CARINGO SWARM

Accelerating file transfer into a massively scalable object storage cluster with **PORTrockIT**

BRIDGEWORKS

EXECUTIVE SUMMARY

As "big data" continues its evolution from a tech-industry buzzword to an everyday business reality, more and more organisations will need to find practical ways to deal with data that is growing at an unprecedented rate.

Object storage technologies such as Caringo Swarm can provide an answer, using clustering technologies and commodity hardware to scale storage environments into the petabytes, without adding complexity.

However, even the smartest, most cutting-edge storage technology is of limited value unless you can get your data into it quickly enough for the business to use it. In situations where large amounts of data need to be ingested into a Caringo Swarm cluster across a wide area network (WAN), the network transfer rate is critical. If new data is being generated faster than it can be transmitted to the cluster, backlogs will build up. As a result, it may become impossible to deliver the kind of rapid, real-time analysis that turns big data into business value.

This paper shows how PORTrockIT can significantly improve transfer rates across a WAN by counteracting the effects of latency and packet loss – the two main causes of performance issues on real-world networks. Even in the most challenging scenario we tested, with 360 ms of latency and a 1% packet loss rate, PORTrockIT was able to move data across a WAN into a Caringo Swarm cluster more than 337 times faster than could be achieved with a traditional network architecture.

> *"PORTrockIT was able to move data across a WAN into a Caringo Swarm cluster* ***more than 337 times faster*** *than a traditional network architecture."*

By simply installing a pair of PORTrockIT either as an appliance or as a virtual instance at either end of your WAN, you can instantly accelerate your Caringo Swarm architecture – avoiding any need to invest in faster servers or higher-bandwidth network infrastructure. This will help you gain the true benefits of massively scalable object storage, at a fraction of the cost.

WHY SPEED MATTERS

Once considered a niche technology, object storage has recently seen a surge in popularity, especially because it offers a highly scalable way to manage very large volumes of data without adding ever-increasing numbers of file servers.

As one of the leading innovators in the object storage sector, Caringo can claim to have more customers than all the other object storage software vendors combined. Its Caringo Swarm technology offers a cluster-based approach that can scale to hundreds of billions of files and petabytes of data, without the complexity and performance issues that legacy storage infrastructures tend to develop when operating at scale.

However, since Caringo Swarm removes many of the traditional bottlenecks within the storage environment itself, the performance of the network infrastructure can become the next choke-point. To take full advantage of the speed and flexibility of Swarm, you also need to be able to move data into and out of the cluster at high speed, even across an extensive WAN.

This is particularly important for big data analytics use cases, where the successful delivery of near-real-time analysis can depend on your ability to ingest large numbers of new files into a storage environment quickly. If the network can't keep up with the data that is being generated, the main advantage of having a massively scalable storage cluster is negated. Therefore, maintaining a high transfer rate across the WAN is a critical enabler for obtaining the maximum benefits from Caringo Swarm.

To address network performance issues and get the full value from a Caringo Swarm solution, we first need to answer a key question:

> *"What are the **causes** of slow WAN transfer rates?"*

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 wide area network (WAN) – whether into a Caringo Swarm cluster, or onto any other kind of storage platform.

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 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, 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.

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

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 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 **much faster network transfer performance**, without making any changes to the rest of the network architecture."*

In practical terms, PORTrockIT is installed as a pair of appliances, deployed at either end of the WAN. A host 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 Caringo Swarm cluster.

TURNING THEORY INTO PRACTICE

To demonstrate the kind of results that PORTrockIT can deliver for Caringo Swarm customers, Bridgeworks conducted a set of performance tests at an independent testing facility in the UK. The test infrastructures mimicked a real-world Caringo Swarm architecture, using a WANulator to simulate different levels of latency and packet loss between a host (the system that is sending data to be ingested by Caringo Swarm), and a three-node Caringo Swarm cluster.

The first set of tests were performed on an unaccelerated architecture, where the host and the Caringo Swarm cluster 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, placed on either side of the WANulator, between the host and the cluster (see figure 2).

