PURE STORAGE PRESENTS



Converged Infrastructure

Jon Toigo & Scott D. Lowe

INSIDE THE GUIDE:

- Discover the Power of Converged Infrastructure
- Learn how Converged and Hyperconverged Infrastructure Compare and Contrast
- Acquire Understanding of Convergence-driven Integration, Scalability, and Business Outcomes

HELPING YOU NAVIGATE THE TECHNOLOGY JUNGLE!



In Partnership With

··II··II· CISCO

THE GORILLA GUIDE TO...

Converged Infrastructure

AUTHORS Jon Toigo & Scott D. Lowe

EDITOR Keith Ward, ActualTech Media

LAYOUT AND DESIGN Olivia Thomson, ActualTech Media

Copyright © 2018 by ActualTech Media

All rights reserved. This book or any portion thereof may not be reproduced or used in any manner whatsoever without the express written permission of the publisher except for the use of brief quotations in a book review. Printed in the United States of America

ACTUALTECH MEDIA

Okatie Village Ste 103-157 Bluffton, SC 29909 www.actualtechmedia.com

ENTERING THE JUNGLE

| Chapter 1: Why Converged Infrastructure? | 6 |
|--|-----------|
| Additional Goals of This Book | 9 |
| Chapter 2: Understanding Converged Infrastructure | 12 |
| Moore's Law & House's Hypothesis | 13 |
| The Rise of Shared Storage (SAN and NAS) | 14 |
| The Software and Integration Revolution | 15 |
| The Converged Infrastructure Alternative | 21 |
| Chapter 3: Infrastructure Design Goals | 27 |
| How Converged Infrastructure Can Help Achieve Design Goals | 28 |
| A High Wall to Scale | 30 |
| Emerging Workloads Driving Converged Infrastructure | |
| VDI to the Rescue? | 35 |
| Case in Point: FlashStack | |
| Tight Integration | 41 |
| Chapter 4: Developing a Model Converged Infrastructu Solution | ıre 41 |
| FlashStack vs. Competitors | 43 |
| Growing Pains | |
| Course Correction | 49 |
| Chapter 5: The Future: Going HyperScale? | 49 |

CALLOUTS USED IN THIS BOOK



The Gorilla is the professorial sort that enjoys helping people learn. In the Schoolhouse callout, you'll gain insight into topics that may be outside the main subject but that are still important.

This is a special place where readers can learn a bit more about ancillary topics presented in the book.



Discusses items of strategic interest to business leaders.

ICONS USED IN THIS BOOK



DEFINITION

Defines a word, phrase, or concept.



PAY ATTENTION

We want to make sure you see this!

Why Converged Infrastructure?

Welcome to *The Gorilla Guide to Converged Infrastructure*. The purpose of this guide is to provide the IT planner or server administrator with background information that will assist them as they consider alternatives for hosting their mission-critical workloads and their data.

A byproduct of the software-defined revolution, which has seen servers, networks and now storage technologies deconstructed and virtualized, has been a proliferation of unfamiliar terminology such as *converged infrastructure* (CI) and *hyperconverged infrastructure* (HCI), although hyperconvergence is not the focus of this book.

Instead, we'll focus the discussion around converged infrastructure, which represents a robust area of technology development; wellknown and up-and-coming vendors alike are engaged in ongoing research and development, and constantly delivering new products to the market. CI products may come from an individual vendor or sold as a collaboration between multiple vendors.

These solutions may comprise server hardware, networking hardware, and storage hardware, plus software that provides management, automation and/or orchestration to produce an integrated system. Put more succinctly, CI typically consists of multiple information technology components that are integrated by their respective vendor(s) to form a single, optimized computing system.

There are many benefits that accrue from pre-integrated technology stacks when they're created with careful consideration to both

workload and infrastructure design requirements. Support for specific workloads can be designed into the solution, minimizing the need to fit the workload to what the stack offers. The dictum remains true: "one size fits most" infrastructure rarely fits anyone's workload very well.

IT planners usually need to satisfy general infrastructure design parameters as well as workload-specific ones. Key to these requirements is typically a technology's conformance with management standards. Management of a CI platform tends to be a more straightforward affair than the management of nonconverged resources, including such mundane tasks as ensuring the interoperability of components.

With a properly-defined CI stack, planners may also have greater certainty with respect to performance of the overall system at scale. There's enormous value in understanding in advance of deployment how technology layers are scaled in terms of capacity and performance, how resiliency is provided, and how services are delivered to workload as hardware nodes or other system parameters change over time.

In short, pre-integration of infrastructure simplifies workload hosting, improves the manageability of the infrastructure, and enhances its predictability. These characteristics of CI lend credence to the claims of CI vendors that their technology advances information technology in positive ways.

Here at the outset of this guide is a good time to offer a simple definition of converged infrastructure. CI is a concept of information system architecture in which the hardware and software components providing processing, networking and storage are coupled via hardware pre-integration and centralized software management, administration and orchestration.

