture achieved a transfer rate of 68 MB/s, compared to just 1 MB/s on the unaccelerated architecture.

The results lead to the conclusion that while NetApp SnapMirror is efficient at replicating data across an ideal network, it can struggle to achieve the same performance in a real-world environment where significant levels of packet loss occur. Even where network latency is very low, a small amount of packet loss can degrade performance to a point where replication jobs run so slowly that they are likely to overrun the available window.

In such situations, the adoption of PORTrockIT can have a very significant impact, keeping transfer rates at 68 MB/s or above even in the most difficult scenarios.

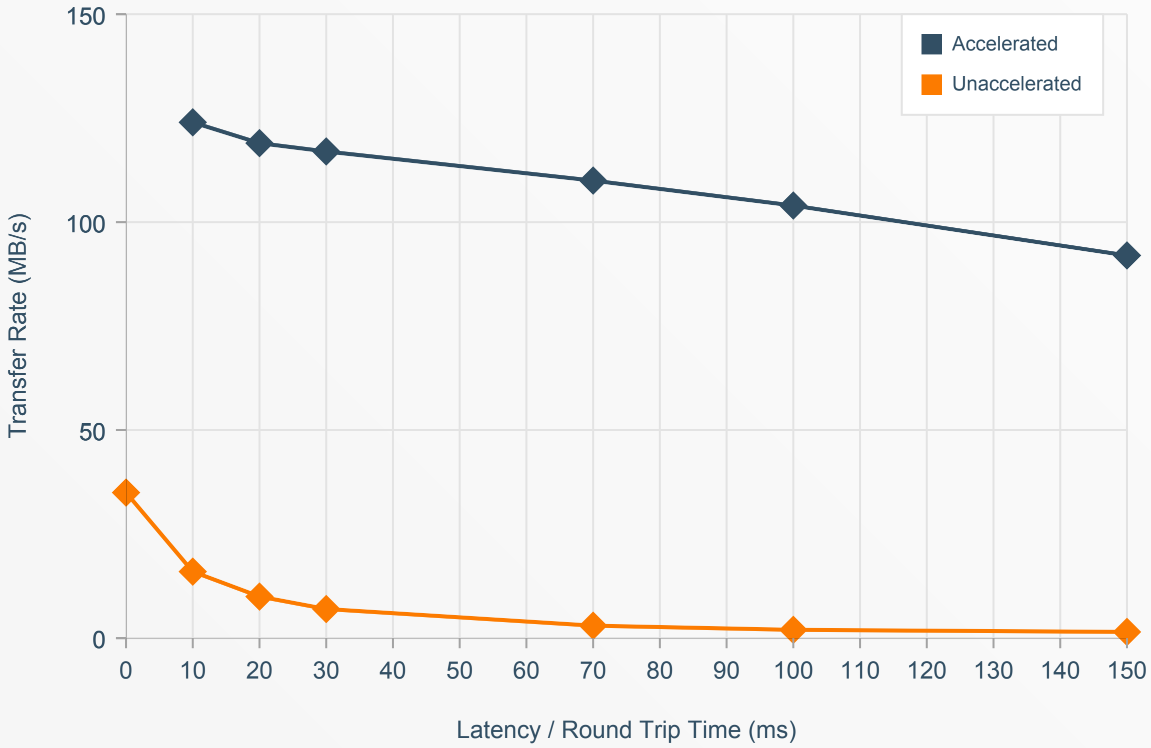

Figure 4: Accelerated and unaccelerated performance with 0.5% packet loss at various levels of latency