TEST EQUIPMENT

SOFTWARE:

- Caringo Swarm v.7.5.4
- Caringo FileJet v.1.5

HARDWARE:

Host
- Dell R710 with 2 x Intel XEON E5506 processors @ 2.13GHz, 8GB RAM and 1 x 240GB SSD
- Windows 2012 R2

3 x Caringo Swarm Server nodes
- 1Gb switch SMC 8508T
- 1 x Sun X2250 with 2 x Intel XEON X5472 processors @ 3.0GHz, 16GB RAM and 1 x 240GB SSD
- 2 x Dell R720 with Intel XEON E5 2603 processor @ 1.8GHz, 16GB RAM and 1 x 240GB SSD

2 x PORTrockIT nodes
- Dell R310

WANulator
- Dell R210

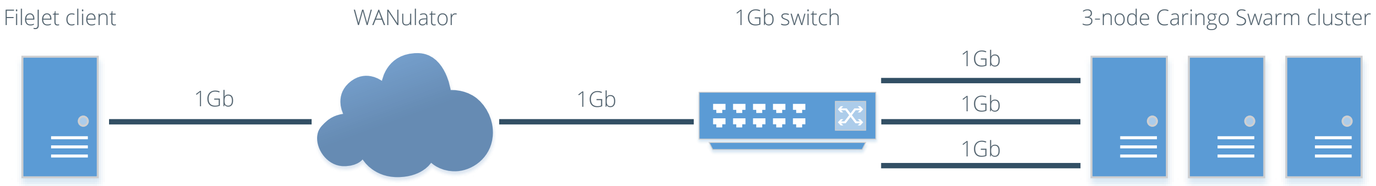

Figure 1: Unaccelerated architecture

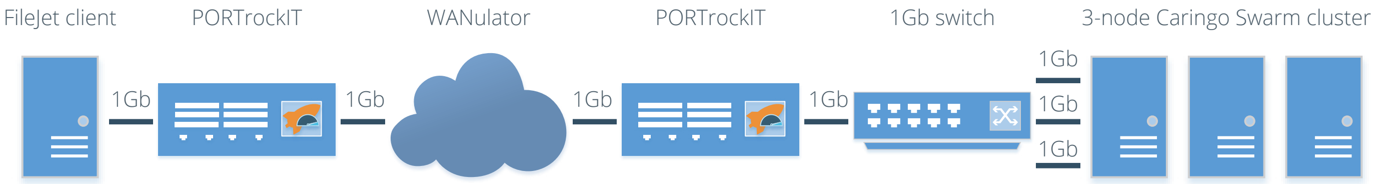

Figure 2: Accelerated architecture with PORTrockIT

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). The Caringo FileJet utility was used to transfer a 1 GB file from the host to the Caringo Swarm cluster, first via the unaccelerated architecture, and then again via the accelerated architecture with PORTrockIT.

From Figure 3, we can see that even small amounts of latency had a significant negative impact on Caringo Swarm's ability to ingest data via a traditional unaccelerated WAN architecture. Moving from 0 ms to 10 ms of latency reduced the transfer rate from 87 MB/s to just 11.4 MB/s – about one-eighth of the speed.

As latency increased further, transfer rates tailed off to almost nothing: once the round trip time exceeded 100 ms, the transfer rate fell below 1 MB/s. At a round trip time of 360 ms, the transfer rate was just 0.38 MB/s. To put this in context: ingesting even a single 1 GB file into a Caringo Swarm cluster at a speed of 0.38 MB/s would take more than 43 minutes.