CONVERGED VS HYPERCONVRGED INFRASTRUCTURE

We don't want to ignore hyperconverged infrastructure since it's a popular infrastructure choice. Although our focus will be on converged infrastructure, you should understand some key differences between the technologies. Whereas converged infrastructure is a conglomeration of first-class technologies that come together to create a whole that is greater than the sum of its parts, hyperconverged infrastructure clusters are formed through the combination of a series of indentical, standalone appliances that each integrates key infrastructure components.

With converged infrastructure, the goal is to provide simplicity, tuneable performance, massive scalability, and a single pane of glass for managing all of your infrastructure components. For hyperconvergence, the goal is to make it easier to scale workload environments in an economical way while driving almost all of the complexity out of the equation. The main reason is that HCI focuses exclusively on Ethernet as the glue that binds everything together.

Early iterations of CI were sometimes pretty complex beasts that had all kinds of different connections and required full-stack engineers to maintain. The current wave of CI products takes a page from HCI's playbook in terms of its reliance on Ethernet, making these products far more affordabale, desirable, and manageable than their early brethren. In this book, you will discover why CI is so important, and we'll even provide some answers up front:

- It fits right into current operating models
- It simplifies management of the full infrastructure stack
- Modern all-flash arrays have replaced the traditional storage portion in CI architectures, leading to incredible workload performance and price-performance gains
- Shared accelerated storage has gained significant traction in a marketplace that demands disaggregated resource deployment while maintaining a single point of management, such as the hyperscaler market
- CI's embracing of all-Ethernet architectures is reducing cost and complexity, making it a highly desirable option for organizations of all sizes

This book will educate you on everything you need to understand to determine the suitability for CI for your own needs.

Additional Goals of This Book

Of course, assisting readers in designing their IT infrastructure requires more than a casual definition of a term. This book will also discern use cases that illustrate common problems that can be addressed effectively by CI.

We'll discuss the impact that virtualization of workloads and consolidation of infrastructure are having on the contemporary data center; we'll also take a hard look at other trends that challenge IT planning, including unprecedented rates of data growth, to figure out the drivers of architectural concepts like CI.

In addition, we'll examine the rationale for deploying CI solutions, including the need for simplification of the infrastructure, cost

What is Converged Infrastructure (CI)?

of information А concept system architecture in which the hardware and software components providing processing, networking and storage are tightly coupled via hardware pre-integration and centralized software management, administration and orchestration.





reduction, improved performance of IT workloads and better agility in IT services. We'll also identify some of the strongest use cases for CI and best practices for implementing CI successfully to achieve IT objectives.

An important differentiator with respect to this book vs. others in the market is our effort to leverage real-world technologies in our reference models. We're fortunate to have technical advice from thought leaders in the CI market, including Pure Storage and Cisco Systems. We'll respect their insights in this book, though any references to specific products should not be viewed as endorsements.

So welcome to this book, and please feel free to contact the publisher if you want to identify something you believe was overlooked or incorrectly stated. There might be future editions in which we can include your insights and experience.

Understanding Converged Infrastructure

Recent research by Gartner and IDC analysts suggest that most of the infrastructure in enterprise data centers today is virtualized and software-defined, converged or hyperconverged. In one recent report, the claim was made that as much as three-quarters of business workloads are running as virtual machines (VMs).

The few remaining workloads operate without virtualization. These tend to be legacy applications delivering acceptable value in their given instantiation, or that have been deemed too costly to re-write, as they're running on infrastructure that would be too challenging to replace cost-effectively and without sacrificing performance. Taking that as a basic "state of the data center" foundation, we can review the history that has brought us to this point.

These data have been interpreted by some vendors to suggest that virtualization, software-definition and hyperconvergence are part of an inexorable or evolutionary trend, replacing less performant, less resilient and more costly legacy infrastructure. As many have discovered, this isn't necessarily true. For example, the reasons Fibre Channel fabrics and shared storage architectures persist, especially in support of non-virtualized workloads, are fundamentally because they continue to provide the best fit to the workload. CI can be viewed as an evolutionary development in these legacy technologies that improve on the foibles of "legacy" without discarding the architectural foundations of traditional IT altogether.

To understand the historical context of CI, we need to go back a few years—first to the factors that led to the undoing of the infrastructure memes that were introduced in the late 1990s, then to the directions that consumers and vendors adopted in the early 2000s (Don't panic—we'll make this review as brief as possible.)

The late 1990s saw the advent of several architectural models deemed innovative at the time of their introduction. Servers, modeled on single-core CPUs, ushered in an era of distributed computing that challenged the centralized computing model inherent in mainframe data centers.

Moore's Law & House's Hypothesis

Server makers capitalized on chip technologies that, in conformance with Moore's Law, doubled their number of integrated circuits roughly every 24 months, and doubled their clock speeds in accordance with House's Hypothesis about every 18 months. This gave rise to an era of distributed computing, aided by robust Ethernet and Internet Protocol networking standards.

