Flash Drives Storage Performance

Non-volatile flash storage, commonly referred to as the solid-state drive (SSD), has dramatically changed the way server storage looks today. And this has had an equally dramatic impact on the types of datacenter infrastructure that have become most prevalent as part of the Cloud Generation.
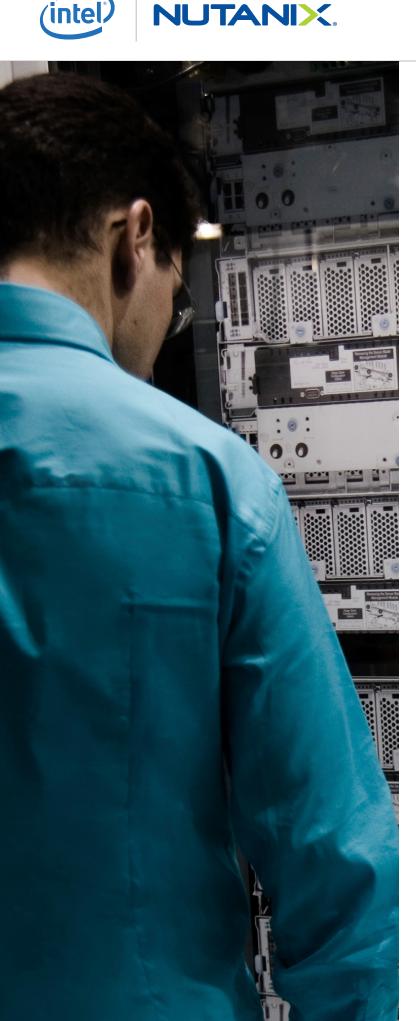
As of the late 2000s, virtualization had been the innovation that spurred breakthroughs in the datacenter. Virtualization, the method by which applications and workloads are abstracted from the hardware layer, paved the way for on-demand scaling. But in order for datacenter storage to keep up and deliver the much-needed throughput to make this scaling possible, datacenters needed centralized network storage systems. The result: costly, complex, and inefficient storage area networks (SAN) and network attached storage (NAS) systems.

The stage was set for the arrival of SSDs. All-flash arrays were introduced to help boost I/O even further, which led to dramatic performance improvements in virtual machine provisioning. SSDs also delivered up to 15 to 20 times greater reliability than standard hard disk drives.[1] The infrastructure became simplified as well, because configurations needed fewer drives to achieve the same level of performance. The bottleneck shifted away from storage and onto the network.

Compared to standard hard disk drives, SSDs deliver up to:

15x to 20x greater reliability[1]

The Cloud Generation has begun. The primary effort of this era is to break down the silos between different pools of resources: compute, network, and storage. The end goal is to deliver IT infrastructure with greater ease of use, rapid self-service, and a seamless experience for end-users. These demands have given rise to a new kind of datacenter infrastructure: hyperconverged infrastructure.

The New, Web-Scale World

Hyperconverged infrastructure combines both compute and storage into a single building block. This is connected over what is now the standard in datacenter 10-gigabit Ethernet. The storage layer is typically based on SSDs to deliver the high read/write speeds that are crucial to making the solution work.

In a traditional configuration, such as those you might find in the era of virtualization, compute and storage resources were segmented. All data was stored on hard drives that were situated behind a pair of hardware-based storage controllers, which handled storage functions including storage pooling, storage protocols, and features such as de-duplication, compression, and point-in-time backups or snapshots.

Because a typical configuration would have only two storage controllers, each controller would be optimized to not exceed a load greater than 50%. If one controller failed, the other could take over to maintain high availability. But this creates a huge bottleneck, and the controllers become a severe choke point even with all-flash storage.

Advancements in processor technology put even more pressure on traditional SAN and NAS storage environments. Processors are now designed with more cores and threads, able to perform more calculations at lower frequencies and lower power for greater efficiency. This, coupled with the potential gain from SSDs, creates an even wider performance deficit in SAN and NAS.

