VEEAM

Accelerating virtual machine replication with PORTrockIT

BRIDGEWORKS

EXECUTIVE SUMMARY

Business continuity solutions such as Veeam offer the ability to recover quickly from disaster by creating a replica of each virtual machine (VM) in your production environment, and enabling instant failover from the production VM to the replica in case of disaster.

However, as the number and size of VMs increase, replicating them across a wide area network (WAN) can be a time-consuming job that puts considerable strain on your infrastructure. The greater the distance between your production and recovery data centres, the worse the problem gets: network latency and packet loss can slow the replication process to a crawl, even on the most expensive high-bandwidth connections.

This paper shows how PORTrockIT can transform the performance of Veeam for VM replication across a WAN by counteracting the impact of both latency and packet loss.

"In every scenario we tested, PORTrockIT was able to increase the transfer rate significantly – in some cases transferring data more than ***80 times faster****."*

In every scenario we tested, PORTrockIT was able to increase the transfer rate significantly – in some cases transferring data more than 80 times faster across the same network connection.

PORTrockIT can help your business replicate VMs faster and more reliably – reducing the risk of disruption to your business continuity plans, and ensuring that you get real value from your network infrastructure investment.

WHY SPEED MATTERS

More and more companies are using VM replication as an efficient, convenient and safe method of encapsulating production servers and allowing them to be spun up instantly on different hardware. This is an excellent strategy for business continuity, as it substantially reduces the recovery time in the event of a disaster at the main data centre.

However, from a day-to-day perspective, VM replication may also create operational challenges. IT teams need to set up a robust process to replicate the production VMs and copy them across a WAN to a secondary data centre, which may be anything from a few hundred metres to a few hundred miles away. This means moving a large amount of data across the network, which generally needs to be done at least once per day to ensure that as little data as possible is lost if disaster strikes.

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 number and size of the VMs increase, this can cause major problems.

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

*"Veeam's advanced features can optimise traffic transmission – but they cannot solve the **underlying causes** of poor replication performance."*

If this process consistently runs over the available window, the business may be faced with tough decisions. It could reduce the number of VMs it replicates (and therefore protect less of its data). It could replicate VMs less often (and lose more data in the event of a disaster). It could extend the replication window (and potentially inconvenience users of other systems). Or it could invest in even more expensive server and network infrastructure to boost performance.

To reduce the need for such drastic measures, VM replication solutions such as Veeam provide advanced features to help mitigate performance issues. For example, Veeam optimises traffic transmission by filtering out unnecessary data blocks and compressing replica traffic.

However, while these traffic optimisation measures might buy the business some time, they do not address the underlying cause of poor replication 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.

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.

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

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

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.

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

In practical terms, PORTrockIT is installed as a pair of appliances, deployed at either end of the WAN. The Veeam source 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 Veeam target 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

To demonstrate the kind of results that PORTrockIT can deliver for Veeam customers, Bridgeworks conducted a set of performance tests at an independent testing facility in the UK. The test infrastructures mimicked a real-world Veeam architecture, using a WANulator to simulate different levels of latency and packet loss between the source (the server that runs the VM that is to be replicated), and the target (the server that will receive the replica of the VM).

The first set of tests were performed 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 PORTrockIT appliances, which sat either side of the WANulator, between the source and the target (see figure 2).

TEST EQUIPMENT

SOFTWARE:

- VMware ESXi 6.0
- Veeam Backup & Replication Server 8.0.0.2030

HARDWARE:

- Source server: Sun Fire x2250, 16GB RAM, 2x Intel Xeon X5472 3.0GHz, Veeam Backup & Replication Server running on Windows 2012 R2 virtual machine, Veeam Proxy running on Windows 2012 R2 virtual machine
- Target server: DELL 2950, 4GB RAM, Intel Xeon X5450 3.0GHz, Veeam Proxy running on Windows 2012 R2 virtual machine
- 2 x PORTrockIT Nodes
- WANulator host

WHAT THE DATA TELLS US

LATENCY

The first test simulated a scenario with no packet loss, at latencies ranging from 0 ms to 360 ms round trip time (RTT). A VMware ESXi virtual machine was transferred from the source to the target, first via the unaccelerated architecture, and then again via the accelerated architecture with PORTrockIT.

Looking at Figure 3, the results show that performance on the unaccelerated architecture degraded significantly as latency increased. By contrast, the accelerated architecture with PORTrockIT provided a stable transfer rate of more than 110 MB/s at all latencies. In the case of a network with 360 ms of latency, the performance gain with PORTrockIT was 99 MB/s – which means that PORTrockIT transferred the data 7.6 times faster.