What Is House's Hypothesis?

Most IT professionals are familiar with Moore's Law, which observed that the number of transistors in integrated circuit systems doubled, on average, every two years. This law was named after Gordon Moore, a co-founder of Intel.



Alongside Moore, David House, an Intel executive, predicted that the performance of computers would double, on average, every 18 months.

It's clear that these two statements are tightly coupled. They're sometimes interchanged with one another, and credit is generally given to Moore regardless of which statement is made. On the storage front, the 1990s saw a movement away from storage that was silo'ed (that is, directly attached to and isolated behind servers) and toward shared resource models.

Some storage systems featured multiple IO adapters that facilitated the sharing of on-board—mostly disk drive—storage resources among multiple client systems. Others featured a thin server bolted to the storage array to create a file server appliance often called Network Attached Storage (NAS).

The Rise of Shared Storage (SAN and NAS)

Toward the end of the decade, Storage Area Networks (SANs) became popular. Built on Fibre Channel, a serialization of the SCSI command language, SANs ran across fiber optic or copper cable "fabrics" (not networks in the strictest use of the term), and later on iSCSI, another serialization of SCSI developed by the Internet Engineering Task Force for use across TCP/IP networks. These SANs had the effect of creating a switched storage infrastructure separate from, and shareable with, all connected servers.

The thinking behind the introduction of a shared, "networked" storage infrastructure that scaled independently of a networked compute layer was a sound one with some significant benefits as well as a few challenges.

On the benefits front, shared storage systems excelled over direct attached storage (DAS) systems in terms of the ability to actually make use of all of an organization's storage resources. With DAS, storage assets were divvied up among numerous server islands, making it impossible to centrally leverage and allocate those resources. The emergence of shared storage in the form of SAN and NAS eliminated these silos and the techncial debt that accompanied them.

Over time, however, shared storage itself began to demonstrate some challenges. First, as a completely separate data center resource, it was managed through a dedicated management console. Early on, the benefit of shared storage was the collapsing of storage islands — and, by extension, the management tools for those individual islands — into a single administrative experience. Eventually, though, as organizations sought even more simplicity in the data center, this separate administration point became a point of concern, particularly since it was often provided by a vendor completely different than everything else in the data center.

Another issue was the introduction of value-add services or functions on storage arrays. De-duplication, thin provisioning, and other features, which were usually added as software instantiated on an array controller, added substantially to the cost to acquire and own storage gear; also, in many cases, they added to the challenges for centralized monitoring and management of the storage infrastructure.

By the early 2000s, workload virtualization and hypervisor computing took hold in the server world, in part as an effort to resolve some inherent stack interoperability issues. The idea was to consolidate workloads and operate them in a more or less uniform way on fewer servers, thereby reducing the likelihood of interoperability problems and support ease of deployment and administration.

However, workload virtualization added new stresses, in the form of consolidated workloads and their IO traffic patterns, that were visited upon SANs and NAS systems. These issues, plus the complexity and cost of storage, made it ripe for revolution.

The Software and Integration Revolution

With the world increasingly turning to software to solve complex technology challenges, it was only a matter of time before this trend hit the storage market. It really got a boost way back when virtualization hit the market. In fact, hypervisor technology was

MODERN SOFTWARE-DEFINED STORAGE IS A SECOND ACT

In reality, software-defined storage predates VMware, Microsoft Hyper-V, Red Hat KVM and the other hypervisor technologies by at least a decade, tracing its origins back at least to IBM System Managed Storage (SMS) in 1993. Still, VMware was ultimately credited with resurrecting the idea; they created a business value case that resonated with consumers and spawned the modern movement.

the tip of the spear in the move toward software-defined storage, which led us to converged and hyperconverged infrastructure.

By virtualizing the application and operating system used on ever more powerful servers as a result of Moore's law, and creating from them a VM that could be cut-and-pasted across servers running a hypervisor, VMware and its competitors found a way to consolidate workloads on a fewer number of servers.

One consequence of software-defining (i.e., virtualizing) servers and consolidating multiple workloads on them was the impact it had on LAN and SAN traffic. According to one vendor in the early 2000s (Xsigo Systems), the creation of 7 to 10 VMs on a single



Figure 2-1: Eliminating physical servers via virtualization



Figure 2-2: Virtualization introduced significant changes to server and storage IO patterns

physical host typically necessitated the addition of up to 16 LAN/ SAN ports per server, and resulted in the concentration of IO across networks and SANs.

This phenomenon created issues for IO efficiency in many shops, which were contextualized as the "IO Blender Effect." This expression refers to the sources of performance degradation and latency in the storage associated with virtual server hosts. The IO blender became something of a rallying cry for a revolution in storage.

