OFF-SITE TAPE REPLICATION

Defining a new paradigm with WANrockIT

BRIDGEWORKS

EXECUTIVE SUMMARY

For organisations that need to move large quantities of data onto tape storage and store it off-site, the prospect of transferring the data electronically, rather than physically shipping the tapes to a disaster recovery (DR) site, is an attractive one. However, the technical challenges of copying large quantities of data to tape across a wide area network (WAN) are formidable.

This paper shows how Bridgeworks' WANrockIT solution can overcome the two major obstacles that prevent most organisations from adopting WAN tape replication.

First, by combatting network latency, WANrockIT can accelerate data transfer rates by a factor of more than 200 – enabling even very large transfers to be completed within a realistic backup window.

Second, by converting between transfer protocols on-the-fly, the solution also enables organisations to use low-cost SAS tape drives at their DR site, while still retaining their Fibre Channel or iSCSI SAN at their primary site.

WANrockIT can help your organisation turn WAN tape replication into a reality – eliminating the effort, cost and risk of shipping tapes to a DR site, without requiring costly upgrades to your network and tape infrastructure.

*"By combatting network latency, WANrockIT can accelerate data transfer rates by a factor of **more than 200**."*

WHY SPEED MATTERS

Most organisations would agree that even as the price of disk storage continues to fall, tape storage is still the safest and most cost-effective way to preserve large quantities of data for backups or archiving purposes. Tape storage media is far cheaper per gigabyte than disk storage, it has no ongoing electricity costs, and tapes are far less likely to fail than spinning disks.

However, creating an efficient tape storage infrastructure does present some challenges. To protect your data in the event of a disaster at your main data centre, the standard best practice is to store at least one copy of each tape at a DR site. But how do you get your data to this second location?

The traditional approach is to copy the data to tape using a tape drive or tape library at the main data centre, and then physically ship the tapes to the DR site. This is a slow, expensive and labour-intensive process – and if a tape goes astray during transit, there is a risk that valuable data could be lost or even stolen.

An alternative option is to transfer the data to the second site electronically, via a WAN, and copy it to tape once it gets there. This is a much more elegant and secure solution, but it is technically much more difficult to accomplish.

> *"The **technical challenges** of copying large quantities of data to tape across a wide area network (WAN) are formidable."*

The first problem is that most affordable tape drives and libraries do not support transfer protocols that are suitable for WAN transfer. For example, the maximum recommended distance between a host and a tape drive connected via the Serial-Attached SCSI (SAS) protocol is just 12 metres.

The alternative is to invest in a Dark Fibre connection between the two sites and purchase a Fibre Channel-capable tape library, but this can be a prohibitively expensive option for all but the largest enterprises.

The second – and more serious – problem is the data transfer rate. Moving large quantities of data from a server onto a tape drive is relatively straightforward within a single data centre; but when we convert the Fibre Channel protocol to an IP-based protocol in order to move data over larger distances across a WAN, it can be a different story. If the WAN connection cannot sustain a sufficiently high transfer rate, the replication process may take many hours, and overrun the available window that the organisation has allocated for backups and archiving.

If this happens, the organisation needs to make a difficult choice: either it waits for the transfer to complete, which will potentially have a severe impact on the performance of other systems that rely on the WAN, or it cancels the transfer and potentially leaves its data unprotected.

Bridgeworks' WANrockIT solution can solve both the protocol problem and the transfer rate problem – but before we explain how, we need to understand why WAN transfer rates can be so slow.

THE PROBLEM: LATENCY

The single biggest issue when moving data across any kind of WAN is network 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 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. TCP/IP, the main protocol used for IP data transfer, is very susceptible to latency, which 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.

Organisations often try to solve TCP/IP performance issues by investing in more expensive network technology (for example, an end-to-end fibre channel infrastructure) that offers a larger maximum bandwidth.

> *"Investment in bandwidth will simply be **wasted** unless latency issues can be addressed."*

However, this does not fix the problem. As we have seen, latency prevents TCP/IP connections from fully utilising the available bandwidth – so any extra investment in bandwidth will simply be wasted unless the latency issues can be addressed.

THE SOLUTION: WANROCKIT

WANrockIT offers a complete solution to network latency issues, and can also convert between different transfer protocols if required.

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, WANrockIT 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.

Moreover, WANrockIT 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 WANrockIT – enabling the appliance to operate optimally at all times, without any need for input from a network administrator.

> *"WANrockIT combines **AI-powered data-flow optimisation** with real-time translation between protocols."*

WANrockIT also provides bridging technology that can translate data between SCSI, iSCSI, fibre channel and SAS protocols on-the-fly. In our tape replication scenario, this means it is possible to use iSCSI to transfer the data across the WAN; once the data has arrived at the DR site, it can then be converted to SAS and sent to the tape drive.

