

WAN-optimized Heterogeneous Replication

Revolutionize Your IT Efficiency

Abstract: Data replication has become part of the fabric of IT due to its multiple uses and benefits. However, as IT and business managers are well aware, data replication can quickly increase your bandwidth and storage costs. WAN optimization technologies can help reduce the amount of bandwidth and storage capacity used by data replication, enabling cost savings while maintaining application performance.

However, in most cases, WAN optimization requires the purchase and installation of expensive hardware. FalconStor solutions, on the other hand, include built-in advanced data replication technology that is many times more WAN-efficient than other types of solutions, providing significant savings in bandwidth, storage capacity, and costs.

The Many Uses of Data Replication

Data replication has become increasingly popular in small and large organizations, primarily for minimizing downtime associated with backup and disaster recovery (DR). Replication also enables organizations to use data in more than one way at a time. For example, if you replicate a database, you can continue using the original data set for production operations while using the copy to load a data warehouse or perform application testing. As a result, you get more value for your investment in information creation and management.

Remote replication can improve the backup process and shrink the backup window to virtually nothing. Many organizations take snapshots of their data, replicate it to a remote location, and complete the backup-to-tape process from the remote site. The primary site can immediately return to normal business functions, reducing downtime and enabling the speedy resumption of production (and revenue-generating) operations. With data replication, backups are no longer a disruptive process. Similarly, remotely replicated data can be used for DR to improve data protection. With data replicated offsite, it is protected from local corruptions and failures. Many organizations count on their replication technologies to deliver the data protection required for corporate and regulatory compliance.

Data replication is also extremely useful for keeping up with maintenance schedules. Some organizations run the risk of ignoring important system maintenance because it is disruptive to take down critical applications such as email. By replicating your email systems, you can bring down your primary system, start up operations on the secondary system, and then install patches, hardware upgrades, and other maintenance tasks on the primary system before bringing it back online. Without replication, these systems would need to be inoperable for long periods of time, interrupting business operations. Data replication is equally helpful for data center moves. You can fail over your systems to a DR site, move your production data center to its new location, and then fail back to the new data center. This type of task is virtually impossible — or at least highly impractical — without replicating over the WAN.

With data volumes growing significantly in remote and branch offices, replication provides another benefit. Since most branch offices lack adequate IT staff and the skills to maintain and protect data, branch office data is often highly vulnerable to loss and corruption. With remote replication, branch office data can be copied over the WAN to a central repository or data center, where IT can manage it according to corporate policies for security, protection, etc.

Clogging the Pipes

Data replication offers business benefits such as improved protection and multiple simultaneous uses of information, to help minimize total cost of ownership (TCO). But those growing data volumes being replicated need more storage to retain them, and their bulk can clog the WAN. Plus, in today's business environment, there are many new applications vying for WAN bandwidth — not only backup and DR, but file sharing between locations, Web 2.0 applications such as YouTube and Salesforce.com, video files, training applications, etc. Furthermore, bandwidth is expensive — few organizations can afford an unlimited supply of it.

Increasing WAN traffic can slow down application performance — a situation for which users have little tolerance. As a result, the only answer for many is to perform expensive WAN upgrades to handle the increased network traffic and maintain adequate application performance.

With so many applications competing for bandwidth, and storage capacity being costly, any way to reduce the amount of data passed across the WAN and stored remotely is helpful. There are several solutions on the market to accelerate WAN bandwidth, but most require additional equipment. The question for users is: Do you want to spend a lot of money implementing a complex solution to solve the problem, or do you want to solve the problem simply and easily with software that already includes WAN optimization?

WAN Resiliency

The nature of IP transport often makes replication problematic. Even the best networks can get noisy or busy, packets can drop, and data has to be retransmitted. Many replication products have difficulty dealing with transport inconsistencies.

FalconStor Software provides proven data protection solutions with built-in quality of service (QoS) tuning to handle network disruptions. FalconStor replication can automatically switch from continuous to periodic replication when transport quality falls below a configured level. Network throttling is also available to ensure that bandwidth is adequately allocated to the replication traffic if the link is shared with other applications without allocating more bandwidth than is necessary.

FalconStor technology is also able to gracefully resume following a service outage. For example, if you are in the middle of replicating 100GB of accumulated data and the WAN connection is lost near the end of the transmission, many replication tools will re-start at the beginning, resending the full 100GB back over the wire, while falling farther behind in light of the constant accumulation of new data. In extreme cases, systems can fall so far behind that they never catch up. With FalconStor replication technology, instead of resending the full 100GB, only the remaining data is sent, allowing the DR process to resume quickly.

