

WHITE PAPER

Infinidat OpenStack Solutions




Table of Contents

Abstract	3
OpenStack Interest and Industry Challenges	4
How Infinidat Can Help	6
Infinidat Values	6
Scalability	6
Reliability	7
Efficiency	7
Simplicity	8
InfiniBox® in an OpenStack Environment	8
Block Storage (Cinder)	9
Object Storage (Swift)	9
Compute (Nova)	9
Image Service (Glance)	10
Conclusion	11

Abstract

Infinidat has developed a storage platform that provides unique simplicity, efficiency, reliability, and extensibility that enhances the business value of large-scale OpenStack environments. The InfiniBox platform is a pre-integrated solution that scales to multiple petabytes of effective capacity in a single 42U rack. InfiniBox's innovative combination of DRAM, flash, and capacity-optimized disk, delivers tuning-free, high performance for consolidated mixed workloads, including object/Swift, file/Manila, and block/Cinder. The InfiniBox™ SSA enterprise storage system takes performance even further and is designed for workloads that require consistent, microsecond latency for every I/O. These factors combine to cut direct and indirect costs associated with large-scale OpenStack infrastructures, even versus "build-it-yourself" solutions. InfiniBox delivers guaranteed 100% availability without resorting to expensive replicas or slow erasure codes for data protection. Operations teams appreciate our delivery model designed to easily drop into workflows at all levels of the stack, including native Cinder integration, Ansible automation playbooks, and powerful Python SDK and RESTful APIs. This paper discusses why Infinidat storage should form the backbone of your OpenStack infrastructure.

OpenStack Interest and Industry Challenges

OpenStack is a growing trend across the IT industry that offers many advantages, including standardization across vendors, accelerated application deployment and increased efficiency. Infinidat believes in the power of OpenStack and as a supporting member of the OpenStack Foundation, we frequently receive inquiries about OpenStack support from a variety of organizations. Those organizations typically fall into one of three categories:

1. Organizations that recognize the agility, flexibility and cost-saving possibilities OpenStack offers, but haven't actually deployed an OpenStack environment in production
2. Organizations that have OpenStack environments in production which incorporate white-box storage hardware—a build-it-yourself type approach leveraging storage software layers like Red Hat Ceph (for larger deployments) or LVM (for smaller deployments)
3. Organizations that have OpenStack environments in production which incorporate branded storage hardware from traditional and startup vendors

The third category of organizations we encounter is dramatically smaller than the others, which speaks volumes about both prevailing client philosophies and the state of OpenStack capabilities among the storage vendor community.

OpenStack is a growing trend across the IT industry that offers many advantages, including standardization across vendors, accelerated application deployment and increased efficiency.

Some OpenStack practitioners, found in organizations from the second category above, adopt near-religious philosophies that all OpenStack deployments should be based on a highly homogenized set of hardware (a “sea of servers”) which supports all the various services required—from database, to general-purpose compute, to storage, to network.

This homogeneous “build-it-yourself” approach has benefits at both the very small and hyperscale ends of the spectrum. At very small scale, it doesn't make sense to separate core services like storage and compute into separate hardware platforms—the costs of adding dedicated infrastructure for different capabilities are too high to see significant incremental return. At hyperscale, the sheer procurement complexities of buying different permutations of dedicated infrastructure thousands of times over impose significant business costs—negotiating with hardware vendors not once or twice but tens or hundreds of times is time-consuming and costly. Moreover, ongoing development and support/maintenance costs associated with build-it-yourself deployment models can be amortized over a vast amount of infrastructure. So the return on build-it-yourself solutions can be positive, but only when the scale is either very small or very large.

Crossing all types of OpenStack deployments is a common thread of efficiency and resiliency challenges.

The cost/benefit forces are different for the vast set of OpenStack users in the middle, with deployment sizes ranging from tens to thousands of servers, usually with tens of terabytes to tens of petabytes of storage. For those midsize to large OpenStack deployments, most organizations can't afford to spend internal engineering and QA time trying to build, integrate, and tune various storage software layers to function with their chosen hardware platforms. In addition, to achieve enterprise-grade resiliency through the build-it-yourself approach, many organizations end up either buying more infrastructure than

they need to support the business, or spending more human capital vetting and qualifying both the hardware and software platforms, obliterating many of the perceived cost savings. Furthermore, the incremental costs to keep a build-it-yourself storage solution resilient as it grows over time add up dramatically—software used in such environments tends to follow a rapid, iterative release process, and new hardware must be qualified regularly. Unfortunately, these challenges are often overlooked until organizations are already heavily invested in the build-it-yourself approach.