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

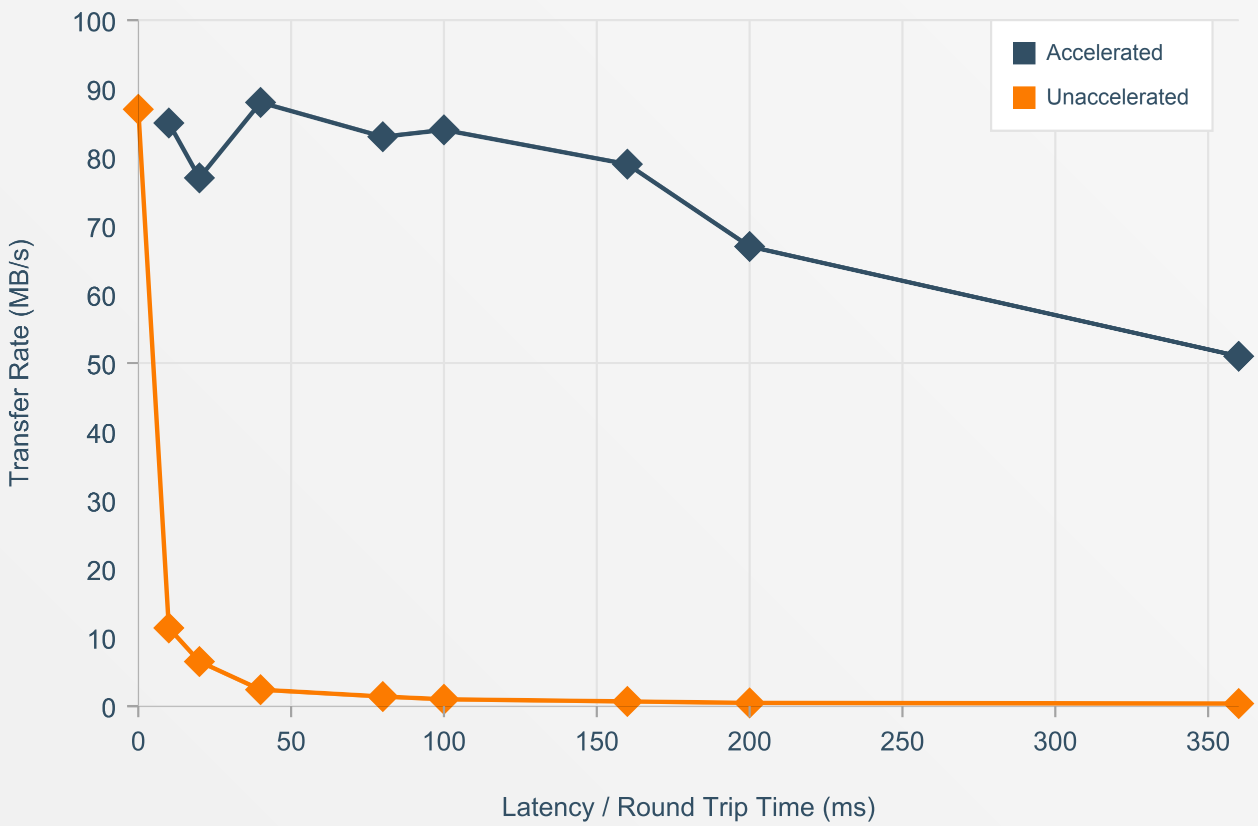

By contrast, the accelerated architecture with PORTrockIT provided much higher performance at all levels of latency. Although transfer rates did begin to dip as the round trip time increased, they remained above 75 MB/s until latency reached 160 ms, and above 50 MB/s in all cases.

In the toughest scenario, with 360 ms of latency, the PORTrockIT architecture was 134 times faster than the unaccelerated architecture. Even with this challenging network configuration, it would still be possible to ingest a 1 GB file into a Caringo Swarm cluster in less than 20 seconds.

"In the toughest scenario, with 360ms of latency, the PORTrockIT architecture was ***134 times faster*** *than the unaccelerated architecture."*