WAN-optimized FalconStor Solutions

FalconStor solutions such as FalconStor® Network Storage Server (NSS) and FalconStor Continuous Data Protector (CDP) include advanced replication technology that minimizes the amount of data sent over the WAN as well as the amount of storage capacity required. There is nothing additional to purchase because FalconStor MicroScan™ technology is built into these solutions. This patented data deduplication mechanism goes beyond other technologies by eliminating redundant, block-level changes that occur due to application and file system inefficiencies.

With most solutions, changing even a single bit of data requires a cluster of kilobytes to be written over the wire — but the bulk of this is redundant data (re-writes) or white space. As a result, when replicating changed data, these solutions will transfer more data than was actually changed. This wastes storage capacity, clogs

the network, and creates unnecessary costs. Imagine that you want to make a copy of a graph that accompanies four paragraphs of text — you'll save paper if you copy only the graph and not the text around it.

That's the difference that FalconStor replication offers. Only the actual changes at the disk-sector level are transferred for replication and storage on the target side. The technology can work with any standard storage array or data source. Regardless of which files or databases your organization wants to replicate, FalconStor solutions will transfer the least amount of data. Compared to other replication methods, users see bandwidth reductions of 80-90% or even more, depending on the characteristics of their data.

The TOTALLY Open™ nature of FalconStor replication technology eliminates vendor lock-in. Array-based systems typically require the same disk array at both sites. All disk resources and capacity expansion must come from the same vendor, limiting corporate buying options and price negotiating. Because FalconStor crosses vendor barriers, lower-cost storage can be used at the DR site in order to meet budget constraints or to extend DR protection to more applications.

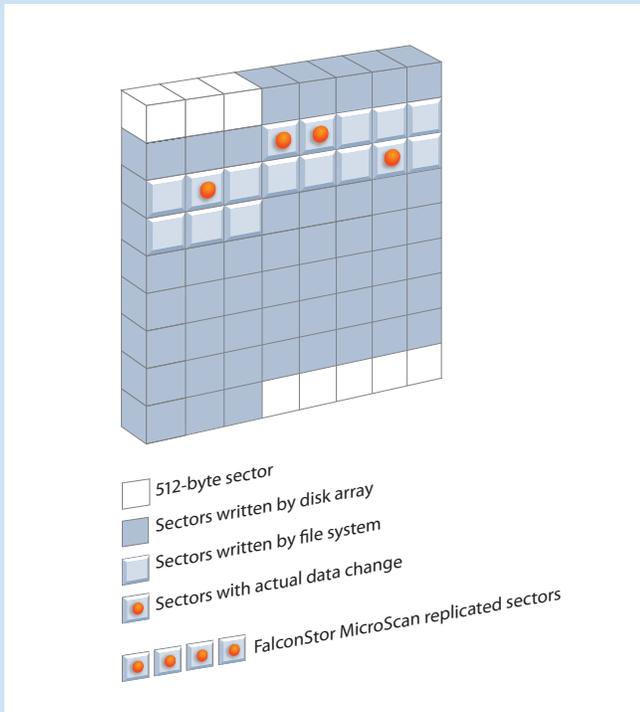
The following graphic illustrates the difference between FalconStor data replication and storage array-based solutions. Each square represents a 512-byte sector on a disk array. This is the standard size for most disk systems and represents the smallest possible segment of data written to a disk. When an actual data change is made, the application needs to write one or more 512-byte sectors. In the illustration, four sectors are actually updated. However, file systems work in larger segments, so the file system may write an 8K section,

or 16 sectors, in order to make the change. The storage array works in even larger segments, so to store the change it must transfer an entire 32K track from the source to the target disk (the shaded blocks). In this case, an array-based replication tool would copy the entire 32K track, or 64 sectors, in order to replicate the four sectors that represent new data. The 60 sectors are simply the same data written again for the sake of disk array efficiency. While this may be efficient from the perspective of the disk array, it is terribly inefficient for replication, as it would send 15 times more sectors over the WAN than what is actually needed to replicate the true data changes.

Compare this to the next graphic. Instead of sending 32K of data over, it only sends the 512 bytes that were actually changed. That makes the FalconStor solution 64 times more efficient than storage array-based solutions on a per-sector basis in terms of WAN traffic and storage capacity. Similarly, FalconStor data replication is significantly more efficient than other types of solutions.

