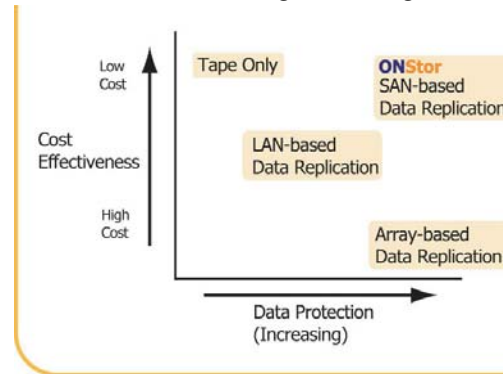


The ONStor NAS Gateway: Effective Solutions for Data Protection

High-availability data delivery

In a recent survey of 400 IT professionals, the majority of respondents identified “improving data protection” as their highest priority storage initiative. Many of them commented that guaranteeing a viable backup and restore process, combined with the requirement to maintain 7x24 data availability, was becoming increasingly challenging.

Given the growing demand for full-time data access, it is remarkable that the technologies most widely employed to protect data have changed little in the past 20 years. In the same survey, over 80% of the respondents reported that tape-based backup remains their primary data protection mechanism. While tape does address the primary objective of delivering a viable backup and restore process, it does not effectively address the objective of ensuring continuous data access. Indeed, the backup process itself often obstructs data access, and the restore process is excessively time consuming.

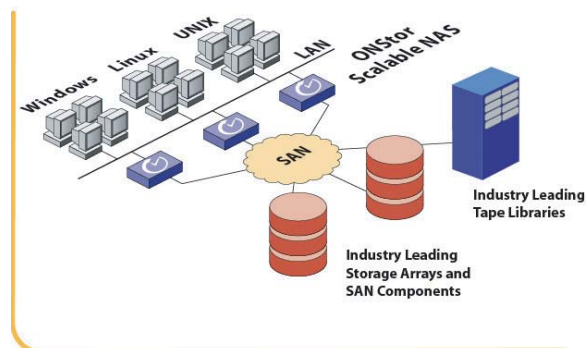


SAN-based data replication offers the security of array-based replication for less than the cost of many LAN-based replication solutions.

Fault tolerant storage environments

Most mission-critical storage environments augment tape solutions with other technologies such as data mirroring and server clustering to ensure continuous data access. Mirroring delivers high data availability by replicating data between two arrays. This technique, when combined with point-in-time copying, provides full-time data access to a known-good source. Server clustering ensures high system availability by providing multiple paths of access to the data.

Mirroring and clustering together provide a fault-tolerant environment that delivers both data and system availability, but at a price. The capital cost and management complexity of these solutions has limited their use to only the most mission-critical classes of data. NAS Gateway technologies from ONStor will change this. ONStor’s combined software and hardware approach will provide these capabilities at dramatically lower cost and with far simpler implementation, making non-stop storage practical for a much broader class of data.



To clients and servers on the LAN the NAS Gateway appears as a file server. All data is stored on arrays and tape libraries from third party vendors.

Introducing the NAS Gateway

The core of ONStor’s solution is the NAS Gateway, the first hardware and software solution designed from the start for one task: storing and managing files within a multiple-vendor storage infrastructure. Like a conventional NAS device, the NAS Gateway maintains file systems and makes files available to users on the LAN. But unlike a conventional NAS device, the storage itself is not a

closed, single-vendor solution. All data is stored on conventional Fibre Channel-attached disk arrays and tape libraries. The NAS Gateway is, in essence, a diskless server. Three technologies enhance the overall data availability of the NAS Gateway architecture:

- NAS Gateway clustering
- Disk-to-disk mirroring
- Open SAN storage.

The combination of these creates ONStor's unique, always-on storage environment.

