

8 Strategies for Building a Modern Datacenter

Nutanix White Paper





8 Strategies for Building a Modern Datacenter

Enterprise datacenters are straining to keep pace with dynamic business demands, as well as to incorporate advanced technologies and architectures that aim to improve infrastructure performance, scale and economics. Meeting these requirements, however, often requires a complete rethinking of how datacenters are designed and managed. Fortunately, many enterprise IT architects and leading cloud providers have already demonstrated the viability and the benefits of a more modern, software-defined datacenter.

This Nutanix white paper examines eight fundamental steps leading to a more efficient, manageable and scalable datacenter.

- ▣ Switch to Modular Building Block Architecture
- ▣ Converge Infrastructure Wherever Possible
- ▣ Let Software Drive the Datacenter
- ▣ Overcome Fear of Commodity Hardware
- ▣ Make End-User Computing a First-Class Citizen
- ▣ Break Down Silos for Increased Flexibility
- ▣ Build a Hybrid Cloud
- ▣ Move Beyond Disaster Recovery to Service Continuity

Switch to Modular Building Block Architecture

Enterprise datacenter designs are not for the faint of heart. They've never been simple, but each successive year they become more complex and more costly to maintain. Datacenter architects, of course, did not chart a course leading to greater complexity. Initial infrastructure designs, in fact, tried to rationalize the choice of technologies and products to make the overall environment easy to manage and support.

To keep pace with business, however, the datacenter has incorporated new technologies to drive increased performance and scale. As a consequence, it has become dependent on multiple generations of products and an incompatible mix of technologies across network, server, and storage silos. Throw in an unwieldy set of management consoles, dashboards, and frameworks and it is easy to see why today's datacenters are so difficult to manage, costly to upgrade, and nearly impossible to scale affordably.

Many architects are now embracing a modular design to expand existing complexes or to create new datacenters. These modular designs use homogeneous "building blocks" that can be added incrementally to reach the desired level of scale. Each module integrates fundamental datacenter components: networking, storage, and computing resources. Modular datacenter designs have many benefits.

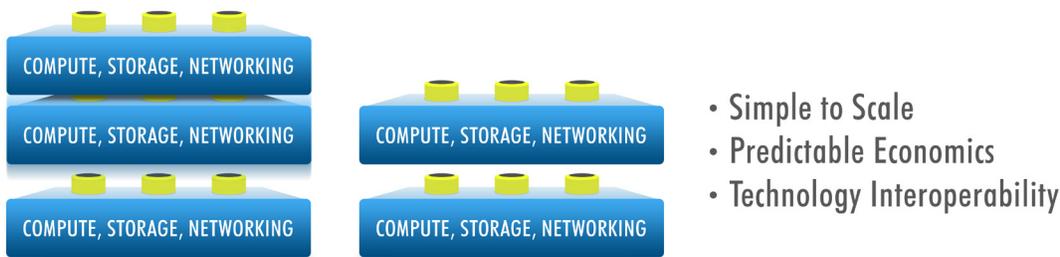
- Greater simplicity compared to traditional do-it-yourself approaches
- More predictable performance when scaling the infrastructure
- Better interoperability among technologies

Modular datacenters are not a new concept. One of the earliest attempts to modularize the infrastructure was to equip 40-foot shipping containers with 20 or more racks of server, storage, and networking equipment, pre-cabled and ready to go. The container approach technically fit the definition of a module, but few enterprises could afford or consume infrastructure in such massive increments.

More recently, modularization has been reduced to the rack level. Products like Virtual Computing Environment's (VCE) vBlock are delivered as pre-engineered, fully cabled racks containing servers, network switches and storage devices. These allow organizations to build out their datacenter one rack at a time, but there are several notable downsides.

- Solutions are still priced too high for many organizations. It's common to spend \$500,000 or more for even basic infrastructure.
- Fixed, vendor-defined ratios of compute, storage, and networking capacity make the systems inflexible and unable to efficiently match datacenter needs.
- Legacy components from multiple vendors in a single rack keep overall administration and management complex.
- The architecture makes it difficult to aggregate resources from individual racks, so it's not truly scalable.

