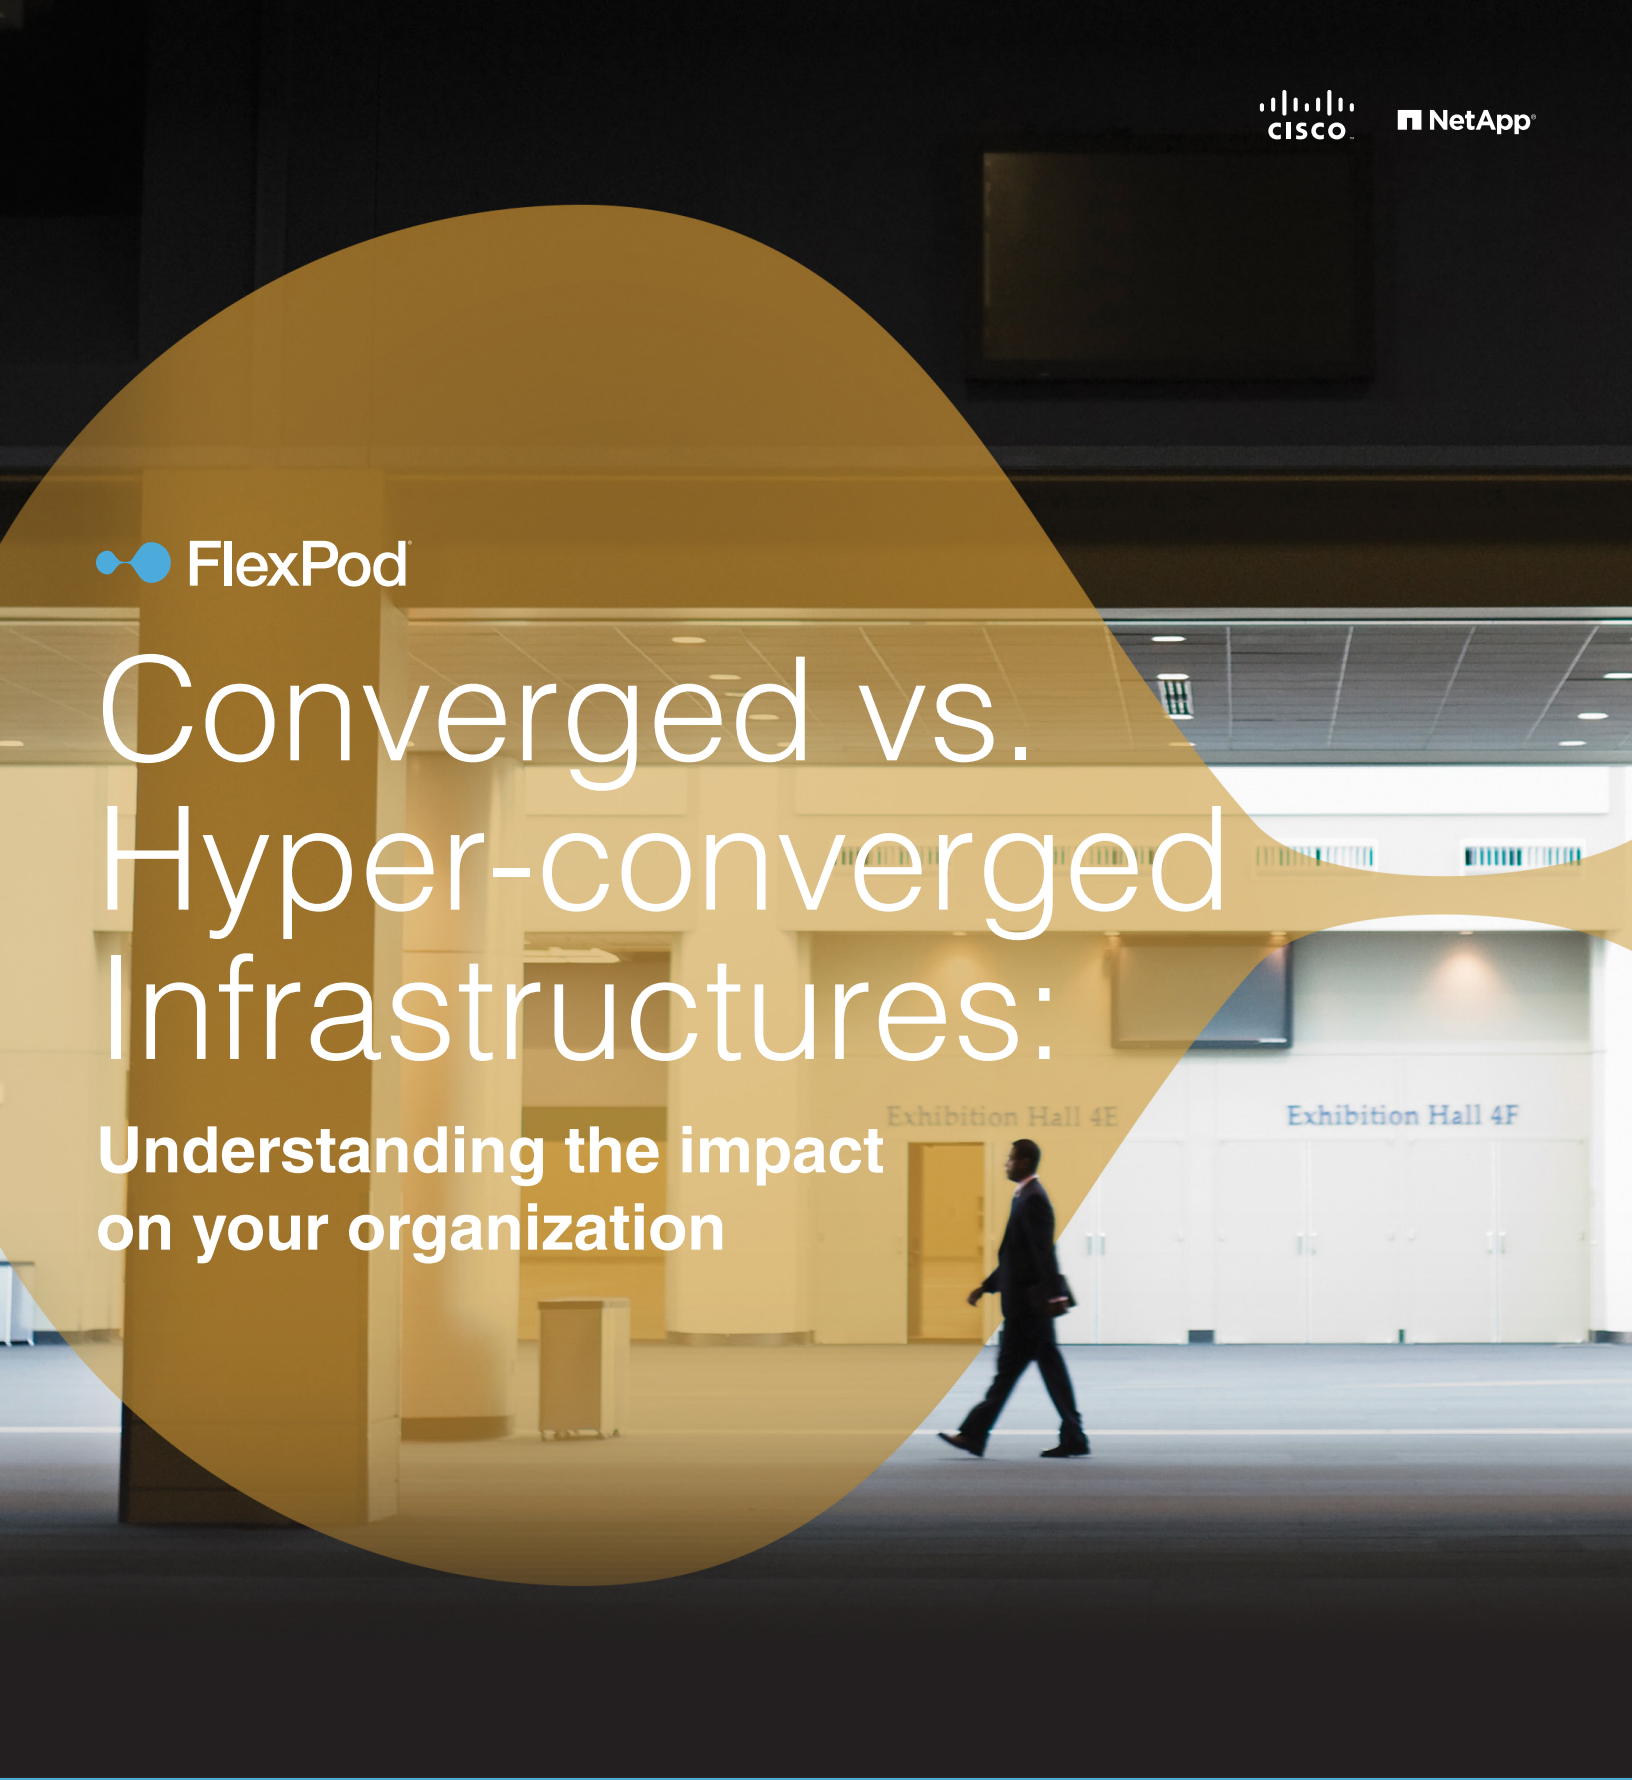

FlexPod

Converged vs. Hyper-converged Infrastructures:

Understanding the impact on your organization

More and more companies wanting to consolidate, upgrade and modernize their IT are moving to converged infrastructures. The question for many organizations is whether to choose converged or hyper-converged infrastructure.

Why convergence, why now?

In the past, data centers evolved without much foresight or planning as new business demands emerged, resulting in complex integrations of hardware, software and data. This was never a formula for maximized infrastructure performance or streamlined operations and required extensive expertise to establish and maintain. Now convergence – in the form of converged and hyper-converged infrastructures – offers a superior model in terms of both efficiency and business outcomes.

Convergence unites the four core aspects of a data center – compute, storage, networking and server virtualization – into a single framework. This powers organizations to achieve greater value and realize increased efficiencies by enabling more workloads to operate on fewer physical servers.