Figure 1: Unaccelerated architecture

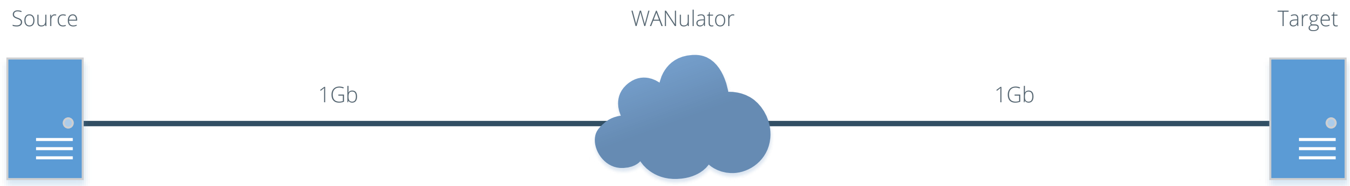

Figure 2: Accelerated architecture with PORTrockIT

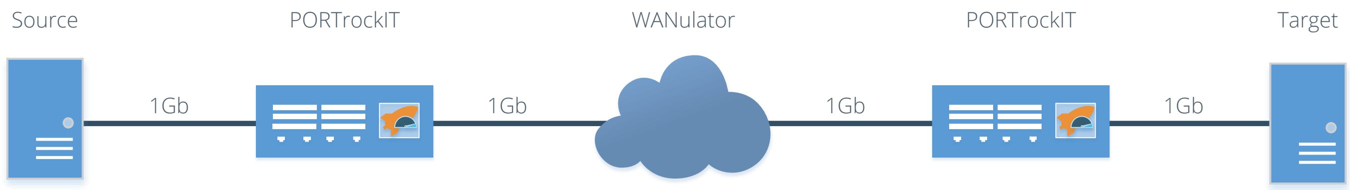

Figure 3: Accelerated and unaccelerated performance at various latencies with 0% packet loss

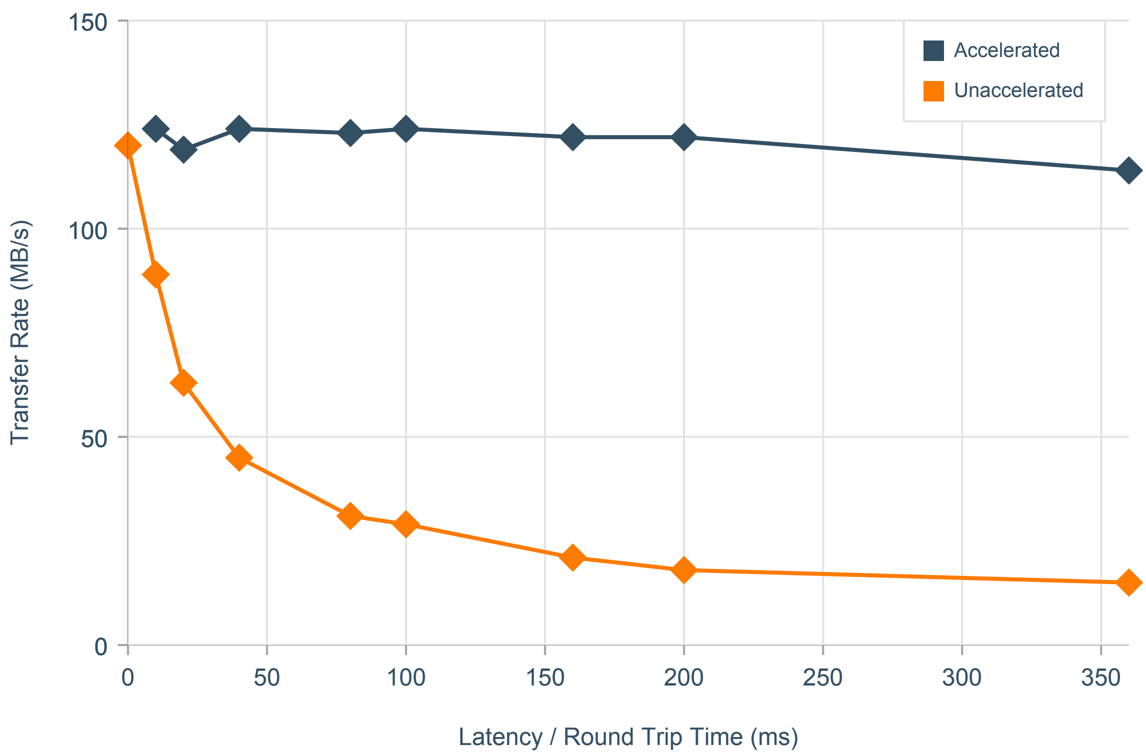

PACKET LOSS

*"With 3% packet loss, the PORTrockIT architecture was **6.2 times faster**."*

The second test investigated the performance of the two architectures on a network with zero latency, but with various levels of packet loss. Again, a VMware ESXi virtual machine was transferred from the source to the target using both the unaccelerated and accelerated architectures.