True datacenter modularization occurs when building blocks can be quickly added to or removed from an infrastructure. A modular building block allows organizations to avoid expensive over-provisioning. The enterprise can start small with a single rack-mount appliance and later add platforms or appliances.



The Nutanix Virtual Computing Platform delivers an ideal building block that consolidates the compute and storage tiers into a single, integrated 2U appliance. It also leverages existing Ethernet networking investments out of the box. The modular building-block design lets organizations start with small deployments and grow incrementally into very large clusters. Nutanix eliminates the need for a dedicated storage network or array and streamlines overall datacenter administration with a single, intuitive management console.

Building an infrastructure with appliance-based modules that integrate networking, storage, and computing resources is a simple and economical means of scaling enterprise datacenters.

Converge Infrastructure Wherever Possible

Many enterprise IT managers have realized a multitude of benefits from converging datacenter infrastructure. They can typically deliver converged services more economically and more efficiently with fewer dedicated resources.

The evolution of storage strategies illustrates why convergence is so powerful. More than a decade ago, hard disk drives migrated from application and database servers to centralized shared storage arrays to improve overall capacity utilization. High-speed networks, such as storage area networks, connected servers to this monolithic storage.

More recently, flash memory has been added to enterprise storage devices to create hybrid storage solutions. When they are engineered correctly, these solutions yield up to 100 times faster

performance than legacy architectures. Some SAN architectures have been successfully retrofitted to accommodate flash-based storage, but often the adaptation creates new network bottlenecks. For example, 100 flash-based PCIe-based storage cards added to an 8 Gb/s SAN will overwhelm most fabric networks and array controllers.

IT leaders like Google and Facebook concluded that technologies like SAN and NAS have no place in a modern datacenter. They wanted compute and storage resources to be co-resident to deliver data I/O with very little latency and with the ability to scale with fewer constraints.

One more convergence step was needed to successfully eliminate complex storage networks from enterprise datacenters and to efficiently combine data and compute in one system. Architects had to achieve a clean translation from the various I/O stacks in the myriad operating systems to a few well-defined standards protocols (e.g., NFS, iSCSI, CIFS). VMware used virtualization to enable that translation and others like Microsoft with HyperV, Xen, and KVM have followed suit.

By adding virtualization to the convergence equation, the SAN/NAS software stack could sit on the server itself. With virtualized controllers it was finally possible to support SAN and NAS features in a single implementation for all operating systems. Now, all virtual machines (VMs) can communicate transparently with storage resources, just as they would across a SAN fabric. The difference is that storage and servers can scale together, without having to invest in high-speed, dedicated networks.

The Nutanix Virtual Computing Platform converges compute and storage into a single system, eliminating the need for traditional storage arrays. A single appliance contains four independent nodes, each optimized with high-performance compute, memory, and storage. Each node runs an industry-standard hypervisor, and a Nutanix controller VM, which handles all I/O operations for the local hypervisor. Storage resources are exposed to the hypervisor through traditional interfaces, such as NFS. A global storage pool aggregates storage capacity from all nodes and all hosts in the cluster can access that pool.

Converging compute and storage resources into a single appliance increases overall datacenter flexibility and minimizes latency. You retain all of the benefits but none of the complexity of separate servers and storage arrays that are interconnected by dedicated storage networks.

Let Software Drive the Datacenter

Traditional datacenters depend on specialized hardware, often for just a single function. Devices might include high-performance load balancers, hardware-based storage arrays, or any of the multiple security appliances that litter most datacenters. These hardware platforms have limited flexibility and portability. Because many of them are powered by FPGAs and ASICs, the systems cannot easily support new software capabilities. The runtime logic for some applications is embedded in the platform, and there is a tight coupling between the platform and the application. The resulting infrastructure is application-specific and is difficult to manage, support, and scale.

A software-driven datacenter that de-couples policy intelligence from the underlying hardware is a better and more extensible solution. It allows logic and policy management to be abstracted into a distributed software layer that can be automated and controlled centrally. Because the services are software-based, datacenter management teams can provision new services on-demand anywhere in the infrastructure without adding hardware.