Hyperconverged systems are starting to replace traditional 3-tier architectures that include SAN and NAS, and here's why: the elimination of hardware-based storage controllers with their own networks can deliver higher read/write speeds without the bottleneck of storage controllers.

In a hyperconverged system, data storage is spread out across individual nodes so there are no centralized storage operations. This enables the processor to work in tandem with SSDs to deliver high performance that's also balanced; neither component is holding the other one back. The result? Truly distributed, scale-out storage. This is what makes hyperconverged infrastructure possible.

Nutanix Hyperconverged Infrastructure

The Nutanix Xtreme Computing Platform converges compute and storage-network resource pools (including Fibre Channel switches, etc.) into an easy-to-deploy, single, integrated appliance. Based on high-performing industry standard Intel® Xeon® processor-based architecture, this platform gives datacenters the flexibility to scale capacity, one node at a time. IT managers are now empowered to upgrade and expand at their own pace with linear, predictable scale-out and pay-as-you-grow expansion.

Nutanix is the right choice because it reduces datacenter complexity and the associated cost, deployment time, and management time that goes with it. An Intel®-based Nutanix infrastructure delivers 8x faster deployment at 50-60% lower total cost of ownership (TCO).[2] And, because it gets rid of the SAN and brings compute and storage together for virtualized environments, it eliminates network bottlenecks to ensure that users get the fastest possible speeds consistently, while getting the most out of their investments.
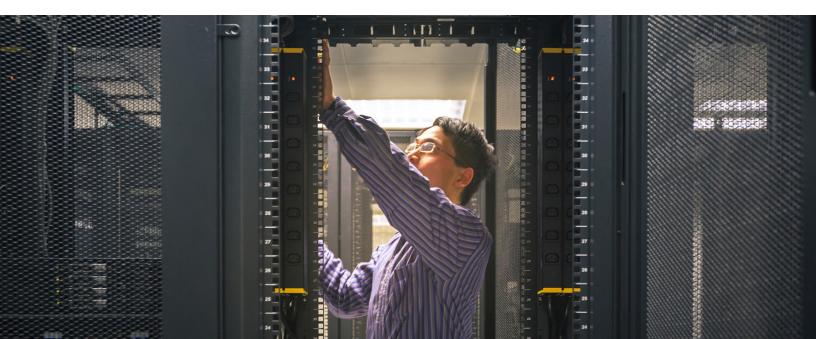
Nutanix Xtreme Computing Platform, based on Intel:

8x faster deployment[2]

50-60% lower TCO[2]

One of the key elements of hyperconverged infrastructure is its ease-of-use and integrated software, and the Nutanix Xtreme Computing Platform is no different. Administrators will save time and hassle with Nutanix Prism, an intuitive graphical user interface (GUI) that simplifies provisioning, updating, upgrading, and maintenance. Nutanix Prism also makes it simple to integrate third-party cloud management systems.

With Nutanix, datacenters can also seamlessly take advantage of the latest innovations in hardware platforms. As Intel launches the latest version of Intel® Xeon® processors, these become immediately available in Nutanix systems. Customers who make planning decisions well in advance of a purchase don't have to worry about getting stuck with legacy architecture. The systems they purchase and deploy will make use of the most up-to-date hardware, delivering an extra boost in performance.

A Convergence of Compute and Storage

Innovations in these key areas have made it possible to create hyperconverged solutions with Intel®-based commoditized hardware:

- **Processors:** The Intel® Xeon® processor E5 family delivers more cores and more threads. Embedded technologies enhance technical compute and provide hardware-level security that hardens platforms and helps protect systems from security compromises.

- **Solid-State Drives:** The Intel® SSD Data Center Series is fine-tuned for datacenter applications, delivering consistently amazing performance and reliability. SSDs continue to be refined with the inclusion of Non-Volatile Memory Express (NVMe) technology, designed specifically for flash.

- **Ethernet:** 10GbE Intel® Ethernet enables high-throughput for major applications such as big data and cloud computing. As the technology continues to evolve, Ethernet continues to get faster with 40GbE speeds for the most demanding environments such as HPC, and for whatever the future holds in store.