Economics also played a key role. As previously stated, from the late 1990s forward, vendors had been adding features and functions in the form of software on monolithic arrays. This practice enabled them to differentiate their wares from that of competitors, to meet certain customer requirements, and to provide their products with room to grow in terms of profitability. One example was a storage array technology, de-duplication, which was acquired by a leading vendor and added to a controller of one of their arrays. The problem was that vendors were charging exorbitant sums for some of these features, leading to a situation that was off-putting to many customers.

By leveraging this antagonism between customers and their legacy storage hardware vendors, and adding in the story of the IO Blender, the stage was set for a storage revolution. Vendors made a compelling



Figure 2-3: The life of a raw IO in a storage system

case that, just as hypervisor technology had reduced server brands to commodity kits, so too would software-defined storage reduce storage to a commodity resource.

This would be done by migrating most or all of the feature/function software currently hosted on an array controller to a software-defined storage (SDS) stack running on a server.

The Emergence of Convergence

At about the same time, some vendors began "purpose-building" hardware/software stacks behind specific workloads. Some worked to develop standardized reference models to make implementation less complex—similar to the approach of the HCI vendors, but with more individual resource flexibility.

Their innovations in terms of storage abstraction, provisioning and management gained these products the collective name of "converged infrastructure," and their performance and scalability, which generally surpassed those of HCI, made them the preferred choice for more demanding workloads.

These advances also changed the storage administration game. In fact, according to analysts, the average amount of storage measured in raw terabytes that could be managed by a single administrator had increased from 132 TB to 344 TB over a five-year time frame.

This advance has come just in time.



Figure 2-4: Hardware abstraction in action in a software-defined storage system

Storage firms are already hard at work planning and designing storage infrastructure in anticipation of a data deluge that analyst firm IDC recently pegged at approximately 163 zettabytes of new data by 2024. Given that the entire annual output of the flash and disk storage industries will deliver less than 2% of the required capacity to store all the new data, there will be a continually increasing need to maximize the activities for infrastructure professionals that have to manage storage.

This need puts converged infrastructure in a highly desirable position thanks to its ability to simplify data center infrastructure administration while also providing incredible scalability potential.

But as you've heard already, converged infrastructure has a competitor in HCI.

Hyperconverged Infrastructure Enters the Scene

In recent years, an entire class of appliances, called Hyperconverged Infrastructure or HCI appliances, appeared in the market. HCI cobbled together a server (running a hypervisor), an SDS stack (generally managed either through a virtual storage appliance [VSA] or a hypervisor kernel module),



and a storage array directly attached to the server. These appliances were touted as "building blocks" for future software-defined data centers.

HCI has certainly created something of a firestorm and, for the right workloads, is a great option, but it's not without its hurdles. Most notably, there are some challenges surrounding operation and management at scale. In part, this is a reflection of the problems inherent in isolating storage behind a particular service hypervisor head. A recent analyst report suggested that, with the advent of SDS platforms in which storage is directly attached to servers and controlled by the server, a good news/bad news situation was developing².

Like all challenges, though, hard work is being done to overcome them. Many HCI vendors are continuing to evolve their products, endeavoring to create a means to drive their clusters of HCI appliances to web-scale and beyond. This feat was previously thought to be the antithesis of HCI.

HCI was originally conceived as a deconstruction of the layered relationship between server and storage. What some up-and-coming HCI vendors are finding is that the storage layer does, in fact, need to be decoupled from the server layer if it's to grow to web scale, and if the storage is to be fit for use by different workloads with different needs.

To phrase it more simply, HCI has emerged as an option for organizations that want or need to simplify the stack so that a single operations team can deploy and manage the stack from top to bottom.

Considerable work is being done to develop the ultimate HCI solution, one that acts as an atomic building block of a scalable platform. Time will tell whether these efforts pay off. For the time being, most HCI product offerings appear to fit certain use cases, such as archive or backup or file/object storage, reasonably well, depending on workload dimensions and SLAs.



Figure 2-5: The building blocks of a hyperconverged infrastructure system

The Converged Infrastructure Alternative

Some might say that CI isn't as ambitious as HCI, but that would be a mistake: the two technologies address wildly different needs. In some respects, CI carries over the notion of "purpose-built" systems from traditional infrastructure. CI architects have made engineering decisions about where best to host various storage software elements, and they've emphasized performance, workload awareness and support, and network flexibility over cookie-cutter simplicity and Lego-like building block deployment, which is where HCI shines.

The focus of CI development has been on consolidation. CI developers have embraced the idea of centralized management of IT resources on an enterprise-wide basis, rather than on a per-system basis.

Instead of focusing only on pre-defined and integrated hardware/ software offerings, the goal of CI designers has been to create resource pools—servers, networks, and storage—that can be shared



Figure 2-6: Cost for storage is decreasing but, in some ways, inefficiency is creeping in, which, over time, can increase total cost of ownership

by multiple applications and managed collectively using policybased automation.