Abstracting services logic into software frees organizations from dependence on specialized hardware. Incorporating REST-based APIs and other software interfaces into a software solution

creates a programmable infrastructure. Software-defined services can be defined, provisioned, and managed as part of a larger cloud management and orchestration strategy.

Nutanix incorporates an advanced, distributed software model. The software leverages economical commodity hardware, such as Intel CPUs and Flash-based data storage. As a Virtual Computing Platform, Nutanix delivers enterprise storage as a service to the virtualization layer. It makes storage a programmable, on-demand asset to all virtual workloads and applications in the environment. REST-based APIs allow the Nutanix infrastructure to automate the management of storage and compute resources.

Organizations that embrace software-driven services delivered on commodity hardware platforms are more agile and economical than datacenters powered by expensive, single function appliances.

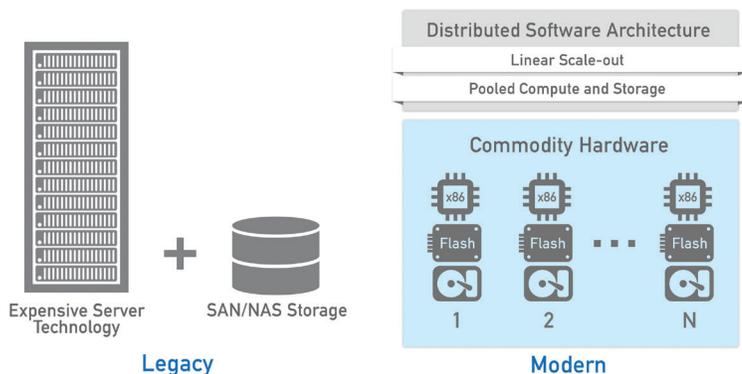
Overcome Fear of Commodity Hardware

If greater datacenter performance was needed in traditional operations models, the prescribed remedy was to either upgrade existing hardware or replace it with new hardware. Even today, many datacenters have designated refresh cycles for most hardware. Every three to five years, IT teams replace servers, networking appliances, and storage systems with more advanced technology. This strategy drives capital expenditures by discarding perfectly functional hardware and replacing it with more expensive equipment.

Not all organizations remain on this hardware-replacement treadmill. Google and other leading cloud providers found that the secret to massive scale without high capital expenses was to break its dependence on expensive, high-performance hardware. Instead, they achieved datacenter scale using a distributed software architecture that could run efficiently on low-cost commodity hardware.

With a software-driven strategy, additional commodity hardware is added to the infrastructure only if it's needed for datacenter scale or performance. Because the system pools resources across a cluster of machines, the servers and storage nodes do not all have to be best-in-class. Consequently, organizations are freed from regularly refreshing expensive proprietary equipment as resource demands increase. Additionally, replicating data and metadata across a single namespace cluster means no single server becomes a single point of failure.

Organizations that commit to commodity hardware have to determine what constitutes "commodity." For instance, cloud and large-scale web providers have determined that as long as the server is Intel x86-based, it generally doesn't matter what brand of server is used. So, vendor selection can be based on price and delivery timeframes.



Building the datacenter using “vanity free” servers radically alters IT economics. It begins to level the playing field between mainstream enterprises and leading cloud providers. Fortunately, the same original design manufacturers (ODMs) that supply the world’s largest and most reliable datacenters are now selling directly to large enterprises. So enterprise IT organizations, and IT solution vendors can lower hardware costs without compromising performance or reliability.

The Nutanix Virtual Computing Platform operates according to the same tenets as public cloud providers: Achieve massive scale and high reliability by running distributed software on clusters of commodity hardware. Nutanix delivers an economical 2U appliance, built from off-the-shelf components. It uses high-performance Intel processors and advanced, non-volatile flash-based memory.

Make End-User Computing a First-Class Citizen

Increasingly, employees bring their smartphones, tablets, and web-based laptops into the enterprise workplace. They expect seamless access to corporate data and applications. This “consumerization” of IT is a driving force behind the end-user computing (EUC) trend. EUC centralizes desktops, applications, and data within the datacenter and delivers these services to any device, from any location.

Many organizations ventured into EUC by virtualizing desktops. More recently, full EUC “stacks” have allowed enterprises to fully realize the promise of BYOD (bring your own device) by delivering applications, desktops, and data to any user device. These solutions also often incorporate self-service portals modeled after consumer web sites to accelerate end-user adoption.

