

The Anatomy of a Consolidated Storage System



by George Crump, Lead Analyst

It has long been a goal of IT professionals to consolidate the storage infrastructure to a single storage system that would meet both the performance and long-term capacity needs of the data center. This goal however has proven to be challenging. Workload requirements seem to be too varied to create a one size fits all storage system. Storage systems are either strong on performance but expensive, or weak on performance and affordable. INFINIDAT's InfiniBox promises to provide the data center with a storage system that will allow them to achieve the consolidation goal, without sacrificing performance or increasing costs.

As Storage Switzerland discussed in its recent article "What are the Requirements to Consolidate Storage?" consolidating the various workloads in the data center to store their data on a single storage system has unique requirements. The storage system has to deliver performance to applications that demand it, and it has to provide cost-effective long-term storage of less active data. Most importantly, since the failure domain includes all the data center's data, the storage system should be well protected from downtime.

What is InfiniBox?

InfiniBox is a storage system based on three active-active compute nodes dedicated to running the storage software. Each node caches inbound writes and the most active data in both RAM and Flash. There is up to 3.2TB of RAM cache and 48TB of flash. The chances of data being read from a cache are incredibly high. The high capacity, memory-based storage front end allows the InfiniBox system to meet the performance challenge. INFINIDAT claims that the system can deliver greater than 900K Random IOPS.

High-performance storage needs high-performance networking to make sure that hosts can efficiently transmit and receive data. Also, because of the potential for a wide variety of environments connecting to the system, a consolidated storage system needs a high number of mixed connection types. InfiniBox meets this requirement by providing 24 x 8 Gbit FC Ports 6/12 x 10 GbE Ports.

For cost effective long-term storage, INFINIDAT leverages a high number, up to 480, of high capacity hard disk drives. Today INFINIDAT solutions can scale to 2PB of usable



HDD storage in a single rack. After data has been stored in the RAM/Flash tier, it is then flushed systematically and sequentially to a hard disk tier made up of eight, sixty drive hard disk shelves. Active data remains in cache, but inactive data is removed from the cache, as space is needed.

The system uses the best attribute of each tier (RAM, Flash, and HDD). For example, writes are written redundantly to memory and then sequentially written to hard disk drives, making for efficient HDD write and read performance. Also, each HDD volume consists of 16 drives from the potential 480 drives, which means that there is not the performance delta between HDD and Flash tiers that is found in some systems.

Recently, INFINIDAT began shipping a new entry-level system for mid-range environments with 250TB of storage in an 18U configuration. It leverages the same software code and provides the same degree of availability as its bigger brother.

Connectivity for Every Occasion

An important aspect of performance, especially when trying to consolidate down to a single storage system, is the ability to support many different types of workloads. The consolidated storage system may be providing storage for a variety of environments including virtual servers, virtual desktops, legacy scale up databases, and next generation NoSQL applications. It may also serve as a data lake for analytics data serving Hadoop projects.

Another important aspect of connectivity is providing the various protocols for the above operating environments. The needs of these environments range from NFS and CIFS to various block protocols (iSCSI, Fibre) and object storage. A consolidated storage system should also provide mainframe connectivity since, especially in the enterprise; mainframes continue to play an important role. InfiniBox is designed to support all of these protocols. Block and file protocols are available today with the others being released over the next year. These protocols are all provided via a single unified architecture; there are no bolt-ons or gateways involved. As a result, each protocol gains full data services like snapshots, clones, and dual parity RAID.

Consolidated Availability

One of the concerns when consolidating storage at this scale is the size of the failure domain. Essentially all the data center's eggs are in one basket. That basket has to adhere to a higher standard of resiliency from failure. For InfiniBox, availability starts at the core of the architecture. Data is served continually, even if there are two compute node failures thanks to the 3-node active-active architecture.

As for data protection, the InfiniBox leverages dual parity RAID, meaning that three drives within a volume would have to fail prior to data loss occurring. InfiniBox, despite using high capacity drives, leverages intelligent drive rebuilds so that only written capacity is



rebuilt, not the entire drive. They also leverage massively parallel RAID rebuilds and claim a rebuild time of fewer than six minutes after the failure of two drives.

Finally, the system also includes end-to-end data verification that provides protection against lost writes, bit rot and writes to the wrong location. Write verification features are important to make sure the information presently being written is correct, but also to ensure long-term data quality.

The result of all of this attention to availability is that the InfiniBox, in both configurations, delivers 99.99999% (7 nines) uptime. While this level of uptime may seem like overkill when consolidating to a single system, it is critical. The storage system simply cannot go down. It is clear that INFINIDAT has gone to great measure to make sure that theirs will not.

Disaster Recovery

The final element in availability is protection from a data center disaster, INFINIDAT provides Near Sync Replication to allow organizations to recover from this worst-case scenario. Near Sync is an asynchronous replication solution that provides unified consistency groups and delivers a recovery point objective of fewer than 4 seconds. The InfiniBox system can provision, automatically, a replication volume on the remote system, eliminating the most common error in DR replication design; forgetting to provision remote storage.

[INFINIDAT](#) provides Near Sync replication at no additional charge.

StorageSwiss Take

A Storage Consolidation project can be a process at odds with itself. Indeed, consolidating to a single system should simplify storage operations and lower overall storage costs but, in theory, it puts the data center at risk. If that storage system fails all data center operations cease. Also, there are concerns about balancing performance and cost.

InfiniBox addresses these concerns head-on. Its unique three-node active-active compute layer combined with its intelligent use of three types of storage allows it to meet the price/performance concern. Its extreme focus on reliability, evidenced by its seven nines of uptime, means that data centers can be confident that the system will always be available.