This is partly because the goal of early CI pioneers was to eradicate IT sprawl, which they saw as stifling agility and increasing CAPEX and OPEX. But it's also, in part, a response of the vendors to the requirement underscored in analyst reports that showed up to two-thirds of annual IT budgets in large organizations going to operations and maintenance, rather than new IT initiatives that might benefit the business.

The immediate problem that the CI mavens sought to address was to find a more efficient way to share pools of resources efficiently to build the various system infrastructure that workloads required. Instead of "one-size-fits-most," CI architects went for flexible frameworks that could adjust and scale for a customized fit to the workload.

By pre-integrating hardware where appropriate, CI designers sought to tackle cost and complexity issues. They virtualized networks and FC fabrics to enable scaling as needed, and provide for easier centralized management. They took delight in referring to their infrastructure designs as "cloud ready."

A Zettabyte Apocalypse in the Offing?

Data Growth Projections Are Increasing



Originally pegged at 10ZB by 2020, IDC revised its projections of data growth to 60ZB by 2020, then to 160ZB by 2024. Given annual output of the flash industry (measured in total storage capacity) of approximately 500 exabytes, and of the HDD industry of 780 exabytes, total capacity will fall well short of any of those three demand projections. This happens to be one reason why tape technology is enjoying a bit of a renaissance at present.



CI has been widely regarded since its inception as the enabling platform for private and public clouds. Efficient resource pooling with policy-based allocation and management are must-haves in the context of dynamic cloud data center operations, and CI offers these advantages.

The tradeoffs are several, of course. Some CI frameworks and management layers in CI designs are proprietary, and perpetuate older tendencies of hardware vendors to use *de facto* standards to differentiate products, lock out competitors and lock in customers.

This is not dissimilar to the efforts of hypervisor vendors to use their presence in HCI appliances to silo entire stacks of compute, networking and storage behind a particular hypervisor model. Given data showing that since 2015 IT shops have been diversifying their hypervisor vendors to avoid lock-in, the infrastructure produced in that model is likely to be even less efficient than the infrastructure produced using CI stacks.

Going forward, the real challenge will likely be how to bring disparate elements of converged and hyperconverged infrastructure under common management. That may ultimately provide the case for cognitive data management.

In the next chapter, we'll look at common infrastructure design goals and how CI can be put to use to make IT services simpler to deploy and easier to manage. We'll also look at how CI might align to common use cases in business verticals.

Infrastructure Design Goals

When designing IT infrastructure, the first step is to determine what we're trying to achieve with that infrastructure. Clearly, applications themselves dictate some of the parameters of the design.

Processing, networking and storage components must work together in a manner that best suits the workload, the profile of the input and output of data from the application, and accesses that are made to the application and its data by other applications and by users.

Just as important as the matching of infrastructure to application requirements, planners need to think about a broader set of requirements. Usually first among these is cost containment.

IT is expensive, and business managers in organizations often seek to assure themselves that money spent on IT is just as scrupulously accounted for as money spent on other aspects of business operations. They want to be assured that what's being purchased is the right technology for the job; the right investment from the standpoint of internal rate of return; is protected from early obsolescence; and that it doesn't unnecessarily lock the firm into a particular vendor.

The infrastructure choice should also contribute to an improvement in the business's profitability by increasing productivity, minimizing errors, enabling expansion into new markets, etc.

Cost-containment, risk reduction, and improved productivity are the three components of a traditional business value case that all IT planners must address in their proposals to senior management whenever they're seeking funding for an IT initiative. Infrastructure is no exception.

Business Value: A Three-Part Narrative

Senior management may not understand performance metrics or other technical dimensions of converged infrastructure. However, they do understand (and expect) that any proposal for technology acquisition must offer a compelling narrative in the the areas of cost containment, risk reduction and improved productivity. The good news is that CI has a great story to tell in each category.

In companies that have embraced cloud philosophy as a context for all IT-related activities, infrastructure technology must also fit the broader objectives of agility, elasticity and resiliency. In practical terms, this means that the infrastructure:

- Should be efficient in terms of use and management, and capable of being deployed or provisioned rapidly in response to fastchanging business requirements (i.e., agile)
- Must scale with workload automatically, if possible (i.e., elasticity)
- Must be sufficiently redundant and otherwise imbued with availability and continuity guarantees to ensure availability throughout the lifecycle, and that users can be assured of its permanency (i.e., resilient)

Experience also guides good infrastructure design. Experienced infrastructure planners know the value of standards, both as a hedge against vendor lock-in and limited flexibility in scaling options.

Compliance with open standards (*de facto or de jure*) should be front of mind in technology selection. Conversely, proprietary technology that limits design options or sources of supply or that commits the design to the whims or financial fortunes of a particular vendor should be avoided if possible.

Given the tribal nature of many technology markets, designers need to guard against "marketecture" and eschew analyses and opinions that don't pass dispassionate, logical tests of veracity or authenticity.