From Figure 4, we can see that for both architectures, performance degraded as packet loss increases. However, in all cases, performance was higher with the accelerated architecture, and the difference in performance increased as the level of packet loss rose. Where the packet loss rate was higher than 0.5%, the advantage of PORTrockIT becomes increasingly clear; and at 3% packet loss, the accelerated architecture was 6.2 times faster.

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

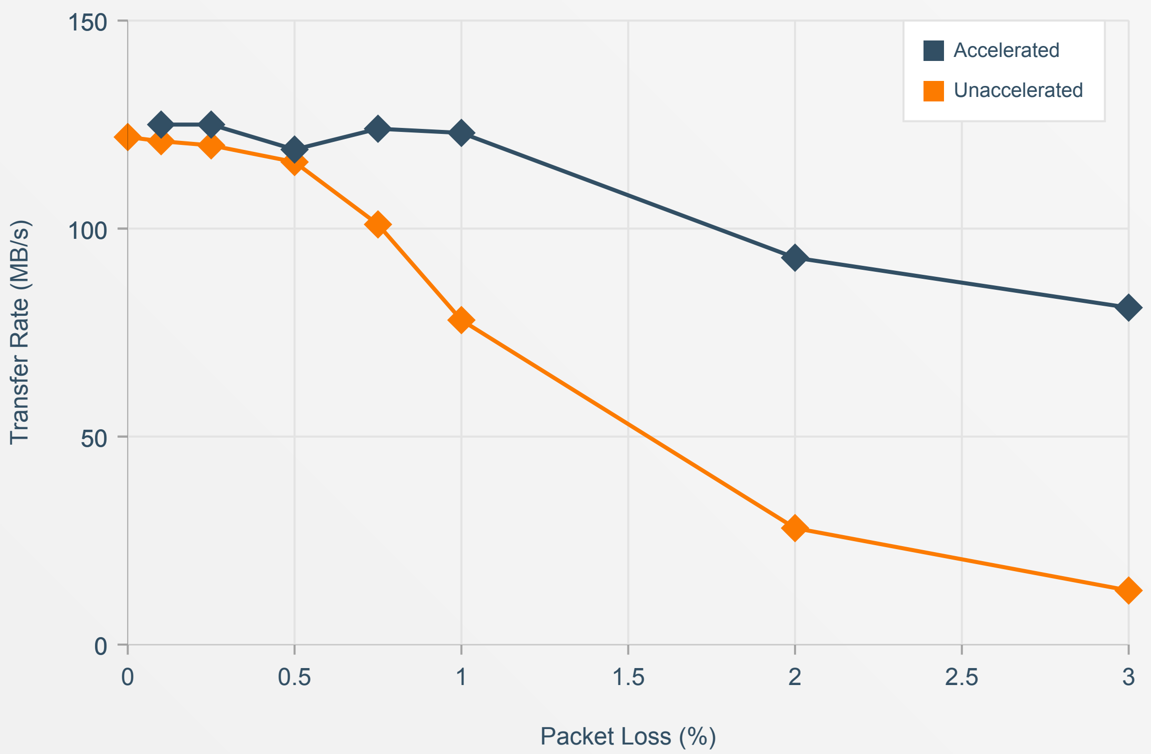

COMBINED EFFECTS OF PACKET LOSS AND LATENCY

Finally, the team decided to test three different packet loss scenarios (0.1%, 0.5% and 1%) at various levels of latency, to assess the combined effects of packet loss and latency on both architectures.

Figures 5, 6 and 7 all show that the unaccelerated architecture saw severe performance degradation from the combination of latency and packet loss. In all three scenarios, the accelerated architecture performed considerably better.