Gateway clustering for system availability

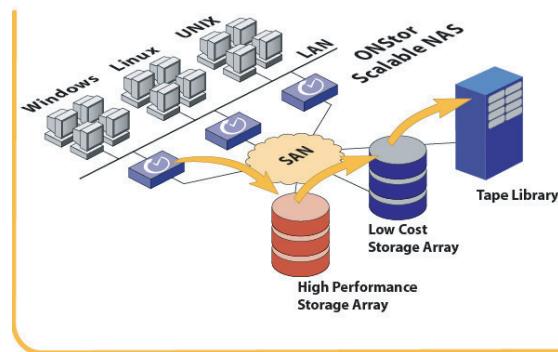
The concept of Gateway clustering is central to ONStor's high system availability. It eliminates single-points-of-failure in two ways: First, it permits users or servers to access data through any NAS Gateway in a multiple-Gateway environment. If a NAS Gateway is taken off-line or is experiencing an excessive workload, users may be easily and transparently migrated to another NAS Gateway with no changes in IP addresses or server names required.

Second, Gateway clustering means that any Gateway can access data from any storage array. In the NAS Gateway environment, all data, including file system directories and metadata, are stored on shared devices accessible over the SAN. Any NAS Gateway can access the data regardless of which NAS Gateway stored it. Because NAS Gateways offer petabyte addressability, each Gateway has essentially unlimited ability to directly access large pools of data. Unlike most virtual file system implementations, no redirection by another Gateway or metadata server is required. By eliminating both single-points-of-failure and performance bottlenecks, this architecture creates a highly robust storage environment.

Disk-to-disk mirroring for data availability

to-disk mirroring addresses this by maintaining a data copy that may be instantly brought on line if needed. In addition, the mirror copy acts as a source for backup data, allowing tape backups to proceed completely independently of the primary data. Compared with traditional backup methods, ONStor's disk-to-disk mirroring has these advantages:

- **Eliminates restore time:** Because the mirrored data can be instantly accessed, there is no need for a lengthy restore process. In conventional backup environments, this process may extend to hours or days.
- **Eliminates the "backup window":** Tape backup copies are drawn from mirrored data rather than primary data, so the primary data remains unaffected during the process. There is no performance impact and no need to take data off-line.
- **Allows more frequent updates:** When compared with traditional backup, mirroring is a straightforward process that requires relatively little server bandwidth or administrator attention. As a result, updates can occur much more frequently than is otherwise possible, thus reducing the exposure to data loss. Because ONStor's asynchronous disk-to-disk mirroring includes snapshot capability, it is possible to revert to earlier data versions if needed.



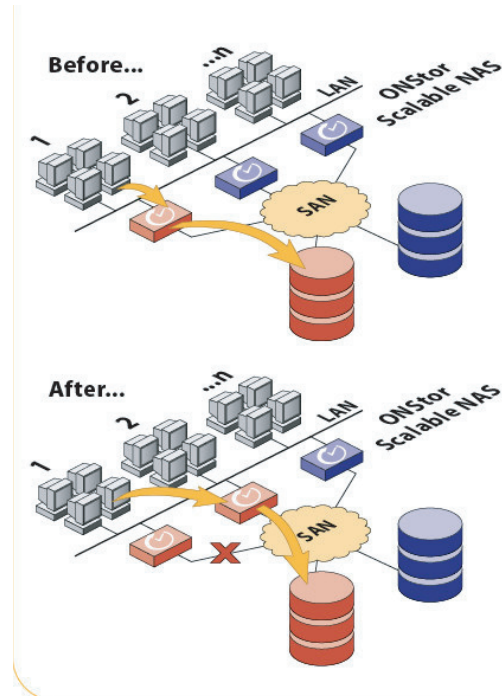
SAN-based asynchronous mirroring replicates data to a second array. When needed, any NAS Gateway may instantly access this copy. The copy also acts as a backup-to-tape source.

High-Availability Solutions Compared

When compared with other mirroring technologies, NAS Gateway-based mirroring offers several advantages that increase availability, reduce cost, and simplify implementation.

Compared with LAN-based mirroring When mirroring conventional NAS devices or servers over the LAN, the target storage device is usually a second NAS device or server. The mirror device is often similar to the primary device, which means that the associated processing power, operating system, and storage media are replicated. This virtually doubles the overall cost and management complexity of the solution.