Converged infrastructure has advanced to the point where it can support IT organizations' rapidly emerging cloud-centric models, in which agility, scalability and security are in high demand. Converged infrastructure platforms are typically dedicated to improving business outcomes and increasing business value. As aging data centers are refreshed, and organizations migrate more workloads to the cloud, convergence is increasingly the approach of choice.

Convergence unites the four core aspects of a data center – compute, storage, networking and server virtualization – into a single framework.

Doing More With Less

Business is only accelerating, with IT bearing much of the burden of bringing new products, services and capabilities on line in radically reduced timeframes. To keep up, IT organizations must continuously evolve in order to deliver high levels of service and security that address these increasing business demands. At the same time, mobility and the proliferation of connected devices, along with disparate users accessing cloud-based applications, is creating massive complexity and new security challenges.

Meanwhile, IT organizations are challenged to make the most of every dollar spent. The majority of their budgets are allotted to maintaining current operations, which are increasingly expensive. Many enterprises opt not to hire experienced specialists with deep understanding of storage, networking, or compute environments. As a result, complex tasks are now being handled by IT generalists, or even business users. Clearly, increased efficiency and simplicity are paramount.

Handling Today's Workloads While Preparing For the Future

Legacy infrastructure is fundamentally challenged to provision enterprise applications at scale. Today's infrastructure must be able to deliver functionality including policy, quality of service, and traffic prioritization, and to optimize resources in the process. This effectively means meeting the discrete needs of any application, any time, anywhere, and at any scale. Also, enterprises need solutions that are simple and automated enough for IT generalists to successfully manage. Infrastructure convergence can meet these demands.