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

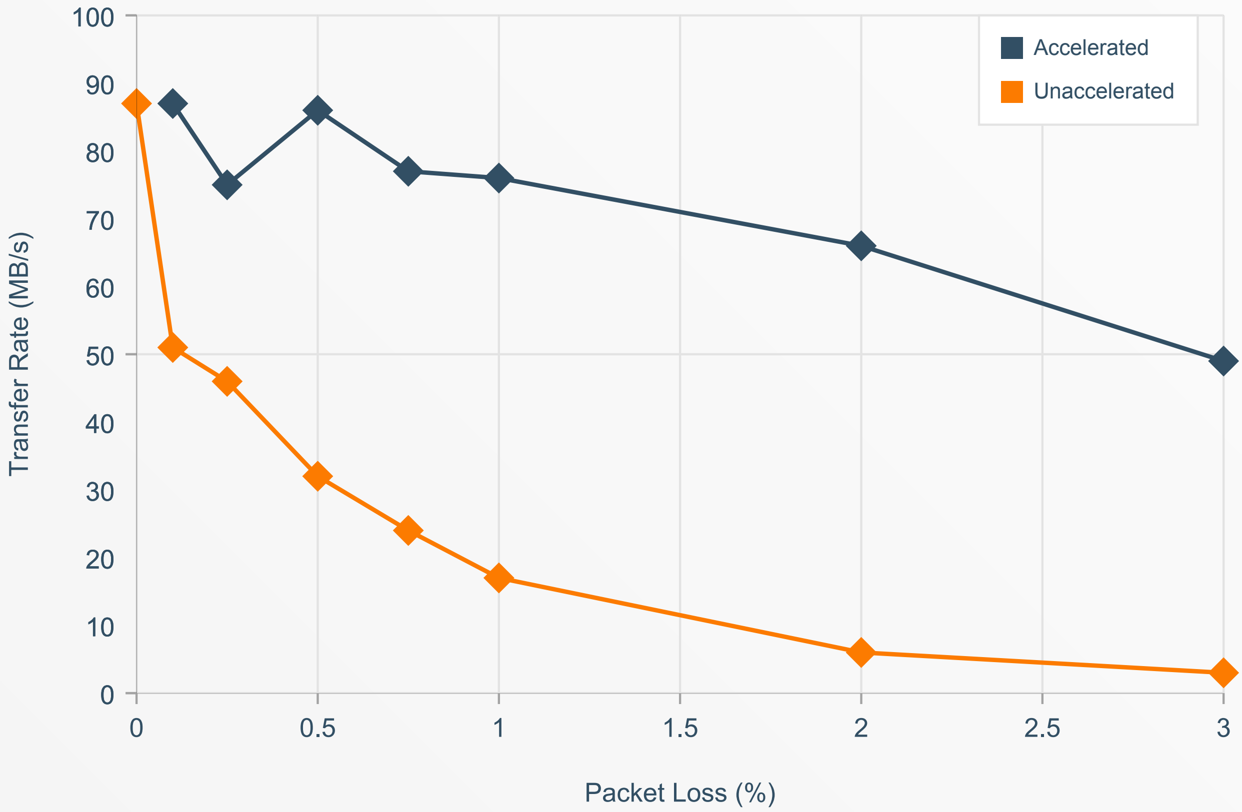

PACKET LOSS

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 1 GB file was transferred from the host to the cluster using both the unaccelerated and accelerated architectures.

Figure 4 shows that for both architectures, performance degrades as packet loss increases – but in all cases, the transfer rate is considerably higher with the accelerated architecture.

Even in the extreme case of a network with 3% packet loss, the accelerated architecture still delivers a transfer rate of 49 MB/s – more than 16 times faster than the 3 MB/s rate achieved by the unaccelerated architecture.

> *"In the extreme case of a network with 3% packet loss, the accelerated architecture is still* ***more than 16 times faster*** *than the unaccelerated architecture. "*

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 impact of the combined effects of both problems. All wide area networks are subject to at least some degree of both latency and packet loss, so these final scenarios give an indication of how Caringo Swarm might perform with and without PORTrockIT acceleration across a real-world WAN.

Figures 5, 6 and 7 all show that the unaccelerated architecture saw severe performance degradation from the combination of latency and packet loss. With a packet loss rate of 0.1%, the best transfer rate achieved was just 51 MB/s, and the rate fell below 1 MB/s once latency hit 100 ms. When packet loss rose to 0.5%, the best rate achieved was 30 MB/s, and the rate fell below 1 MB/s before 80 ms of latency was reached. Finally, with 1% packet loss, the highest rate achieved was just 17 MB/s, and performance fell below 1 MB/s when latency rose above 40 ms.