In practical terms, the WANrockIT solution comprises a pair of appliances, deployed at either end of the WAN. A host server simply passes data to the WANrockIT appliance on the near side of the WAN, which manages the virtual connections to the second WANrockIT appliance on the far side of the WAN. Once the second WANrockIT appliance begins receiving packets, it converts the protocol to SAS, and passes the data on to the tape drive.

The result is 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 WANrockIT can deliver for organisations that want to accelerate their off-site tape replication process, Bridgeworks conducted a set of performance tests at an independent testing facility in the UK.

The test infrastructures mimicked a real-world architecture, using a WANulator to simulate different levels of WAN latency between a host server and a tape drive.

To manage the replication itself, the Bridgeworks team chose Veritas NetBackup – one of the most widely used backup products on the market – as a representative solution.

Figure 1: Unaccelerated architecture

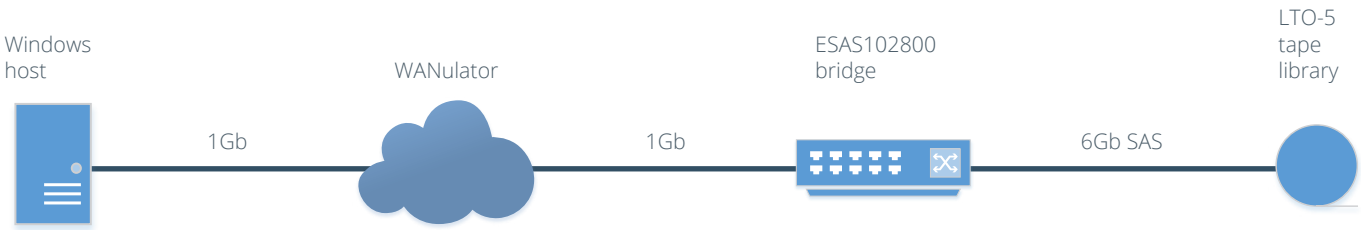

The first set of tests were performed on an unaccelerated architecture, where the host was connected directly to the WANulator, and the tape drive was connected to the WANulator via the bridge (see figure 1). The bridge in the first test case is used to convert the iSCSI protocol to an SAS protocol.

The same tests were then repeated on an architecture that was accelerated by introducing the full WANrockIT solution, with two WANrockIT appliances placed on either side of the WANulator between the host and the bridge (see figure 2).

TEST EQUIPMENT

HEADING:

SOFTWARE:

- Veritas NetBackup v.7.6.0.3

HARDWARE:

Host
- IBM x3250, 4Gb RAM, Intel Xeon E31230 3.2GHz
- Windows 2012 R2

Bridgeworks WANrockIT Series 200
- 2 x WANrockIT nodes

Other
- 1 x ESAS102800 bridge
- IBM Ultrium HH LTO-5 SAS tape drive
- WANulator

Figure 2: Accelerated architecture with WANrockIT

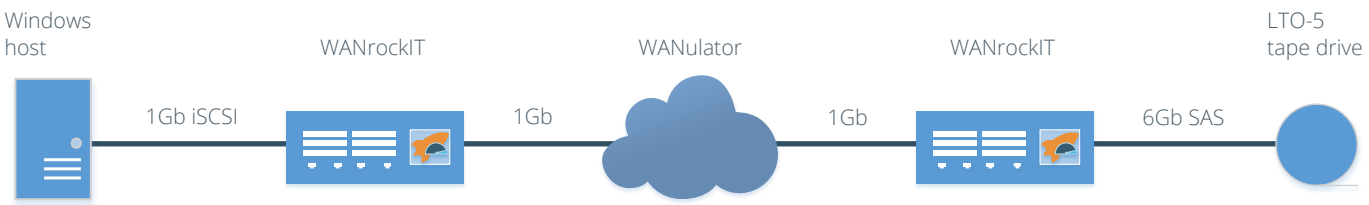

WHAT THE DATA TELLS US

The tests simulated a tape replication scenario with latencies ranging from 10 ms to 360 ms round trip time (RTT). Veritas NetBackup was used to transfer a file from the host to the tape library, first via the unaccelerated architecture, and then again via the accelerated architecture with WANrockIT.

From Figure 3, we can see that even at relatively low levels of latency, the unaccelerated architecture performed poorly. With an RTT of 10ms, the average speed was 15.9 MB/s – which means that it would take more than 62 seconds to copy just 1 GB of data to tape across the WAN.

As latency increased further, transfer rates tailed off to almost nothing: once the round trip time exceeded 160 ms, the transfer rate fell below 1 MB/s. At a round trip time of 360 ms, the transfer rate was just 0.4 MB/s. To put this in context: transferring 1 GB of data at a speed of 0.4 MB/s would take more than 41 minutes.