*"In the most extreme scenario of latency and packet loss, PORTrockIT was **87.3 times faster** than the unaccelerated architecture."*

Even in the most extreme example (360 ms of latency with 1% packet loss) the accelerated architecture achieved a transfer rate of 69 MB/s – 87.3 times faster than the unaccelerated transfer rate of 0.79 MB/s

Figure 5: Accelerated and unaccelerated performance with 0.1% packet loss at various levels of latency

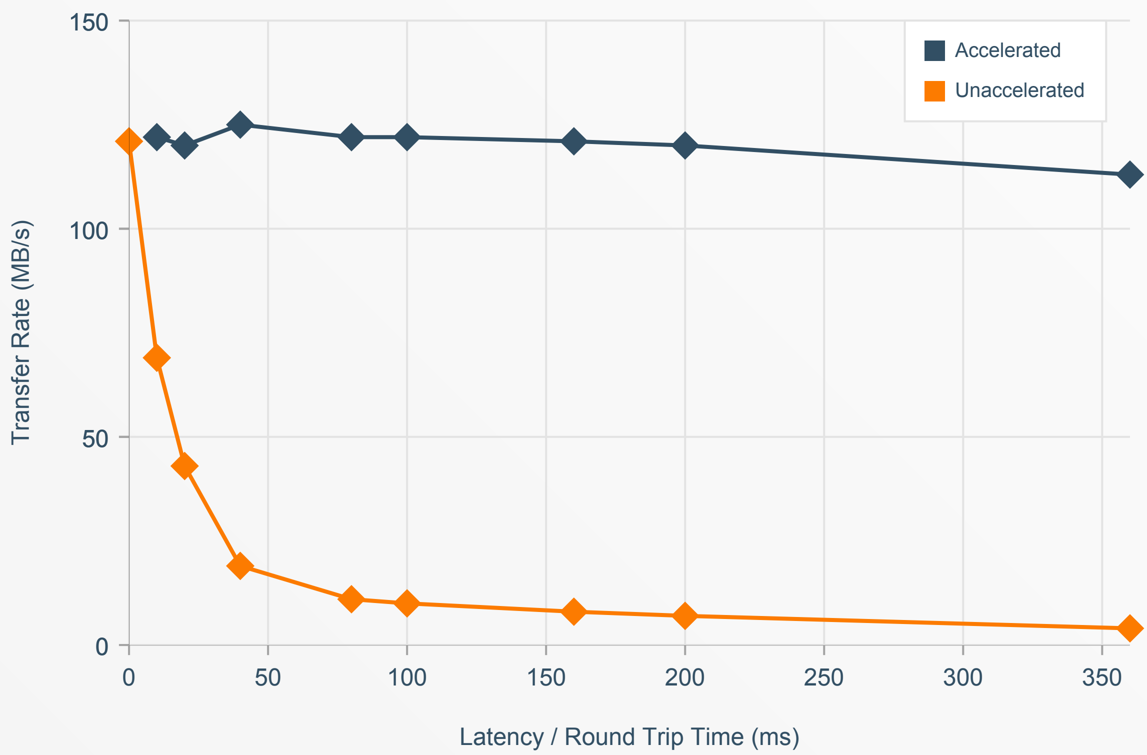

Figure 6: Accelerated and unaccelerated performance with 0.5% packet loss at various levels of latency