ONStor NAS Gateways provide a simpler and more cost-effective approach. With NAS Gateways, data maintained on one disk array is asynchronously mirrored to another array. The NAS Gateway's unique multi-tasking architecture supports this mirroring process with no degradation to the performance of ongoing file services. A storage array acts as the mirror target, and a single NAS Gateway directs the mirroring process. No additional hardware or third-party software is required, providing a simpler, less costly, and more reliable solution.



N-way clustering provides better availability than traditional clustered pairs.

If the application demands high *system* availability as well, two or more NAS Gateways may be used to create multiple paths to the data. This approach also has advantages over LAN-based mirroring. Because all NAS Gateways have access to all connected data, when a switchover from the primary to a secondary array is required, the change can be nearly transparent to users. They will have access to the identical data set before and after the migration.

LAN-based mirroring, by contrast, requires users or servers to be migrated to a separate mirror device that does not have access to the primary server's storage. Depending on when the primary and mirror servers were last synchronized, users may see discrepancies in the data set after the migration has occurred. Furthermore, when the primary server is brought back on-line, it will again have to be synchronized with the mirror server, a time consuming process.

Finally, LAN-based mirroring may increase LAN traffic congestion and security concerns. The NAS Gateway method achieves both cost savings and reliability by minimizing data traffic and overall complexity, and by virtually eliminating user disruption.

Compared with array-to-array mirroring High-end array-to-array data replication solutions usually employ block-level mirroring. Data copies are maintained by the storage arrays themselves, independent of direct server control. While this technique does provide effective data protection, it tends to be costly for two reasons. First, both the primary and mirror arrays must be from the same vendor and must generally be of the same class. Since the arrays that perform this function are typically at the upper end of the price scale, acquiring several of them, plus the associated control software, is an expensive proposition. Second, because the mirroring process is controlled by the array, the process must be managed separately from the servers themselves. This adds significant complexity that contributes to administrative workloads and IT department skill set requirements. The combination of these usually restricts homogeneous array-to-array mirroring to only the most mission-critical of data.

ONStor NAS Gateway mirroring provides the data protection of these higher-cost solutions while substantially reducing the capital expense and administrative workload. NAS Gateways reduce capital

expense with a multi-tiered approach to mirroring that allows storage types to be matched with application requirements. The primary and mirror arrays may represent different classes of device, and may even be from different vendors. NAS Gateways also reduce administration workload by centralizing the management of the Gateways and the mirroring function within one framework.

Open storage for unsurpassed ROI

The NAS Gateway's broad interoperability significantly boosts the return-on-investment for the total solution. Unlike systems that are built around one vendor's storage devices or infrastructure, NAS Gateways are compatible with a wide range of arrays, switches, and tape libraries. This interoperability has powerful implications for both reducing the cost of high-availability storage and simplifying its integration.

Leverage existing storage & infrastructure To enhance the return on investment, NAS Gateways can leverage storage investments already in place. An existing SAN environment can be shared among application servers and NAS Gateways. Alternatively, components can be redeployed to create a dedicated file storage environment that is accessed by NAS Gateways. Either way, existing infrastructure can become an integral element of the future file storage solution.

Multi-tiered storage flexibility Not all applications demand the same level of performance and data availability, and it makes sense that data storage systems should have the flexibility to meet these varying requirements. But most file-storage systems are designed around proprietary storage and have little or no ability to include other vendors' solutions. NAS Gateways have the flexibility to store data on arrays ranging from high-end, high-performance subsystems to the emerging cost-effective SATA-based subsystems.

Non-disruptive integration The NAS Gateway's interoperability extends beyond infrastructure and arrays, to storage and device management software as well. This allows NAS Gateways to integrate with existing procedures and practices without disruption. From data backup processes to SAN management, NAS Gateways provide a solution that works with existing procedures, rather than replacing them.

Summary

NAS Gateways will dramatically increase an IT manager's opportunities to accomplish more with limited resources. This architecture provides comprehensive data and system availability protection at a cost comparable with far less robust solutions. The key attributes of NAS Gateways: **multi-tiered data protection** for data availability; **Gateway clustering** for system availability; and **open storage** for return on investment, combine to make non-stop storage a compelling solution for a wide range of environments.