By contrast, the accelerated architecture with WANrockIT provided much higher performance at all levels of latency, and the transfer rate remained more or less stable even as the round trip time increased.

While a slight downward trend can be observed in the graph, the decrease in performance between 10 ms and 360 ms

Figure 3: Accelerated and unaccelerated performance at various latencies

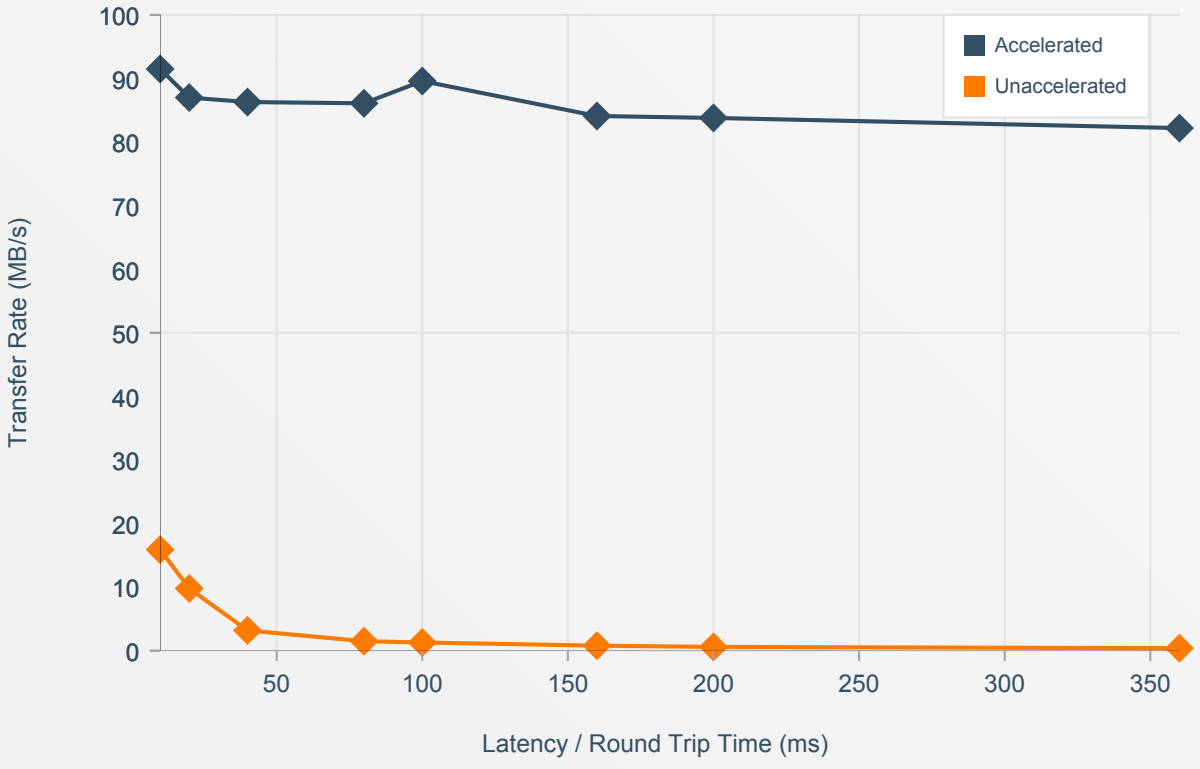

*"WANrockIT performed significantly better at all levels of latency. In the toughest scenario, it was **205 times faster**."*

of latency is less than 10 MB/s, and in all cases, the transfer rate remains above 80 MB/s.

In the toughest scenario, with 360 ms of latency, the WANrockIT architecture still achieved a transfer rate of 82.2 MB/s – 205.5 times faster than the unaccelerated architecture. Even with this challenging network configuration, it would still be possible to copy a 1 GB file to tape across the WAN in less than 13 seconds.

REALISING THE BUSINESS BENEFITS

By significantly accelerating data transfer rates and enabling the use of low-cost SAS-attached tape hardware across the WAN, WANrockIT makes off-site tape replication into a realistic and attractive proposition for organisations of all sizes. Investing in WANrockIT will eliminate the need to physically move tapes to a DR site, reducing operational costs and avoiding the risk of tapes being damaged, lost or stolen in transit.

WANrockIT 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, WANrockIT also reduces the need to invest in expensive high-bandwidth network infrastructure, fibre channel-capable tape libraries, or more powerful host servers – enabling further 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 and other smart networking solutions from Bridgeworks, please visit www.4bridgeworks.com, or call us on +44 (0) 1590 615 444.

Copyright Bridgeworks 2018 / www.4bridgeworks.com / +44 (0) 1590 615 444