There are two approaches to building infrastructure convergence – converged infrastructure (CI) and hyper-converged infrastructure (HCI). The simplest distinction between them is that CI is hardware component-based, and HCI is hypervisor-based infrastructure.

IT organizations often step up to infrastructure convergence when:

A data center infrastructure and data center refresh is due

It's time to upgrade/migrate enterprise applications (Oracle, SAP, SQL)

Initiating a hybrid cloud model and application self-service

Virtualizing workloads on VMware, Microsoft or OpenStack Hypervisors + 'bare metal'

Addressing branch operations

What is hyper-converged infrastructure?

Hyper-converged infrastructure adds tighter integration between more components through software, to the point that the solution cannot be broken out into separate elements. HCI is typically deployed on a single hypervisor through virtual machine translation technologies and APIs.

In hyper-converged infrastructure, the storage controller function runs on each node in the cluster as a service, which can improve scalability and resilience. All of the local storage across the cluster is configured as a single storage pool. Data that needs to be localized for speedier response can be kept locally, while lesser used data is stored on a server with spare capacity.

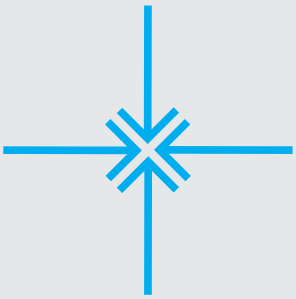

Who is buying HCI and why?

Hyper-converged solutions reduce the entire stack of products used to manage, store, compute, network, and protect data.

The primary advantage of a hyper-converged system is simplicity. A single capable IT person can manage a hyper-converged cluster, for a significant reduction in the cost of labor. This makes it a compelling choice for an SMB or SME environment. In addition, niche applications within the enterprise can also be well served by HCI within the existing environment. Ultimately, a hyper-converged infrastructure provides an efficient way to support third platform applications at the lowest cost per unit.

Simplicity Gained At The Expense Of Scalability And Control

HCI is relatively easy to deploy, with less expertise required to execute than standard infrastructure. And it recovers quickly from technical error. Since network and storage are all represented on each node, losing a node means simply moving on to the next one, rather than having to bring the failed node back to operability. Key use cases include application development environments and virtual desktop infrastructure.

When workloads are relatively predictable and fixed, HCI can get the job done. But because of its ultra-integration, expanding singular IT resources can be problematic, limiting flexibility and scalability. You can't expand storage without adding compute and networking. Simply put, you can't get one without the others, which is not always an efficient way to respond to emerging business needs.

In addition, HCI provides little to no means for tuning the storage and compute environments to optimize for a particular workload or environment. Many people refer to HCI as a solution for organizations that don't want to manage storage.

Converged infrastructure – what is it?

Converged infrastructure (CI), combines the four data center functionalities – compute, storage, networking and virtualization – into a single unit. This minimizes compatibility problems between storage systems, servers and network devices while reducing costs related to floor space, cooling, cabling and power. Storage can attach directly to the physical servers, through a network fabric. Flash storage typically supports high-performance applications and can also cache storage from an attached disk-based storage system.

CI can be supported by a single vendor for end-to-end support for all technology components. It also provides for a single end-of-support date, simplifying maintenance strategies. Through CI, organizations can save time and money on deployment, while streamlining their overall management by combining data center controls, and reducing the footprint within the data center.

Who is Buying CI and Why?

Typically, CI is used to support the big stuff – second platform applications including Enterprise Resource Planning, Customer Relationship Management (CRM), database grids, and enterprise messaging. Enterprises opt for CI to gain the best functionality elements such as blade servers, modular storage arrays, virtualized networks and multiple hypervisors, all interacting for optimum performance. In this way, CI delivers the best of compute, networking and storage through a predefined hardware and software solution.