Replication method	Minimum data amount moved for 512-byte change	How FalconStor compares
FalconStor	512 bytes	
Array-based	32 kilobytes	64 times more efficient
Appliance-based	16 kilobytes	32 times more efficient
Storage-appliance-based	8 kilobytes	16 times more efficient

512-byte sectors on disk array



Save Bandwidth, Storage, and Money with FalconStor

By optimizing WAN replication, FalconStor replication solutions can help you boost network performance without having to purchase expensive hardware. This improved network response time works in both directions — whether you are replicating data to a target, or you need to recover data after a corruption or failure. With significantly less data to move, your recovery operations occur much faster, speeding your organization's return to productivity. Because this technology is built into our solutions, your investment is minimal and the return on investment (ROI) is immediate.

To demonstrate the potential ROI, we completed a case study with a large law firm with multiple data centers. Over an 83-day period, this organization replicated data changes that took up 1.2TB of capacity per day. However, the FalconStor solution removed 1TB of redundant data on a daily basis, for a whopping 84% savings in both WAN bandwidth and storage requirements. If they were to copy all the data changes made at the file system level, it would have taken 68 hours to replicate each days' worth of data: an impossible situation. FalconStor technology reduced this organization's daily replication time from 68 hours to 10.7 hours and saved them \$600,000 annually in WAN costs. In addition, FalconStor software was able to reduce storage requirements at the DR site from 101TB to 16TB — saving 630% in storage capacity.

FalconStor solutions with built-in data replication are available to organizations of any size or budget and facilitate the use of existing remote locations for replication. In the past, many organizations had to implement expensive, third-party DR solutions. With FalconStor technology, they can replicate to a branch office by simply using their current WAN. This makes it easier to comply with increasingly stringent regulatory requirements for data protection.

Without FalconStor MicroScan bandwidth reduction



DS3 = 30 hours, 9 minutes
OC3 = 8 hours, 28 minutes

Costly OC3 is required to meet an 8-hour replication window

With FalconStor MicroScan bandwidth reduction

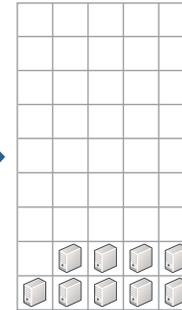
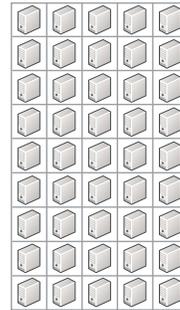


DS3 = 6 hours, 1 minute
OC3 = 1 hour, 41 minutes

A DS3 circuit completes replication in less time than an OC3 without FalconStor MicroScan replication

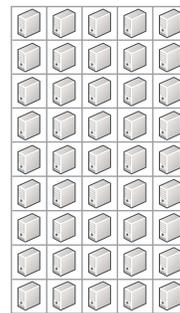
FalconStor MicroScan technology greatly reduces the amount of bandwidth needed for replication. If your systems generated 500GB of data per day at the file system level, it would take an expensive OC3 circuit to replicate all those changes within an eight-hour window. Using FalconStor MicroScan, over 80% of file system changes are filtered out (the ratio is even higher with certain data types). This means you can replicate the same daily changes using a much less costly DS3 circuit while actually finishing two hours before the eight-hour window.

DS3 without FalconStor MicroScan



Replicate 9 servers in 8 hours

DS3 with FalconStor MicroScan



Replicate 45 servers in 8 hours

Because of bandwidth costs, it is common to leave many systems unprotected by replication. Often, protection is limited only to the most critical systems. FalconStor MicroScan technology lets you extend coverage without increasing bandwidth. In this example, 45 applications create an average of 15GB of new data per day, per application. Using a DS3 circuit, you can only protect nine servers within an eight-hour replication window. By adding FalconStor MicroScan technology, you can protect all 45 servers within the same time window.

Summary

With cost savings of paramount importance in today's economy, simple, cost-effective WAN-optimized replication solutions such as those from FalconStor Software can ease multiple IT tasks. FalconStor technology enables you to easily improve network performance, reduce contention for network bandwidth, and meet your recovery point and recovery time objectives (RPO and RTO). Now, at last, you can complete backup, DR, maintenance, remote office consolidation, data center moves, and other replication tasks at a minimal cost.

For more information, visit www.falconstor.com or contact your local FalconStor representative.

Corporate Headquarters
USA
+1 631 777 5188
sales@falconstor.com

European Headquarters
France
+33 1 39 23 95 50
infoeurope@falconstor.com

Asia-Pacific Headquarters
Taiwan
+866 4 2259 1868
infoasia@falconstor.com