Once organizations realize the challenges of scaling OpenStack deployments from initial proofs-of-concept to full production scale, they tend to move to the third category identified earlier, and start considering OpenStack storage solutions based on branded vendor hardware. Consultants or “managed” OpenStack distributions are often considered at this stage, too. Most storage array vendors offer OpenStack solutions that sound simple—traditional arrays updated with OpenStack Cinder support, for example. The challenges really start to appear when clients dive deeper and start noticing that among other things:

- ▶ Legacy array platforms usually require a separate management server to translate to OpenStack APIs
- ▶ There can be significant restrictions on the features that can be used with legacy storage platform OpenStack drivers
- ▶ Authentication layers can be varied and complex

OpenStack storage must address three key challenges—simplicity, efficiency, and resiliency—all with an eye toward total cost of ownership (TCO).

Suddenly, clients in the third category find themselves in a bind: what initially sounded like the easier, cheaper option ends up requiring dramatically more effort and significantly more cost. The net result is just another layer of complexity on top of their already complicated storage environments. This reality is the exact opposite of the OpenStack philosophy.

Crossing all types of OpenStack deployments is a common thread of efficiency and resiliency challenges. From the build-it-yourself perspective, typical expectations for scale-out, software-defined storage solutions include:

- ▶ Several replicas of data to try to achieve four-to-five nines resiliency
- ▶ A forced choice between expensive high-performance or low-cost high-capacity alternatives—never both
- ▶ An inefficient high ratio of controllers to storage capacity

Branded vendor storage typically does not require so many replicas to get decent resiliency, but the other factors remain—and of course, there is a massive price hike when putting one of these vendor’s logos on the box.

In today’s world, costs of raw storage capacity and compute still add up dramatically, workload profiles can vary dramatically minute-by-minute, and four-to-five nines availability is unacceptable, especially at scale. An ideal solution for today’s large-scale OpenStack environments needs to challenge the traditional expectations associated with both ends of the spectrum. Infinidat believes the lack of storage solutions that address all of these challenges for large-scale OpenStack environments is a major reason why so many organizations are:

- ▶ Stuck dreaming of the values of OpenStack—Category 1
- ▶ Struggling with scaling build-it-yourself solutions—Category 2
- ▶ Just plain frustrated with the clash of old vendor architectures trying to adapt to new platforms—Category 3

Fortunately, the advent of truly OpenStack-optimized storage systems like Infinidat InfiniBox gives mid-to-large scale OpenStack users a better path.

How Infinidat Can Help

OpenStack storage must address three key challenges—**simplicity, efficiency** and **resiliency**—all with an eye **toward total cost of ownership (TCO)**, because technology decisions are usually economic decisions in disguise. With over ten years of development and hundreds of petabytes of real customer deployments, the Infinidat storage architecture is proven to excel at all three factors across a variety of use cases, and the Infinidat business model enables clients to achieve unprecedented economics at multi-petabyte scale. Infinidat's comprehensive approach to OpenStack integration makes it easy for OpenStack users to consume these advantages across their entire OpenStack infrastructures.

INFINIDAT VALUES

The Infinidat storage software architecture², as implemented in InfiniBox, embodies the next generation of highly optimized storage for multi-petabyte scale deployments, at a fraction of the cost of traditional solutions. InfiniBox addresses diverse workloads including traditional line-of-business applications, virtualization platforms like VMware, high-performance computing (HPC) environments, cloud platforms like OpenStack, and many other data-intensive activities. This section describes the fundamental platform capabilities and values that enable this kind of large-scale, mixed workload consolidation for our clients.

The scale of the Infinidat architecture can both consolidate existing workloads, as well as support new workloads like OpenStack, all on the same platform.

SCALABILITY

The development of InfiniBox started with scale in mind. Today, the most popular InfiniBox configuration includes over a petabyte of storage capacity, and larger clients choose to deploy more than 10 petabytes of InfiniBox storage. InfiniBox scales to multiple petabytes of effective capacity in a single, standard 42U rack. The scale of the Infinidat architecture can both consolidate existing workloads—clients have consolidated as many as 30 legacy arrays to a single InfiniBox—as well as support new workloads like OpenStack, all on the same platform. Clients can also choose to start small and grow, with starter configurations available in the 250TB range. Flexible consumption models, including elastic pricing and FLX STaaS offerings, allow clients to deploy a larger system than their current requirements and pay for capacity as they consume it.