In all three charts, we can also see that the accelerated architecture delivered considerably better performance at all levels of latency. Figures 5 and 6 show a slight downward trend as latency increases. However, in both cases, transfer rates remain above 70 MB/s until round trip time hits 200 ms, and above 59 MB/s in all cases.

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

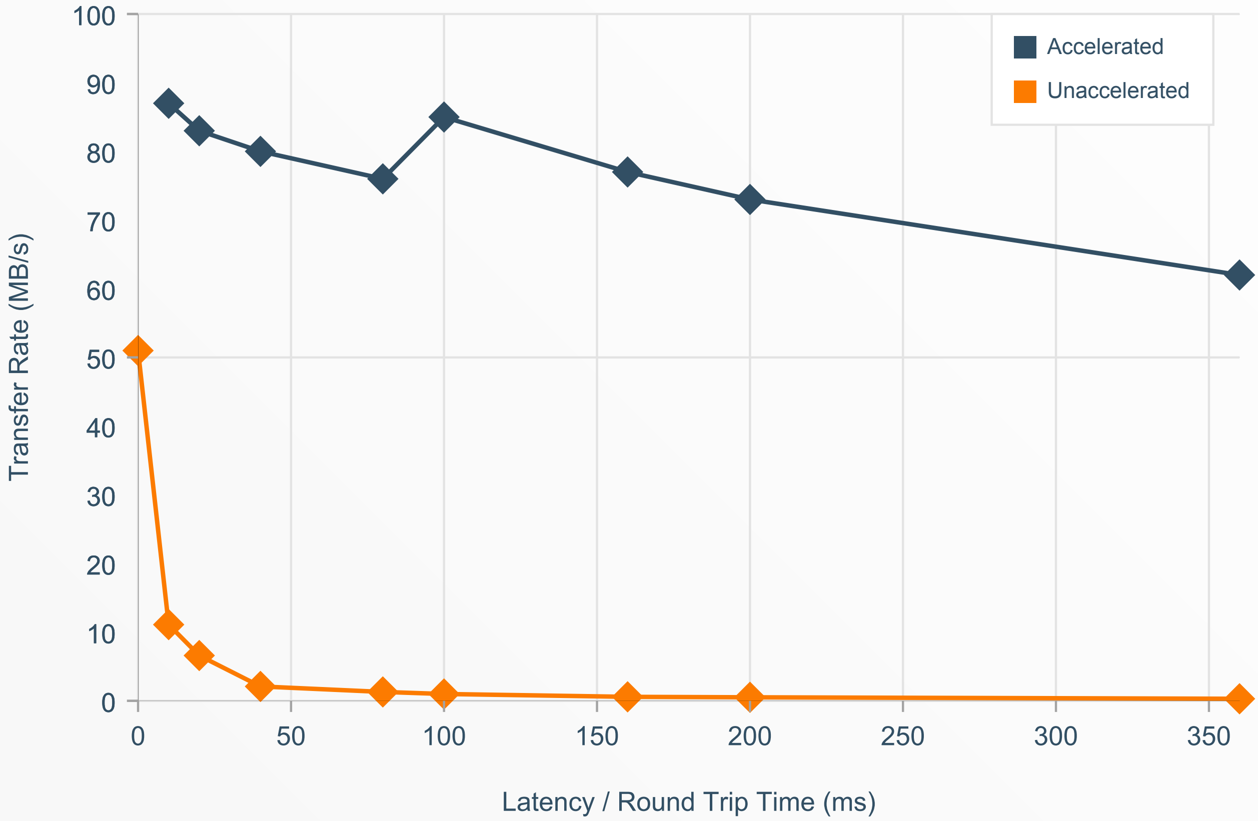

Figure 6: Accelerated and unaccelerated performance with 0.5% packet loss at various levels of latency