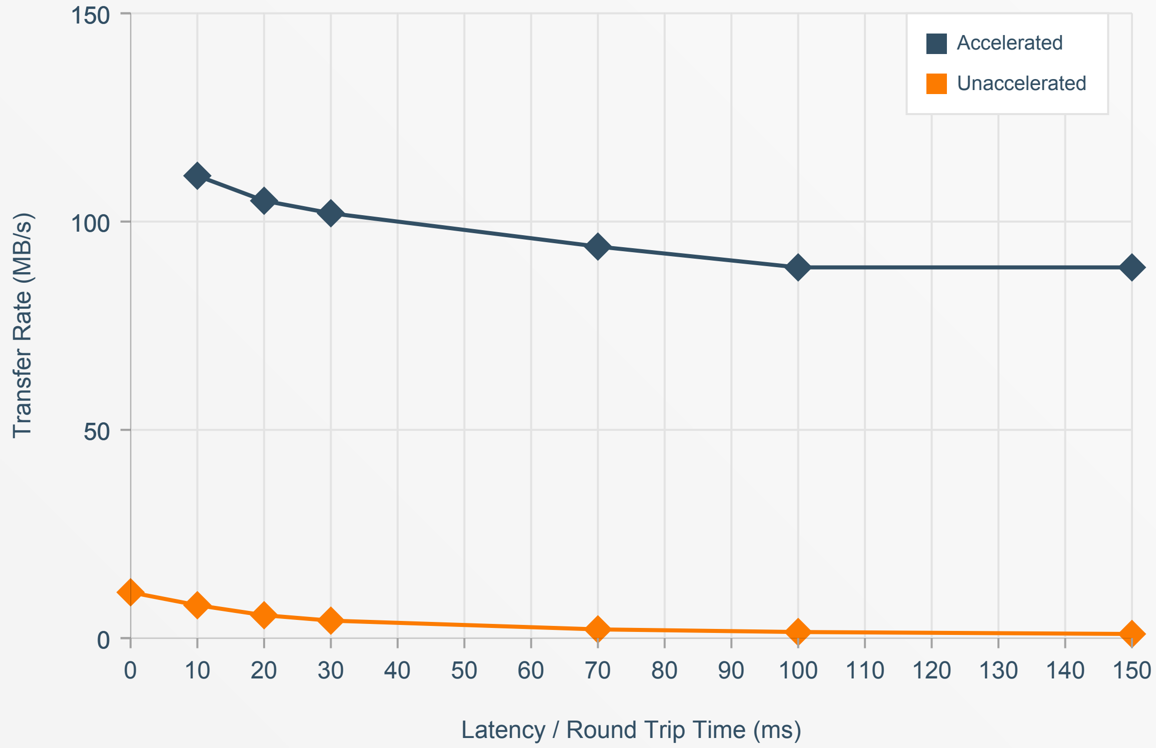

Figure 5: Accelerated and unaccelerated performance with 1% packet loss at various levels of latency

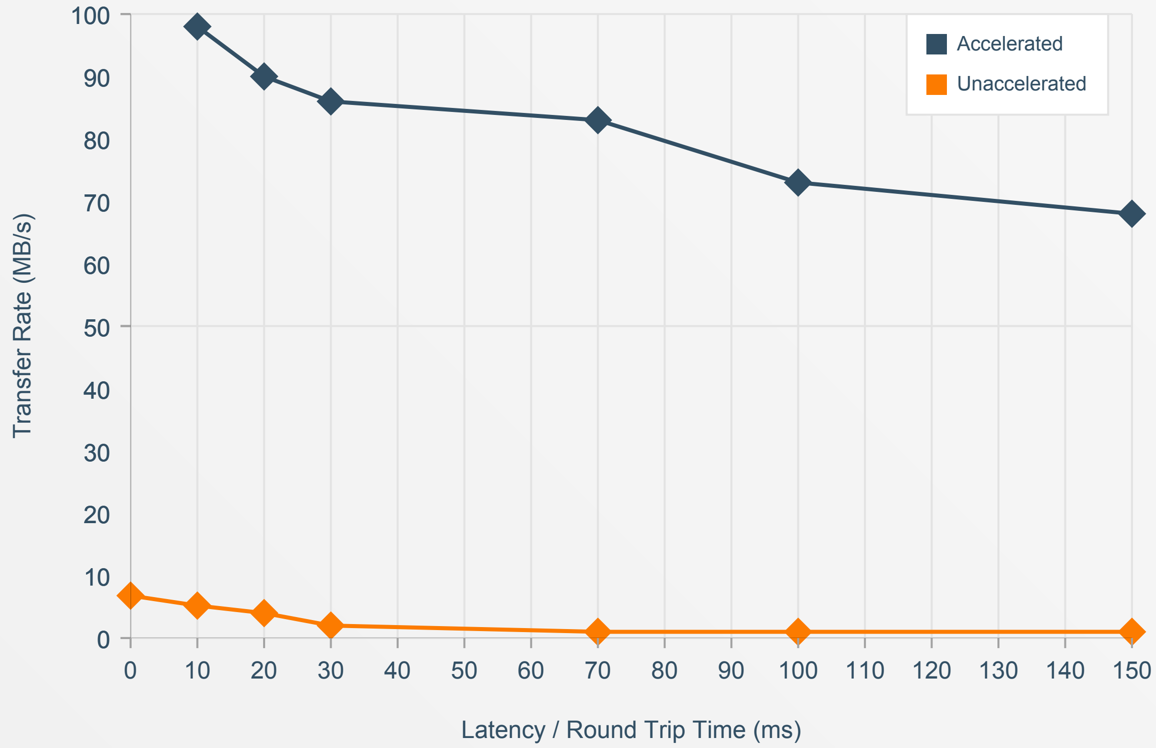

Figure 6: Accelerated and unaccelerated performance with 1.5% packet loss at various levels of latency

REALISING THE BUSINESS BENEFITS

The proof-of-concept exercise enabled Trace3 and Agilesys to help their client make a strong business case for investing in PORTrockIT. The tests showed that the Bridgeworks technology would make it feasible to replicate large volumes of sensitive, highly compressed and encrypted clinical image data across a WAN. This would help the healthcare provider meet its data protection, business continuity and disaster recovery goals, as well as facilitating compliance with laws and regulations around the safe and secure storage of patient data.

Looking beyond the specific context of the proof-of-concept, the results also suggest that PORTrockIT can transform performance for any organisation that wants to use NetApp SnapMirror for replication across a WAN. If replication jobs are threatening to overrun the available window, or if it is desirable to reduce the replication window to free up server and network resources for other important jobs, PORTrockIT provides an elegant solution.

PORTrockIT offers plug-in-and-go technology that can be implemented quickly with minimal impact on the rest of your IT infrastructure – keeping deployment cost and risk to a minimum. By maximising the performance of existing infrastructure, PORTrockIT also reduces the need to invest in expensive high-bandwidth connections or more powerful backup servers – enabling significant cost-avoidance.

"PORTrockIT delivers ***significant business value*** *by transforming WAN replication performance."*

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