Case in point: in the early days of software-defined storage, the case was made by hypervisor vendors that the reason applications newly encapsulated in VMs were less performant was IO congestion. However, this story was rarely supported by objective data, such as measurements of IO queue depth. A designer needs to test such assertions before selecting a design.

Most importantly, designers need to consider management and administration. Analysts have repeatedly touted Total Cost of Ownership (TCO) graphs indicating that the cost to acquire technology is only a small fraction of the cost to own and operate technology.

Storage provides a good example. Depending on the analyst one reads, storage hardware represents between \$0.33 and \$0.70 of every dollar spent on IT hardware generally. However, this acquisition cost is only about 20% to 25% of the annual TCO of storage infrastructure. Most of the cost of storage is not found in capital expense (CAPEX), but in operational costs (OPEX) that range from management and administration costs (labor) to backup (labor) to environmental and facility expense.³ However infrastructure is designed, both CAPEX and especially OPEX cost containment should be kept in mind. In the case of management and administration, the quest is for greater simplification and greater automation.

³ Storage TCO is a much discussed and debated concept. One of the best guides for evaluating storage TCO is provided here: https://www.business-case-analysis.com/total-cost-of-ownership.html. Since that time, Gartner and IDC have repeatedly published reports on the TCO of various platforms from a CAPEX and OPEX perspective. Invariably, their findings confirm that the annualized cost of ownership from an OPEX perspective is between 4 and 6 times more than the cost of CAPEX (acquisition and deployment costs).

Simplification is a mantra in contemporary infrastructure design. Standardization is part of the story; using standards-based technologies usually supports (but doesn't always guarantee) the interoperability of standards-compliant components, regardless of the source. This, in turn, may benefit rapid deployment, ease of management, and smoother, more predictable scaling.

Automation is another holy grail. When developing new technology, most vendors try to reduce the number of manual steps taken to perform routine tasks with their products or the infrastructure which they help to form. Management is key: planners should seek to ensure that monitoring and management of all core infrastructure components is possible using a standards-based approach such as RESTful APIs. This is the foundation of automation.

Then, designers should pay attention to how vendors are automating functions or processes known to be labor-intensive—migrating data across storage tiers, for example. Automating such tasks makes them less prone to human error (in theory, at least), and also makes possible the administration of infrastructure with fewer staff. With IT staff sizes reduced as a function of cost-containment or "cloud strategy" or a shortfall of available qualified resources in a given geography, automation may be essential to achieve any sort of operational efficiency.

How Converged Infrastructure Can Help Achieve Design Goals

Converged Infrastructure, in addition to its architectural nuances, can actually be viewed as an effort to realize or embed the design goals enumerated above in a platform that can deliver *business value* (costcontainment, risk reduction, and improved productivity) and *cloud enablement* (agility, elasticity, and resiliency) right out of the box.

Vendors offering CI typically contextualize their value in terms of *digital transformation*. This is marketing code for IT simplification,

moving to cloud concepts in data center design and information service delivery, standardizing remote office computing (the mainstay of first-generation HCI) with its plug-and-play appliance model), and moving to standards-based networking, mostly IPbased.

Analysts such as Enterprise Strategy Group have been evaluating the trends in CI taking place over the past eight years. Initially, vendors were responding to complaints from customers about the complexity associated with managing three or four separate vendors, each providing products for different infrastructure layers that often times fell out of synchronization with one another as products were modified or improved. Managing multiple products from multiple vendors was a challenge, and vendors responded by joining together in ecosystem relationships and pre-integrating products for release as a one-stop shop.

Early CI integrations were still operated like hardware-defined infrastructure, but challenges such as patch management became less painful for consumers. HCI arguably went further toward software-defining the integrated compute and storage stack but not necessarily the network component—for top-down management. First-generation efforts, however, introduced other sorts of problems in many HCI environments, including impaired scalability and inefficient resource utilization, on an enterprisewide basis.

HCI received a lot of media attention because its appliance model attacked the challenges that organizations were having with complex infrastructure at the edge of the business—complex server and storage topologies were too expensive and complex to deploy in branch offices and remote offices where IT staff resources were generally minimal.

HCI also simplified networking requirements by eliminating "northsouth" traffic (communication between layers of servers in LANs and storage devices in SANs) and replacing them with "east-west" traffic (traffic between appliances in a LAN, much like network-attached storage of the prior decade). However, with all of the notice paid to small-scale implementations of HCI, CI—which was being drafted to do the heavy lifting of core data center application and data hosting—tended to fall off the radar.

But given the priority placed on manageability as a foundation for realizing the simplification and automation objectives of contemporary on-premises and cloud-based data centers, CI has been the development effort to watch.

Because CI perpetuates the idea of independently-scaled technology layers, it needs to do a more efficient job of serving up resources to applications across a potentially complex north-south internetwork of component layers. This has taken modern CI into directions that HCI didn't really address, with management and orchestration being a priority for CI developers.