While multiple software vendors race to meet this burgeoning demand with EUC solutions, the datacenter infrastructure is straining to satisfy the new requirements. Many IT teams model their implementations on standard server-virtualization methodologies. This may yield successful proof of concepts, but full-scale production deployments place such significant load on the network, storage, and server resources that scaling challenges quickly emerge.

VDI solutions, for example, drive very high IOPS requirements and need high-performance, low-latency infrastructures to ensure a high-quality experience for end users. If virtual desktops are run on the same general-purpose server-virtualization infrastructure, the resources needed for delivering the desktops at high speed may not be available. The lack of compute, memory, and networking resources at the right time jeopardizes the entire VDI investment.

The newer EUC suites exacerbate the problem by integrating enterprise data services, (e.g. ‘Dropbox for the Enterprise’), to bring “follow-me data” to end users. These EUC-enabled data services demand even more storage capacity from the infrastructure. Expanding infrastructure to support companywide EUC initiatives is not the answer.

Deploying a dedicated, standalone environment to support just EUC technologies, however, is also not a good solution. Although IT may have greater control of the EUC resources, there are substantial downsides.

- Independent datacenter silos are expensive to build and maintain, and they are frequently underutilized.
- EUC specialists are often designated to maintain the environment, which increases overall operating expenses.

A new approach to building datacenter infrastructure has to be agile and elastic enough to meet a range of workload demands. The datacenter must be able to handle compute-intensive VDI, storage-intensive enterprise data services, and the existing portfolio of virtualized enterprise applications.

Leading companies and government agencies around the world are adopting the Nutanix solution to power their EUC deployments and avoiding the common pitfalls. A single Nutanix 2U appliance delivers the physical compute and storage resources to power up to 400 virtual desktops in a single appliance. As the VDI environment grows, more Nutanix blocks can be added to dynamically scale the infrastructure.

The Nutanix Virtual Computing Platform provides high-performance, tiered storage with PCIe-attached Flash and more than 20 TB of storage. It dynamically pools storage resources across all nodes in the cluster to present a single namespace to all VMs. This distributed, software-defined storage strategy supports the storage-heavy data repositories required for the “follow-me data” component of EUC.



To successfully support EUC, IT must create a datacenter that is powerful enough to handle compute-intensive VDI and flexible enough to scale quickly with the business. To satisfy users who tolerate no downtime and a business that’s agile enough to compete, you need a robust and simple-to-operate datacenter.

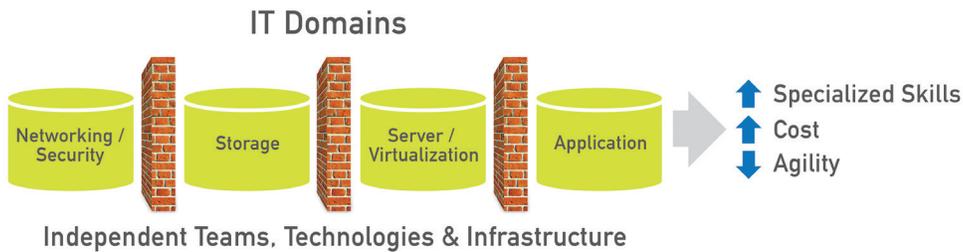
Break Down Silos for Increased Flexibility

Datacenters are designed to run and deliver applications for the business. While this mission is straightforward enough, the underlying infrastructure has devolved into discrete technology silos, each managed by a team of specialists. For example, most datacenters have a storage silo staffed by a dedicated team handling data management and information archiving. This team defines and manages low-level, storage-specific constructs such as LUNs, volumes, and RAID groups to support overall datacenter operations. The networking, server, and virtualization teams have separate silos, with additional discipline-centric technologies.

Segregating individual technology-specific capabilities into separate silos complicates ongoing management. It also requires that scale-out operations be managed independently, within each isolated silo. This organizational rigidity prevents traditional datacenters from being flexible and able to scale efficiently.