Capacity and performance scalability typically go hand-in-hand. InfiniBox is designed from the ground up for unmatched performance, leveraging as much as 3TB of DRAM and 368TB of flash cache to achieve aggregate performance of over 25 GB/s and 2M IOPS and microsecond latency—depending on workload, a double-digit performance advantage over competing platforms. Infinidat storage software maintains a “heat index” for all stored data, and dynamic machine learning processes ensure that the most active data resides in DRAM or flash, while the permanent copy of all data is stored on cost-effective hard drives. This data movement process happens proactively and continually, in contrast to “tiering” approaches, with no need for user intervention or excess data churn within the system. All storage systems implement some kind of caching or tiering layer, but doing so effectively at scale is extremely difficult, and Infinidat holds several key patents in this space.

While scalability across performance and capacity dimensions traditionally requires multiple storage solutions, InfiniBox is a single storage software platform that can serve any OpenStack workload with extremely low TCO and a 100% availability guarantee.

²<https://www.infinidat.com/en/resource-pdfs/infinidat-storage-architecture.pdf>

RELIABILITY

InfiniBox delivers 100% availability guaranteed. InfiniBox achieves mainframe-class availability without incremental costs from proprietary hardware or excess hardware redundancy.

InfiniBox is a single storage software platform that can serve any OpenStack workload with extremely low TCO and seven nines availability.

InfiniBox protects data with a patented dual-parity, declustered RAID-like approach called InfiniRaid®, which automatically distributes data across all the system's enclosures and disk drives. InfiniRaid supports a concurrent loss of two drives without any disruption or performance impact. Since InfiniRaid uses all drives in the system at all times, recovering to a fully redundant state from a double failure of even the largest capacity drives takes only a few minutes. Even though it is designed for performance, InfiniRaid requires minimal overhead to achieve this outstanding resiliency.

Some clients choose to replicate InfiniBox systems for data center-scale disaster recovery purposes. The InfiniBox replication engine uses our innovative snapshot technology, minimizing data transfer across replication links and allowing systems to perform most replication activities leveraging fast cache alone. To replicate volumes, the primary InfiniBox ships snapshot logs to the secondary InfiniBox periodically based on a guaranteed RPO as low as four seconds. The resulting asynchronous relationship is responsive enough to replace synchronous replication in many client environments.

EFFICIENCY

Overall InfiniBox efficiency comes from several factors:

- ▶ Highly efficient InfiniRaid data protection—supporting overall 100% platform availability while improving raw capacity utilization versus typical data protection implementations
- ▶ Optimized use of DRAM, flash and hard drives—effectively combining solid-state performance and capacity-optimized hard drive economics at multi-petabyte scale
- ▶ Data reduction software features without compromises—in-line compression, space-efficient snapshots and thin provisioning

While all data reduction features are important, the values of InfiniBox space-efficient snapshots and in-line compression particularly stand out for many workloads. InfiniBox snapshots offer a combination of performance and scale (100,000 snapshots per system) that is particularly attractive for modern DevOps workloads that are designed around many copies of data. InfiniBox snapshots can also be easily integrated into environments from backup solutions like Commvault to active data layers like OpenStack Cinder. Leveraging snapshots instead of writing redundant data dramatically enhances overall system efficiency. Meanwhile, InfiniBox real-time data compression increases system capacity efficiency regardless of workload layout, without impacting performance. Compression occurs as data is destaged to disk, after data is protected in mirrored controller memory. As a result, writes get acknowledged quickly, but the compression engine has more time to run before data is destaged, yielding higher compression ratios. These advanced data reduction features multiply InfiniBox efficiency for many workloads.

InfiniBox has designed an efficient platform that minimizes environmental impact and physical footprint during both the manufacturing process and operations. A fully configured InfiniBox with multiple petabytes of effective capacity occupies a single 42U rack and consumes 8 kW of power under peak load, for a green energy footprint of under 4W per usable TB. Other storage systems offer some combination of efficiency features, however, none provide all as effectively at scale as InfiniBox.

SIMPLICITY

InfiniBox has been designed to be easy to use and integrate into existing data center workflows, regardless of desired consumption model, deployment environment, or level of storage expertise. The Infinidat data path already abstracts away many of the challenges of traditional storage and even new, build-it-yourself architectures, with no need to keep track of different storage media or physical data layout through the life of the system. Users also don't need to worry about protocol choices, as InfiniBox delivers truly unified storage designed natively for multiple workloads and protocols, including Fibre Channel, iSCSI, NAS and SMB. Clients can easily deploy OpenStack via iSCSI, VMware via NFS, and bare-metal workloads via Fibre Channel, all in the same InfiniBox.

Infinidat's end-to-end simplicity allows today's companies to focus more resources on innovation and less on storage infrastructure.