Versatility for Multiple Use Cases

Not only can Nutanix hyperconverged infrastructure help simplify datacenter architecture at lower CAPEX/OPEX, but it also enables efficient scalability with pay-as-you-grow economics.

Private Cloud and Server Virtualization – Deliver storage quality of service (QoS) for multi-workload and multi-tenant environments. Run popular applications such as Oracle*, SAP*, Microsoft SQL Server*, Exchange* and Sharepoint*.

Virtual Desktop Infrastructure (VDI) – Support end user computing, including persistent and graphics intensive desktops. Ready to run VMware Horizon Suite*, Citrix XenDesktop* and Citrix XenApp*.

Big Data – Dramatically reduce rack space and hard costs while running applications adjacent to other services with no performance degradation. Scale-out applications such as Splunk* work hand-in-hand with scale-out architectures such as Nutanix.

Data Protection and Disaster Recovery – Efficiently manage VMs and data replication. De-duplicate and transmit data with byte-level granularity for maximum replication efficiency. With the Storage Replication Adapter (SRA), you can also support VMware Site Recovery Manager* and third-party automation.

Enterprise Branch Office – Rapidly deploy standardized infrastructure to any remote or branch office location. Affordable entry-level platforms are also available for small enterprise locations.

A Hyperconverged Future

Nutanix helps pave the way to the Cloud Generation for enterprise datacenters everywhere. With the Nutanix Xtreme Computing Platform, businesses will benefit from VM-centric and scale-out storage based on industry standard Intel® SSDs and Intel® Xeon® processors for exceptional performance, efficiency, and reliability. For businesses of any size, a fast, responsive, and easy-to-use infrastructure starts with Nutanix.

Contact us at **info@nutanix.com**, follow us on Twitter **@nutanix,** or send us a request at **nutanix.com/demo** to set up your own customized briefing or live demonstration of Nutanix hyperconverged infrastructure powered by Intel.

About Nutanix

Nutanix delivers invisible infrastructure for next-generation enterprise computing, elevating IT to focus on the applications and services that power their business. The company's software-driven Xtreme Computing Platform natively converges compute, virtualization and storage into a single solution to drive simplicity in the datacenter. Using Nutanix, customers benefit from predictable performance, linear scalability and cloud-like infrastructure consumption.

1. Elerath 2007: Elerath (NetApp) & Pecht (U. Maryland), "Enhanced Reliability Modeling of RAID Storage Systems," 37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'07)

Jiang 2008: Weihang Jiang, Chongveng Hu, Yuanyan Zhou (all U. Illinois), and Arkady Kanevsky (NetApp), "Are Disks the Dominant Contributor for Storage Failures? A Comprehensive Study of Storage Subsystem Failure Characteristics," ACM Transactions on Storage, Vol. 4, No. 3, Article 7, Nov. 2008

Pinheiro 2007: Eduardo Pinheiro, Wolf-Dietrich Weber and Luiz Andr´e Barroso (Google), "Failure Trends in a Large Disk Drive Population," 5th USENIX Conference on File and Storage Technologies (FAST'07), February 2007

Schroeder 2007: Bianca Schroeder and Garth Gibson (Carnegie Mellon), "Understanding Disk Failure Rates: What Does an MTTF of 1,000,000 Hours Mean to You?" ACM Transactions on Storage, Vol. 3, No. 3, Article 8, October 2007

"Hard Drive Reliability Update – Sep 2014," Backblaze, 2014. http://bit.ly/1HyA693

"Hard Drive Reliability Stats for Q1 2015," Backblaze, 2015. http://bit.ly/1BrYtDt

"Les taux de retour des composants," Hardware.FR, 2010. http://bit.ly/1JXxSoT

2."The Path to Hyperconverged Infrastructure," Nutanix, 2015. http://go.nutanix.com/hyperconverged-infrastructure-for-the-enterprise.html