Rapid technology advancement within IT disciplines also means that each siloed team must constantly update its skill sets to effectively manage their responsibilities. Policy management illustrates the difficulty of operating a silo-based infrastructure: The networking team may rely on HP OpenView or Cisco Datacenter Network Manager (DCNM) to monitor the networking environment, while the virtualization team uses vCenter Operations Manager or Microsoft System Center.

Although specialization benefits the organization somewhat, the ultimate net effect is datacenter inefficiency. A multitude of dependencies between teams with specialized technical skills requires tight coordination to make even the most minor changes. For example, provisioning server resources to host a new virtual machine typically requires the networking team to configure IP addresses and VLANs. The storage team provisions LUNs, which must then be zoned on the SAN to enable devices to see one another. The storage team has to allocate additional capacity, while the server team must assign proper UIDs for the right services profile. Efficiently orchestrating even such a mundane task is difficult for most organizations.



Breaking up these technology silos requires a different strategy, one that exploits flexible, multi-purpose solutions. The ideal is a converged solution that collapses interdependent resources into a single scalable unit. By consolidating multiple technologies, converged solutions abstract the low-level intricacies within each silo so that technology management is aligned with datacenter-level objectives.

Converged solutions also streamline policy and resource management so that a single datacenter team can handle them without requiring a myriad of specialized skills. They simplify overall datacenter-growth strategies by scaling equally traditional IT domains.

By converging storage and computing into a single appliance, the Nutanix Virtual Computing Platform delivers fundamental datacenter infrastructure that helps organizations move away from a silo-based IT model. Nutanix provides a holistic services platform that incorporates all of the core functionality to run enterprise applications.

A converged solution that pools resources into a single scalable unit reduces the need for highly specialized skills. It helps organizations break down IT silos and more quickly respond to dynamic business needs.

Build a Hybrid Cloud

Public clouds provide Internet-accessible compute, storage, and network resources that are shared among multiple tenants. They have become an accepted part of an overall enterprise IT strategy.

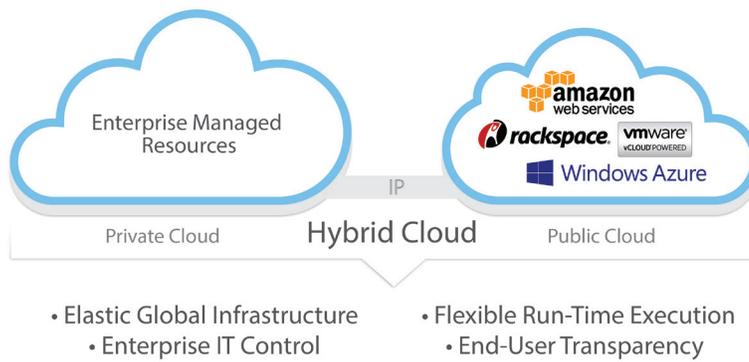
Many enterprises have identified applications that work well in public clouds, such as Infrastructure as a Service (IaaS). If an application has uneven and unpredictable demand characteristics, it might work best in a public cloud that can deliver true global elasticity. Public clouds are built to deliver self-service resources, making the infrastructure ideal for application developers who need rapid access to compute and storage capabilities.

Enterprises understand, however, that not all applications should be run in a public infrastructure, nor should all data be stored in a shared environment. Business-critical applications with confidential company or customer data, for example, are still run within the enterprise datacenter, often to meet regulatory mandates.

These enterprise datacenters are now designing private clouds that provide many of the attributes of public-cloud infrastructure, such as on-demand provisioning, resource-sharing across applications, or business units that remain under the control of the datacenter team. These private clouds can run on-premise or they can be hosted offsite by a provider using dedicated, non-shared infrastructure.

Internal datacenters and private clouds allow more control over security, and enable IT teams to take full responsibility for performance and security service level agreements (SLAs). So, internal systems will continue to co-exist with public cloud and SaaS-based services.

Both public and private clouds have emerged as viable choices for datacenter architects. But most experts agree that neither approach alone satisfies all of the performance, security, scale, and SLA requirements of enterprise IT. Modern datacenters must use both public and private cloud infrastructures in a hybrid cloud.