SCENARIOS IN WHICH CI IS LIKELY TO BE THE IDEAL OPTION INCLUDE:

- When deploying new tier-one applications or when virtualizing applications
- When deploying private clouds
- When managing remote or branch office, or ROBO, infrastructure
- When managing remote disaster recovery sites
- When engaging in data center consolidation, whether within data centers or across them
- When implementing virtual desktop infrastructure, or VDI
- When performing data migration within sites or across them
- When testing or developing infrastructure

Are flexibility and scalability offered with simplicity?

CI consists of standalone components, which can be attached or detached as needed. Because individual pieces can be deployed when and where needed, flexibility is built into the model. At the same time, all the IT components are integrated into a single platform optimized for performance so that workloads are distributed and processed for maximum efficiency. Centralized controls keep this agility easy to direct.

In addition, CI is scalable over time, with its plug and play flexibility providing an on-demand growth model. As organizations evolve, they can build out specific parts of a virtualization stack, enabling them to quickly and easily grow the capacity to support emerging business initiatives. The key is to select a vendor who doesn't require you to sacrifice simplicity for the above benefits.

A Look Ahead: Why CI Better Prepares You For Tomorrow

IT investment today must consider the advancements of tomorrow. Organizations will increasingly pursue private clouds, enabled by software and hardware capabilities. Platforms that support open, heterogeneous and dynamic environments will be the most relevant and long-lived. For this reason, CI is the future-proofed approach to infrastructure innovation.

It's important to think through the full implications of the HCI model for the future as well. Because it is software-centric, a highly homogenous environment is created. This achieves cost savings through labor efficiency. But it also eliminates specialization on specific tasks such as storage or compute, creating the potential for inefficiencies of scale, by creating large amounts of unused resources.

In addition, there is the security challenge of including the control software on every node to consider. Because hyper-converged infrastructure combines data and control planes, an unwanted incursion anywhere in the system makes the entire system vulnerable, as you cannot partition off the areas that have been compromised. There is no benefit to network isolation if all data exists on all machines.

There is undoubtedly a place for HCI, ideally in small and medium businesses, and in larger environments for niche applications with limited scalability and security requirements.

Coming Soon: Composable Infrastructure

Converged infrastructure and software defined networking is giving rise to the approach of pooling physical infrastructure resources in the form of composable infrastructure.

In a composable infrastructure, hardware doesn't have to be physically configured. Instead, resources are allocated to software applications on demand, through an automated process. When a given application no longer requires that infrastructure, it's re-allocated to another application. This enables a data center to maximize physical infrastructure usage, while accelerating the deployment of new applications. Composable infrastructure is promoted as a means for IT departments to provision workloads with speed and efficiency comparable to a public cloud service provider, without giving up control over mission-critical application infrastructure.

Next-gen CI

CI continues to evolve, which will undoubtedly empower organizations with more and more possibilities. A range of new options will become available to IT teams, including such elements as GPU, RAM, and networking blocks. These advancements will enable converged infrastructure solutions to rapidly mold to virtually any workload profile as it emerges.

These innovations will support enterprises in further reducing IT costs while addressing emerging business objectives and market priorities. This will mean delivering solutions that:

- Simplify deployment and management of IT infrastructure
- Reduce the risk of capacity shortfalls
- Provide cost-efficient scalability
- Optimize infrastructure usage across in-house and service-provider infrastructures
- Avoid investment in capacity that cannot be used

Introducing FlexPod

FlexPod, a leading converged infrastructure solution from Cisco and NetApp offers many of the advantages of hyper-converged infrastructure without compromising control or scale. Available as both a reference architecture and as a factory-integrated solution, FlexPod leverages market-leading Cisco UCS servers and Nexus switches, along with NetApp All Flash FAS storage. Customers can expand a FlexPod in any dimension to evolve with changing business demands. And, FlexPod includes the granular control enterprises need to optimize their storage environment.