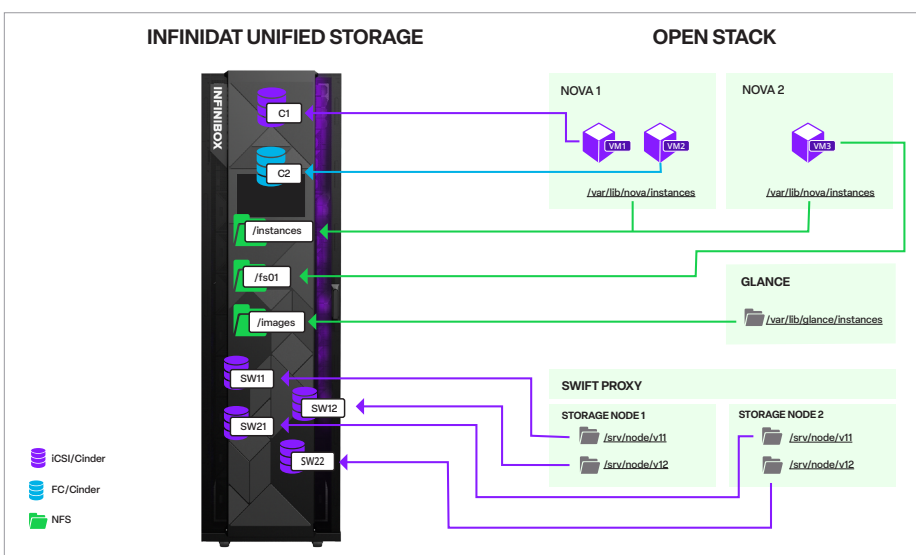
Infinidat exposes that simplicity to the administrator via an intuitive HTML5 GUI that facilitates deploying storage within minutes of installing new InfiniBox systems—no training required. For lower-level automation, all InfiniBox functionality can also be controlled using a native non-blocking RESTful API, Python SDK, and command-line interface (CLI), and Ansible Playbooks are available. For server-driven workflows, Infinidat Host PowerTools³ and related integrations are available for VMware, OpenStack and many operating systems.

Simplicity is also reflected in Infinidat's business practices. All systems include 24x7 on site hardware and software support, and software licenses and pricing are all-inclusive and predicated on usable capacity. Infinidat's end-to-end simplicity allows today's companies to focus more resources on innovation and less on storage infrastructure.

InfiniBox in an OpenStack Environment

OpenStack was designed from the beginning to offer a flexible assortment of services that may be deployed depending on the needs of the particular environment. One of the major strengths of InfiniBox in OpenStack environments is that it can interact with a variety of these services depending on use case and architectural preferences. Typical OpenStack services that can integrate with InfiniBox include:

- ▶ Block storage (Cinder): persistent block storage for virtual machine instances
- ▶ Object storage (Swift): persistent storage for unstructured data via HTTP-based APIs
- ▶ Compute (Nova): virtual machine instance creation and management
- ▶ Image service (Glance): discovery, registration and delivery services for disk and virtual machine images
- ▶ File storage (Manila): persistent file-based storage for virtual machine instances



Infinidat Unified Storage catering to all OpenStack storage needs

Several of these integrations are detailed later in this section.

³<https://www.infinidat.com/en/resource-pdfs/host-powertools.pdf>

Block Storage (Cinder)

Infinidat provides a native OpenStack Cinder block storage driver for InfiniBox which supports both Fibre Channel and iSCSI storage connectivity. This driver provides all necessary Cinder functionality, including volume management, snapshot management, migration and more. InfiniBox Cinder integration directly capitalizes on the capabilities and simplicity of the Infinidat storage architecture, unlike competitive solutions which require ongoing tier management, careful snapshot planning, and performance planning due to legacy back-end storage infrastructures. These capabilities are particularly important as workloads change during the growth of a Cinder deployment.

In addition to the Infinidat Cinder driver, Infinidat offers an OpenStack command-line interface (CLI) utility to simplify OpenStack configuration. Instead of manually editing OpenStack configuration files, the administrator just needs to run one simple command:

```
# infinidat-openstack --commit volume-backend set <InfiniBox-IP>  
<Username> <Password> <Pool-Name>
```

Once InfiniBox storage is added to the Cinder deployment, Cinder users can use standard Cinder tools to provision and manage InfiniBox-based block storage, without disrupting overall storage management activities or the activities of other InfiniBox users using different storage pools on the same system.