Hybrid clouds can handle “bursty” application traffic. When the demand for an internally hosted application exceeds the capacity of the private cloud, the excess demand can spill over into an IaaS environment. All users are served and enterprise SLAs are preserved. This situation is often referred to as cloud bursting.

In the future, more enterprises will split data-storage responsibilities between environments. Low-cost data archiving may be supported in a public-cloud environment, while more active or “hot” data resides in the private cloud to deliver the best possible user experience. This balances storage performance and cost. A hybrid cloud is also a viable solution for building disaster-recovery capabilities.

By integrating compute and storage resources into a single platform, Nutanix provides the infrastructure for secure, high-performance private clouds. It abstracts and pools storage resources across all nodes in a cluster. This pooled information is delivered to the virtual machines in the cluster as a single datastore.

This system enables IT to provision applications on-demand and eliminates manual configuration of storage resources. A Nutanix cluster can also be deployed in an off-premise private cloud, which enables automatic data replication between the enterprise datacenter and the private cloud infrastructure.

Additionally, the Nutanix software architecture is extensible. It will enable the easy consumption of public clouds via standards-based APIs so that a single Nutanix footprint can simultaneously leverage both private and public cloud resources.

Move Beyond Disaster Recovery to Service Continuity

Enterprise disaster recovery (DR) strategies are typically reactive. When natural or manmade disasters occur, application requests from end users simply fail over to a secondary site. After a period of service interruption, which could be minutes or hours, application services resume. Data availability and restoration are the primary focus of this method.

Many DR plans are defined by recovery-point objectives (RPO), which determine how much data can be lost, and recovery time objectives (RTO), which target how long it takes to recover data. Traditionally, this process has kept the business running and met user demands.

Consumerization of IT, however, has radically elevated user expectations. End-users are now accustomed to obtaining services on demand, with near 100 percent availability. If IT teams cannot match this experience level, end-users are likely to go around enterprise IT and procure cloud-based services on their own. To meet these new SLA mandates for always-on service, IT teams have to rethink how they architect, deliver and secure services.

Providing near 100 percent availability means that the legacy, data-centric approach to DR has to evolve to a more proactive and services-focused strategy, which is referred to as service continuity. There are two strategic approaches to achieving service continuity across the organization.

1) Re-architect your datacenters to be highly available. This approach fits most legacy application architectures, which are not meant to be physically or logically distributed. Good examples would be an SAP or an enterprise database. In these cases, datacenter architects often design active-active sites, based on metro clusters, with HA operating between the two sites.

These sites must have a great deal of bandwidth connecting them, as well as low round-trip times (RTTs) so that the infrastructure can handle latency-sensitive operations, such as synchronous replication between separate storage arrays. In the case of a datacenter outage, applications running in datacenter A, for example, would automatically fail over to datacenter B. Such deployments can be complex and expensive, as well as limited by physical constraints such as a maximum 100-km distance between sites.

2) Re-architect your applications to be distributed. Distributed application architectures are increasingly common. Major advantages include the capacity to scale globally, perform well, and deliver unmatched uptime. Facebook, Amazon, and Google are the most notable pioneers of this technology.

Distributed application architectures enable an application or service to be distributed among multiple sites, regions, and datacenters. If they are designed correctly, the result is a single global service, with multiple availability zones and failure domains to maintain service during site failures.



The Nutanix Virtual Computing Platform is built to run distributed applications within a single site and across multiple global locations. Built-in fault resiliency eliminates any single point of failure across the compute and storage tiers. VM-level replication and backup capabilities create “crash consistent” snapshots to preserve the state of a full application stack across all sites. You can manage clusters with a single UI.

Traditional disaster recovery is no longer sufficient for modern datacenters. End users expect 100 percent application availability and reliable performance SLAs. DR is evolving to a service continuity model and IT teams are taking a more proactive approach to managing outages.

Nutanix

Nutanix is the recognized leader in the fast-growing Virtual Computing Platform market. The Nutanix solution converges compute and storage resources into a single appliance, delivering a powerful, modular building block for modern datacenters. It incorporates the same advanced, distributed software architecture that powers leading public cloud providers, such as Google, Facebook and Amazon – but is tailored for mainstream enterprises and government agencies.

To learn more about what it takes to build a modern datacenter, go to: www.nutanix.com.