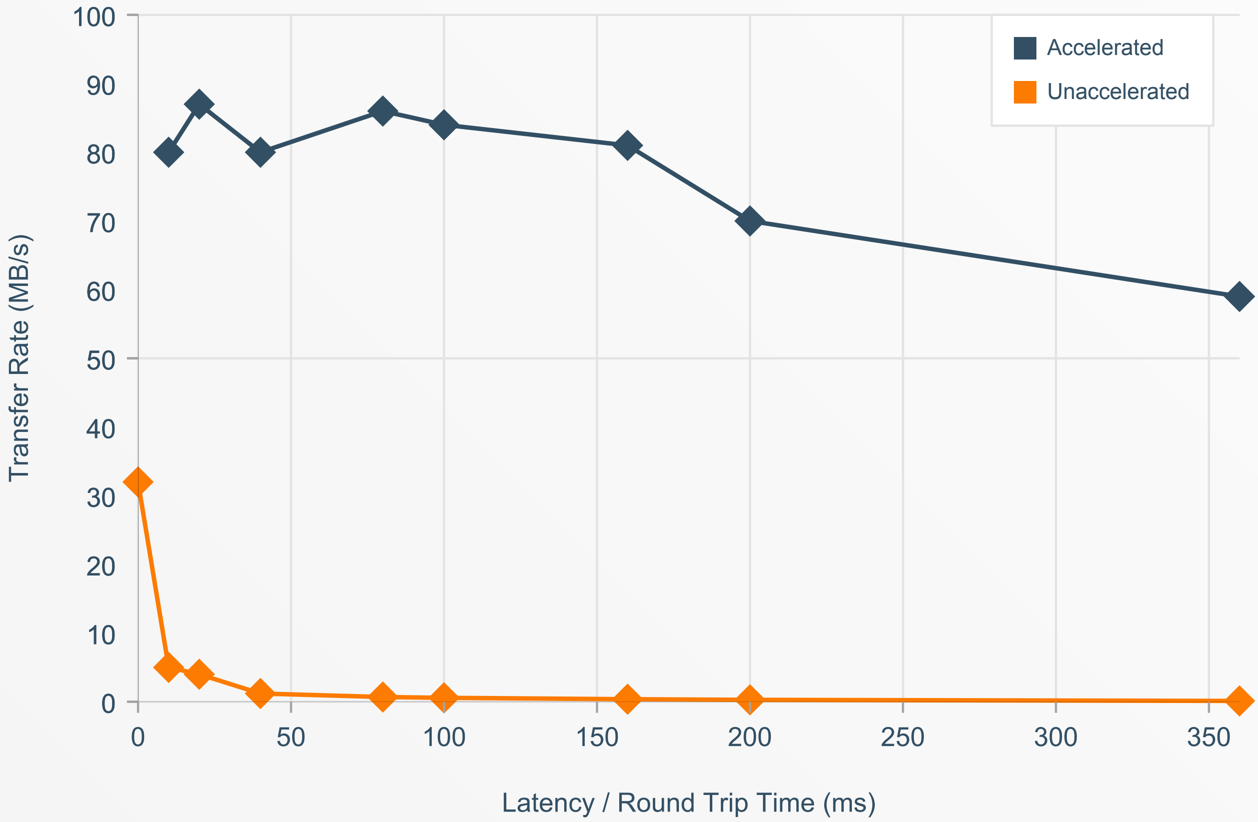

Figure 7 suggests that once the packet loss rate reaches 1%, PORTrockIT's performance does begin to degrade – but only at the highest levels of latency. When round trip times hit 200 ms or higher, there is a significant reduction in performance; but at lower levels of latency, the performance remains good – above 70 MB/s in all cases.

Moreover, even at the highest latency levels, the PORTrockIT architecture still performs considerably better than the unaccelerated infrastructure. Even in the most challenging scenario (360 ms of latency with 1% packet loss), the accelerated architecture achieved a transfer rate of 29 MB/s.

This is more than 337 times faster than the unaccelerated transfer rate of 0.086 MB/s.