Object Storage (Swift)

Many clients consider InfiniBox for their OpenStack Swift (or other object storage environments) because it provides capabilities often not viewed as economically feasible in an object context—particularly high performance (over 2M IOPS) and high-resiliency (100% availability)—without needing the attention required for build-it-yourself solutions. OpenStack Swift can easily consume InfiniBox volumes with iSCSI or FC connectivity to Swift Storage Nodes running XFS file systems.

DRAM and flash cache layers built into InfiniBox provide dramatic performance improvements for Swift deployments. There is no need to carefully tune or balance a mix of different storage media to achieve optimal performance; the Infinidat storage architecture handles that automatically. Capacity economics are also improved for many workloads through native InfiniBox line-speed data compression.

There is no need to carefully tune or balance a mix of different storage media to achieve optimal performance; the Infinidat storage architecture handles that automatically.

Shifting primary resiliency requirements from the Swift layer to the back-end InfiniBox platform improves both availability and efficiency, eliminating the need for:

- ▶ Internal Swift data replication, which wastes storage capacity and multiplies required hardware to reach a particular resiliency level
- ▶ Swift erasure coding, which requires more powerful storage controller CPUs and tends to compromise performance

In addition, InfiniBox provides line-speed data-at-rest encryption for security and corporate governance requirements. Typical Swift deployments are unable to provide that capability without costly or complicated add-ons.

Compute (Nova)

Every virtual machine instance running on an OpenStack Nova compute server has a root volume as well as optional ephemeral (nonpersistent) volumes, often unaccounted for in OpenStack storage planning. In many installations, these ephemeral volumes are stored using local storage on the physical compute node. With this approach, live migration for a running VM may take time and impact performance, when it works; but by default, such migration is not supported at all:

```
# nova live-migration <guid> <second-cluster>
```

```
ERROR (BadRequest): <first-cluster> is not on shared storage:
```

```
Live migration can not be used without shared storage. (HTTP 400)
```

Some InfiniBox customers with OpenStack environments boot virtual machines with root disks stored on iSCSI or Fibre Channel-connected Cinder volumes. However, it's also possible to use InfiniBox NFS volumes for this purpose. Using InfiniBox NFS shared storage for OpenStack virtual machine root disks provides several key benefits that contribute to infrastructure agility:

- ▶ **Stateless compute hosts.** As long as there are no virtual machine instances running on a compute host, users can take that host offline or wipe it completely without affecting the rest of the OpenStack environment.
- ▶ **Easier recovery from compute node failures.** If a compute node fails, virtual machine instances are usually easily recoverable as their data resides on the InfiniBox and not on the compute node's local storage.
- ▶ **Simpler capacity provisioning and expansion.** Without changing hardware configuration or any other aspect of the compute nodes, it is possible to seamlessly provision and increase the capacity of VM repositories stored on the InfiniBox, either leveraging the simple InfiniBox GUI or APIs including Python SDK. Further, built-in InfiniBox compression often helps reduce overall VM storage requirements.
- ▶ **Faster performance.** Highly scalable InfiniBox NAS capability provides optimal performance leveraging DRAM and flash caching layers without the complexity of manually combining different storage media.
- ▶ **Live migration.** With InfiniBox NAS, live migration doesn't involve any data transfer across compute nodes—as long as they're using the same shared InfiniBox storage.

Image Service (Glance)

Basic OpenStack Glance image servers use local storage on OpenStack controller nodes. However, similar to OpenStack Nova VM storage, it is possible to use InfiniBox NFS mounts to store Glance images. The benefits are parallel to those explained previously for Nova: simplifying Glance capacity management scalability, removing the need to plan separate Glance image backup and replication, and taking performance concerns out of the picture.

Conclusion

Cloud computing is fundamentally changing the economics of IT, and open source platforms like OpenStack are leading the charge. The advantages of OpenStack cannot be understated, however, many organizations fail to realize the full potential of their OpenStack infrastructures—or even fail to launch them at all—due to storage constraints. Build-it-yourself infrastructures and typical branded vendor solutions have significant weaknesses for mid- to large-scale OpenStack deployments, from both technical and business perspectives.

Infinidat's innovative InfiniBox platform addresses these challenges and more, with a particular focus on the most critical challenges for growing OpenStack storage infrastructures—simplicity, efficiency and resiliency—all with game-changing economics. InfiniBox can be easily integrated into major OpenStack services, including Cinder block storage, Swift object storage, and Nova compute orchestration.

Infinidat understands that storage is the foundation on which successful cloud deployments are built. Our advantages combine to make Infinidat an ideal strategic partner to help power next-generation OpenStack infrastructures. Contact your Infinidat representative to learn more.