A High Wall to Scale

In comparing CI and HCI, it's clear that there are different scaling goals. CI aims to create highly scalable and eminently flexible pools of infrastructure with the ability to scale individual resources either discretely or in lockstep with one another, all within the confines of the centralized management portal. HCI focuses on bite-sized linear resource scalability, and may not always be able to reach the scaling heights of its CI cousin.

The FlashStack[™] offering from server and networking giant Cisco Systems and storage maker Pure Storage represents what can be done via vendor cooperation to create a highly efficient and manageable CI offering. We'll look at FlashStack in greater detail in the next chapter, but for now it's useful to note that products like FlashStack help to address the needs of the on-premises/cloud data center, especially with respect to platforming "Tier 1" or "SLAdriven" application workloads.

IT planners designing infrastructure for these applications have prioritized performance, reliability and scalability in their design criteria, and are looking to more established vendors working with innovative partners to deliver products that they deem to be "enterprise class."

In FlashStack, Cisco plays the role of the tenured partner, bringing considerable innovation developed over time in their hardware and software offerings. Pure is the innovative storage newcomer with the latest flash storage technologies. Other CI offerings tend to come from similar collaborations of established vendors and newcomers.

Consumers tend to favor the notion that the combination of vendors will help cushion them from the pace of technology change and its potentially disruptive impact on the infrastructure cobble that they deploy from the team. Remember that CI technology layers are less tightly coupled than HCI, so the components at different layers can evolve and scale independently of other layers—an oft-cited problem with legacy infrastructure and something that the new collaborations between old and new vendors need to address.

In the case of FlashStack, Cisco and Pure Storage can host different application deployment models, including bare-metal applications, VMs and the very much in the zeitgeist containers, provisioning to each the appropriate system, network and storage resources required. Considerable attention is being paid to automating the provisioning of these resources and to monitoring performance as a precursor to simplifying and automating administrative tasks.

This was arguably what the "software-defined revolution" was originally all about: driving cost and complexity out and improving the management and automation of legacy infrastructure, especially behind application workloads considered to be mission-critical. Not surprisingly, the primary inroad for software-defined technology in CI has taken the form of improved management and orchestration functionality, rather than the virtualization and unification of server and storage stacks as in the case of HCI. Interestingly, this maps with current trends in workload development.

Emerging Workloads Driving Converged Infrastructure

In the case of FlashStack and other CI stack offerings in the market today, the key drivers are, as noted previously, the improvement of performance, reliability and scalability, particularly with respect to Tier 1 or mission-critical business applications. Not surprisingly, many of these applications are database-driven and instantiated on bare metal rather than operated in VMs or other abstract containers. Even these so-called legacy workloads or systems of record are changing, however.

With the advent of flash memory storage, for example, in-memory databases (IMDBs) are becoming a hot topic in the industry. The core idea is that the performance of a database can be improved, potentially by several orders of magnitude, by eliminating calls to data on slower storage devices and by hosting all data in system memory (DRAM), or in flash storage directly adjacent to system memory, or in a combination of both using DRAM, flash and possibly NVDIMM (which is a hybrid of DRAM and flash storage).

New protocols are being developed to facilitate IMDBs by enabling low-latency internetworking of on-server and on-storage flash capacity for the purpose of capacity scaling. NVMe-over-Fabric and several other IP and Fibre Channel protocols are currently in the works that will likely find their first deployment in CI infrastructure.

NVMe flash, as well as the investments that companies will need to make in high speed/high capacity LAN and SAN interconnects, will



Figure 3-1: A modern application hierarchy diagram Source: The Data Management Institute

unleash the real potential of flash storage going forward. They're seen as the next logical evolution in cloud data centers.

Infrastructure similar to that supporting IMDBs will likely be required for next-generation business analytics as well. **Figure 3-1** depicts a contemporary mobile commerce hybrid data center, showing the major subsets of infrastructure as deployed today.

"Systems of Engagement," applications for order-taking from PCs, tablets, smart phones, etc., are used by customers to peruse offerings and make orders from catalog systems. These so-called "Systems of Interaction" tend to be instantiated as VMs, either in an on-premises server farm or in a cloud provider facility.

These orders are shunted first to "Systems of Insight"—analytics servers clustered and running technology such as MapReduce—so that orders can be scanned for authenticity and upsell opportunities. Once approved, the order is sent to legacy databases that actually control inventory, shipping, financial transaction processing, customer relationship management and resource planning: the so-called "Systems of Record."

While much simplified, this illustration shows multiple opportunities for CI evolution in response to workload requirements. For example, in the not-too-distant future, Systems of Record may require much more performant compute, networking and storage infrastructure; if for no other reason than to handle the so-called *Starburst Effect* in which one approved transaction in the Systems of Interaction generates hundreds of transactions for the Systems of Record.

Similar technology will likely be introduced into the Systems of Insight to expedite the processing of order data, to authenticate purchasers faster and "upsell" them with related merchandise.