> *"Even in the most challenging scenario, PORTrockIT achieved a transfer rate* ***more than 337 times faster*** *than an unaccelerated architecture."*

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

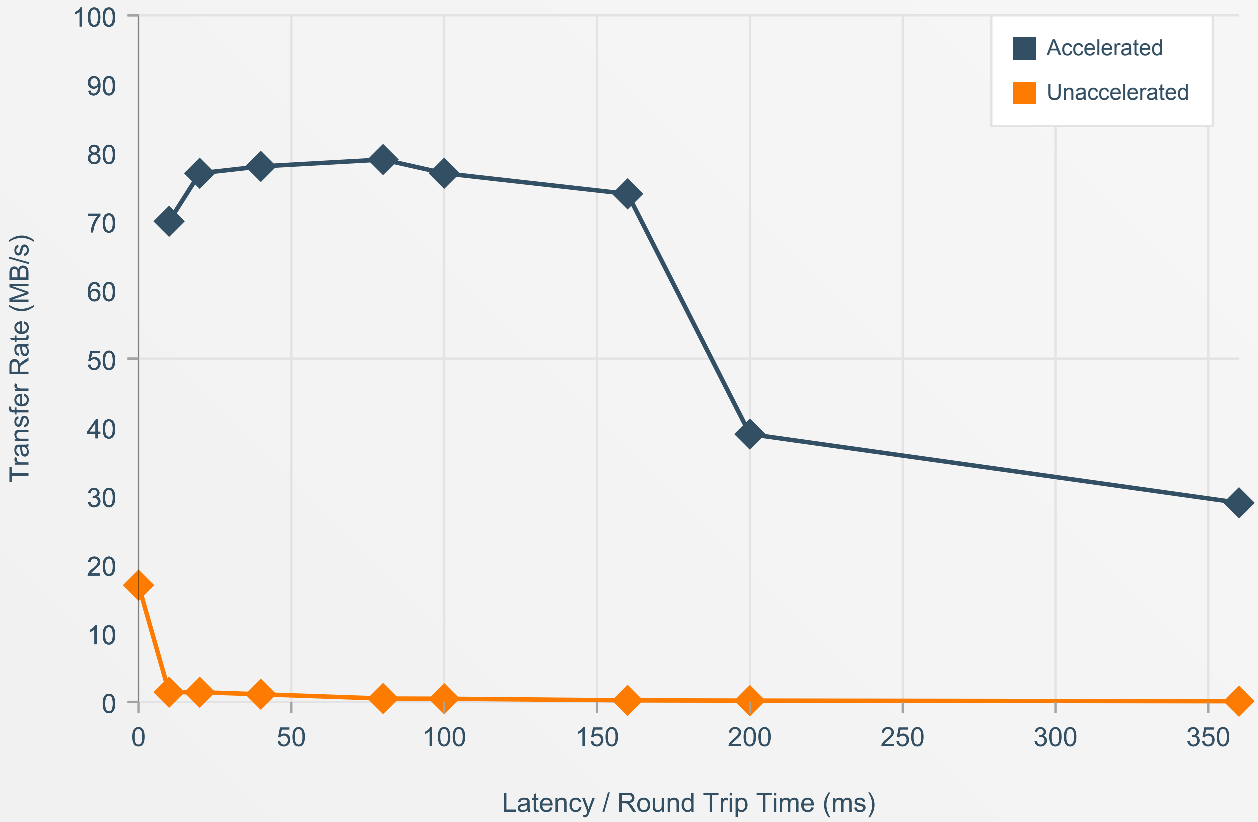

REALISING THE BUSINESS BENEFITS

By substantially mitigating the impact of latency and packet loss on WAN transfer performance, PORTrockIT can help to unlock the full benefits of Caringo Swarm.

Instead of creating a high-performance large-scale storage environment and then finding that network performance becomes the new bottleneck, organisations that use PORTrockIT can expect their file transfer processes to keep pace with their Caringo Swarm storage landscape.

Furthermore, PORTrockIT offers plug-in-and-go technology that can be implemented quickly with minimal impact on the rest of the 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 host servers – enabling significant cost-avoidance.

As organisations seek new ways to capture, manage, process and analyse ever-increasing volumes of data, Caringo Swarm and PORTrockIT provide an ideal combination for massively scalable file storage that offers both speed and cost-efficiency.

"By maximising the performance of existing infrastructure, PORTrockIT enables ***significant cost-avoidance****."*

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, please visit www.4bridgeworks.com, or call us on +44 (0) 1590 615 444.

BRIDGEWORKS