The FlexPod Advantage – Now and in the Future

Superior Performance

FlexPod provides consistent high performance delivery of fast and secure business critical applications. This means optimizing data center usage for the most demanding applications, while eliminating costly downtime and securing data throughout, including the cloud. This caliber of performance has real bottom line implications including enabling superior decision-making through better access to information, seizing increased revenue opportunities from faster business transactions, and increased overall response times to business-critical SAP, Oracle, and Microsoft enterprise applications.

FlexPod Advantage: **Performance**

208%

SQL RESPONSE TIMES THROUGH UCS WITH M4 PROCESSORS

- **100+ world record benchmarks for UCS**
  Including SPECCINT2006: SPECint®_rate_base=1760 Best 2-socket x86-architecture result, representing a 27.5% increase2
- **Up to 4 million IOPS with 1ms latencies on All Flash FAS[1]**
- **Up to 20x faster enterprise application performance on All Flash FAS1**
- **All Flash FAS delivers 685,281.71 IOPS at 1.23 ms average response time (ART)[2]**
- **All Flash FAS was rated #5 in SPC-1 “Top Ten” by performance, with the second best ART (1.23 ms) and SPC-1 LRT (0.48 ms) in the top 52**
- **4x improvement in storage IOPS 4x improvement in SQL Server CPU utilization[1]**

Agility on Demand

The business environment is more fluid than ever, and profitability demands infrastructure that can keep up. That's why FlexPod is designed to allow users to deploy new hardware and software in just minutes. Also FlexPod supports virtually any cloud strategy, a key component to increased efficiency, agility and investment protection. And, it scales seamlessly to meet the enterprise's emerging needs. FlexPod allows IT departments to deliver the responsiveness of public cloud, focus resources on high-value activities and make operational efficiency match the speed of the business. All of which translates into critical operational advantages that position our clients for growth.

FlexPod Advantage: **Agility**

83%

FASTER PROVISIONING WITH CISCO ACI ARCHITECTURE[1]

- **Ready for use in under 60 minutes with infrastructure automation**
- **Integrated and standardized components reduce application deployment from months to weeks**
- **20-30% reduction in application testing time[1]**
- **NetApp engineers stood up a complete VDI solution in just 12 days for ~10,000 virtual desktops with All Flash FAS FlexPod foundation[3]**

Superior Economics

With FlexPod's best of breed approach, the numbers add up. While raising performance, it reduces cost as well as complexity. This streamlined approach to IT and IT management lowers operating expenses by minimizing the need for specialized skills and the maintenance required to tune databases. At the same time, reduced power, cooling and cabling needs afford additional savings. FlexPod empowers the enterprise to focus resources where they are needed most in terms of business performance. This increases the reach and impact of business offerings as well as the customer service supporting it, while decreasing the risk of new deployment through fully tested and validated designs.

FlexPod Advantage: **Economics**

76%

ROI IN JUST 17 MONTHS*

*Source: Forrester TEI, 2015

FlexPod CI is the Answer

If you're ready for converged infrastructure, FlexPod simplifies deployment, reduces the risk of capacity shortfalls and provides cost-efficient scalability, while supporting virtually any cloud strategy. With all-flash speed, simplicity and innovation, FlexPod empowers your enterprise to meet the demands of a dynamic, future-minded business environment.

Learn more about FlexPod at www.flexpod.com and NetApp.com/flexpod

1. Savings compared with previously installed HDD-based ECB storage systems. Source: Gartner, "Solid-State Array TCO Reality Check," Joseph Unsworth and Arun Chandrasekaran, 22 January 2016.
2. Cisco, FlexPod Advantage: Performance, Agility Economics, 2016
3. Cisco, Cisco UCS Servers Claim Nine New World Records on Industry-Standard Benchmarks, 2016

© 2016 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, FlexPod, and ONTAP are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at http://www.netapp.com/us/legal/netapptmlist.aspx.