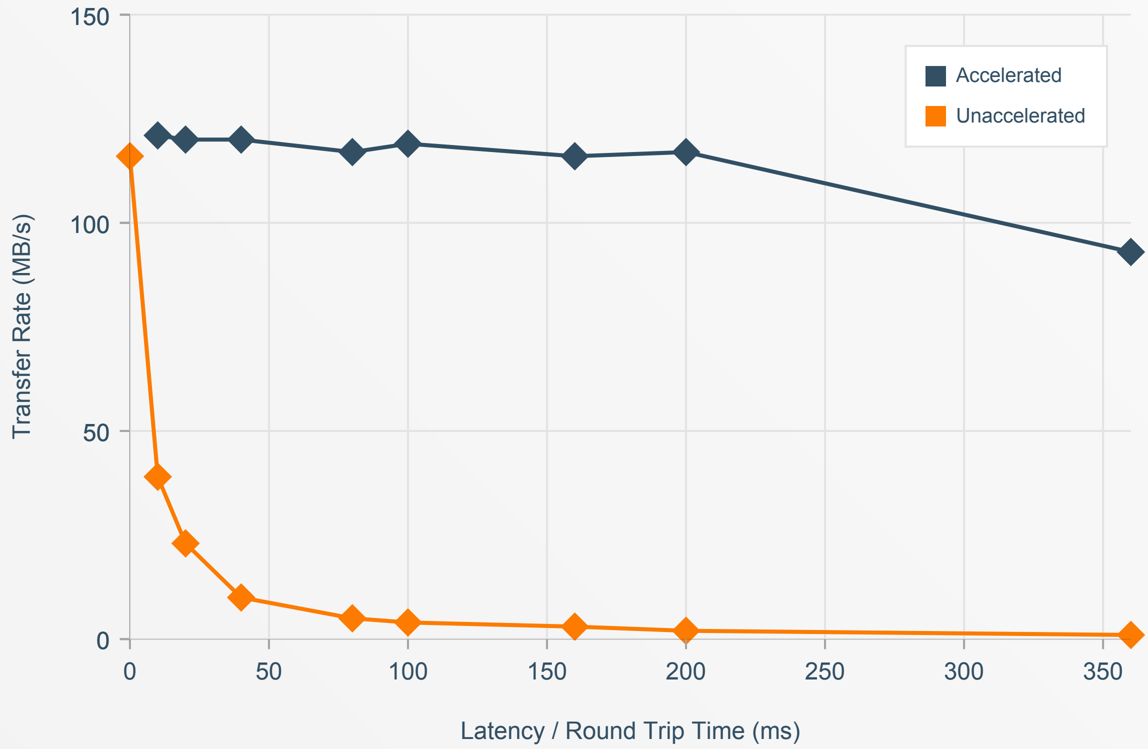

Figure 7: Accelerated and unaccelerated performance with 1% packet loss at various levels of latency

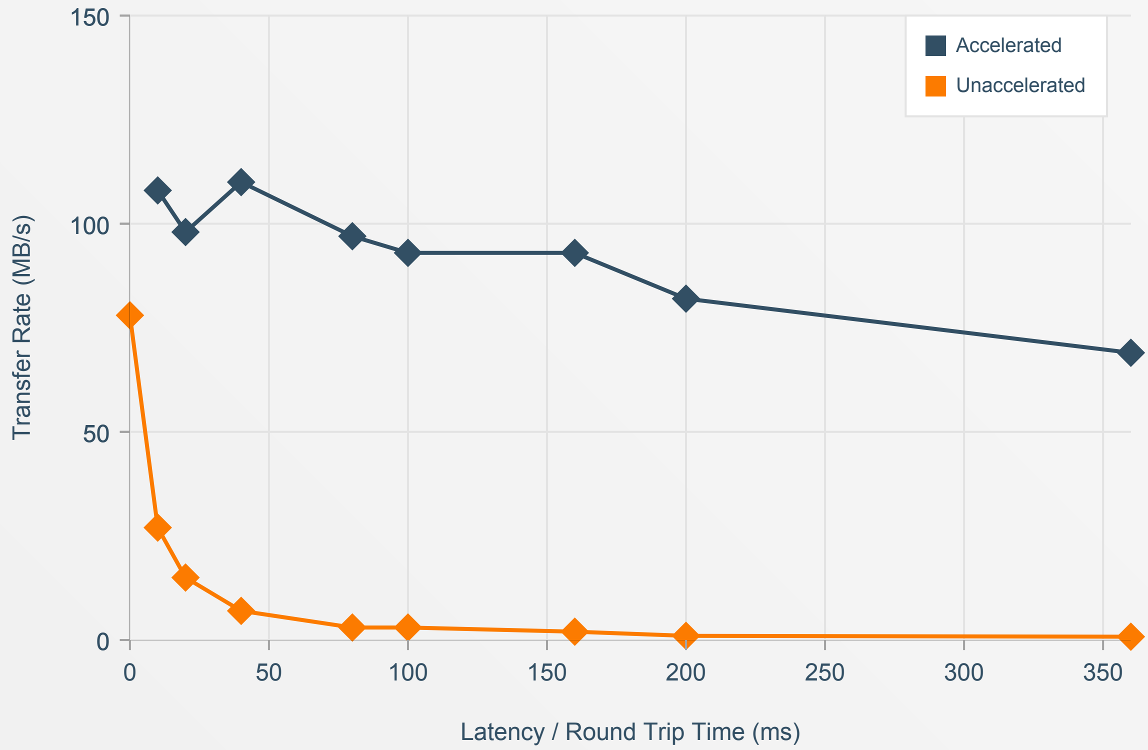

REALISING THE BUSINESS BENEFITS

For companies that use Veeam for VM replication over a WAN, PORTrockIT can transform performance. If the replication process is threatening to overrun the available window, or if it is desirable to reduce replication times 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 source and target servers – enabling significant cost-avoidance.

Most important of all, PORTrockIT removes a key risk from business continuity strategies that depend on rapid failover to a replicated set of VMs at a remote data centre, by helping to ensure that all VMs can be replicated quickly, efficiently, reliably and safely.

Increasing the speed of the replication process can also enable companies to shrink their replication windows, which reduces the amount of time that other systems need to spend contending with replication traffic for limited network bandwidth. By reducing the impact of replication on other systems, PORTrockIT could also make it feasible for companies to perform replication more frequently, helping them optimise their recovery point objectives.

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