

InfiniBox Unified Storage: NAS Capabilities

Common Challenges with Unified Storage Systems Today

As Network Attached Storage (NAS) continues to outpace growth expectations, storage vendors often respond to this demand by creating solutions from limited, inefficient architectures. For example, implementing NAS over SAN (which creates performance bottlenecks), providing NAS-only solutions (which creates storage silos), or providing scale-out solutions (which lack the ability to automatically distribute workloads across all nodes).

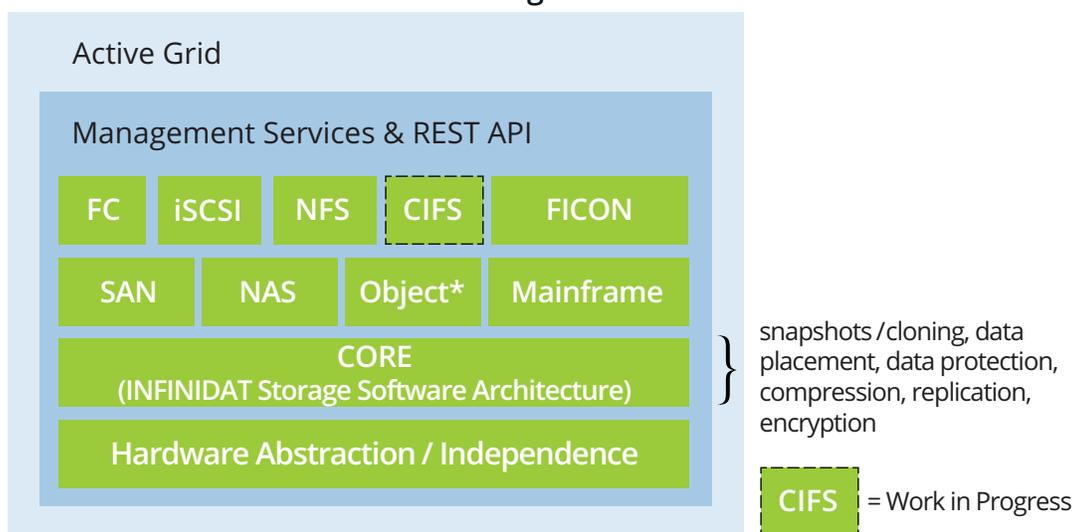
InfiniBox — A Different Approach

Unlike these solutions, InfiniBox was designed from day one to be a unified storage solution. Its core architecture is based on a data layout and metadata structure that caters to the needs of each protocol type (block and file). This core includes all the data services (snapshots, clones, dual-parity data protection, common data pooling, etc.) that both SAN and NAS need, and use. This allows INFINIDAT to offer a similar feature set on both SAN and NAS, simplifying storage administration and reducing footprint, cost and complexity. Our core design and all of our development are based on three priorities — reliability, performance, and ease of use.

A Modern NAS Architecture

Designing a NAS solution from the ground up to scale to multiple petabytes requires a lot of attention to performance optimization. With an extremely efficient Btree-based metadata, InfiniBox NAS leapfrogs the large number of indirect metadata reads required to perform a read from disk in large NAS environments. It always points directly to the iNode, reducing read latency to the absolute minimum. Combined with a thick SSD layer to cache data and metadata, InfiniBox NAS offers an ideal combination of cost and performance.

InfiniBox Unified Storage

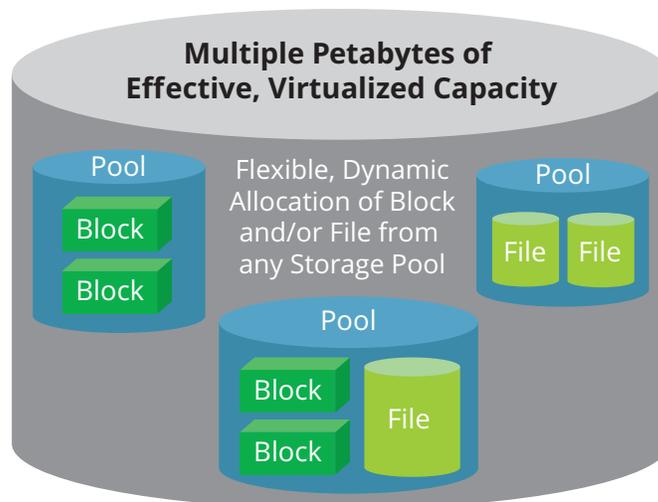


*Native Object protocol is a statement of direction

Unified Block & File Storage at Scale

Our unified block and file solution is designed specifically to provide a complete set of capabilities including:

- Virtual capacity pools
- File systems and block devices are allocated from storage pools
- Pools can be as large as the entire Useable Capacity (for block or file)
- Pools are virtual, and share all physical resources to maximize HW utilization
- Support for NFSv3
- Up to 4,000 file systems or one huge file system, with snapshots and clones in a single array with multiple petabytes of effective capacity
- Granular export management
- Thick or Thin file systems, with sparse file support for efficiency
- Double protection across all layers (software and hardware) for resiliency
- Self-Encrypting Drives (SED) for data security
- Thin, non-locking, smart snapshots and clones for fast backup, recovery and testing
- Run SAN and NAS side by side from the same system, with the same management tools
- Rich CLI/GUI/RESTful API/Python SDK
- VMware HCL certified



A Whole New Level of Reliability

Our unified storage solution has been designed to eliminate all single points of failure and maximize redundancy to provide the reliability needed for today's enterprise applications. An unmatched 99.99999% uptime is delivered through an innovative self-healing architecture, high-performance double-parity RAID, and comprehensive end-to-end data verification capability. This architecture provides no more than three seconds of unscheduled downtime per year. In addition, the system is based upon a fully abstracted set of software-driven storage functions layered on top of commodity hardware. The result is a highly available system that is also very cost effective.

Common Use Cases

With NAS running inside the same array, InfiniBox has the ability to tightly integrate with the Core data services layer, similar to the way our SAN does. These integrations allow data placement on disk to be optimized to support NAS I/O patterns. In addition, SSD caching decisions (made by the backend) get hints from the NAS service, optimizing cache efficiency. The result is a unified platform that can support many common use cases:

Group Shares

- Simple management
- Scale — allows separation of workloads into dedicated, efficient filesystems
- High performance improves user experience
- High resiliency allows mission-critical shares to run non-disruptively
- Fast node failover allows maintenance non-disruptively
- Thin provisioned shares reduce storage waste
- REST API/SDK simplifies automated deployment

NFS for VMware

- High performance from thick SSD layer for mission-critical VMs
- Large datastores reduce administrative overhead
- Eliminates performance-degrading LUN locking operations
- Simplified backup/restore operations
- Thin provisioned datastores reduce capacity waste

HPC Scratch Space

- Optimized for high file count
- REST API/SDK simplifies automated deployment
- Thick SSD layer provides high performance for dynamic workloads
- All system resources are available to a single workload

File Distribution/Archive

- Optimized for high file count
- Allows creation of fewer filesystems for efficient archive administration
- SSD can be disabled for archives to benefit production workloads
- High density allows multiple petabyte archive in a single rack

INFINIDAT

www.infinidat.com | info@infinidat.com