IMDBs may be receiving the lion's share of attention currently, given the interest in the technology from both large enterprises and clouds and the high-performance computing market, but CI's original claim to fame and an ongoing source of interest in the model continues to be its use in consolidating devices. Sometimes pundits and analysts conflate consolidation and virtualization, though they're quite different in terms of their goals.

Virtualization provides a means to consolidate servers in the form of logical abstractions that can be hosted in a multi-tenant server environment. In that way, virtualization focuses on maximizing the value of technology through increased automation, scalability and support for multiple services from a single platform.

Consolidation is different in the sense that its mission is to reduce the cost of a technology by improving its operating efficiency and effectiveness. The mission of data center planners has always been to consolidate infrastructure relentlessly, driving down cost and complexity and standardizing core technologies for ease of management and automation. CI continues this mission, while HCI arguably, finds its solution for cost and complexity by using virtualization. One issue with virtualization is proving to be the proprietary nature of leading hypervisors, and the inability to share resources across different software-defined stacks organized under and controlled by proprietary hypervisors. CI tends to be more hardware-defined, hence more flexible and workload agnostic.

This doesn't mean that CI has no connections to virtual computing. In fact, a mainstay of current CI—an oft-demonstrated use case is virtual desktop infrastructure, or VDI. VDI is a virtualization technology that hosts desktop operating systems in a centralized server environment in a data center. The term, for which VMware is generally credited as coining, sometimes goes by the moniker "server-based computing."

VDI to the Rescue?

Originally, the idea behind VDI was to save money by eliminating a large number of desktop computers. Instead, virtual images of desktops would be stored on a centralized server and downloaded and used as needed. This strategy was also seen as an enabler of the mobile workforce, since users could download their personal desktop from virtually any location where they could access a network.

The touted cost-savings of VDI evaporated quickly, in many cases, when the costs for data center infrastructure and administration to support the strategy were factored in. CI saw VDI as a target workload early on, and has used its strategy for consolidating the administration of separately-scaling compute, networking and storage layers as a means to drive down costs and improve the performance and scalability of VDI.

Heathcare is yet another workload generator seized on by CI developers. Electronic Heath Records (EHRs) are a fast–growing type of storage, incorporating a patient's medical history, diagnoses, medications, treatment plans, immunization records, allergies, imagery and laboratory/test results. They're among the most regulated data in the world, and many countries have established regulatory regimes to govern the manner and duration of their retention.

More and more healthcare providers are seeking to streamline health data exchange with other hospitals and healthcare facilities within growing private networks or between their networks and partners, thereby enabling more complete and longitudinal patient health records and better-informed clinical decision making.

Getting to success requires the elimination of the many silo'ed healthcare systems that have been created around specific software/ hardware platforms over the years. Epic EHR and Cerner are just a couple of the software companies positioning themselves as EHR legacy replacements, with Epic customers reporting that they're hitting their financial targets with the deployment of the new system. Chances are good that the Epic EHR will be hosted on CI, such as FlashStack from Cisco Systems and Pure Storage.

The CI approach to hosting EHR systems makes sense because of the performance, agility and resiliency that can be brought to bear to

Fit to Workload and Fit to Business

CI Meets Design Goals

EXECUTIVE CO

Infrastructure planning must serve three masters: (1) the business, which controls the budget and sets practical limits on choices, (2) general IT goals regarding standards, manageability and suitability to staff skills, and (3) suitability or fitness to workload. Converged infrastructure addresses all three groups of design criteria for core applications. healthcare workloads running on a CI platform, and because of the platform's support for high performance transaction processing, analytics, and scaling in response to the massively-growing patient healthcare data explosion.

In the final analysis, "fitness to workload" determines the suitability of infrastructure. But as the above use cases suggest, infrastructure must also satisfy criteria that are business-savvy and cloud-aware. In the next chapter, we'll take a closer look at a CI stack.

Developing a Model Converged Infrastructure Solution

In the previous chapter, we outlined some design criteria becoming commonplace in CI planning. Some criteria were pragmatic and had to do with the suitability of CI to a specific type of application or workload.

We noted that CI generally aligns with the requirements of Tier 1 workloads, involving bare-metal databases and transactional processing. However, we also noted that in-memory databases and virtualized workloads, including VDI and EHR, were also finding their way onto CI as a means of achieving consolidation and cost containment.

We noted that, in addition to application/workload-specific design criteria, there were perhaps two additional groups of criteria that designers should consider when making infrastructure decisions. One group of criteria is associated with general IT practice and service delivery. The IT infrastructure selected needed to satisfy general goals of improving IT service agility, elasticity and resiliency. These criteria go to the "cloud enablement" or delivery of "IT as a Service" that you may have read about in the trade press or heard discussed by analysts or thought leaders.

Finally, there's the business value criteria that must be met by the infrastructure design. Management may not understand technical metrics or nuances, but they do understand the standards they

apply to evaluating all requests for hard-to-come-by budget dollars. Before they release these monies for use in infrastructure development initiatives (or anything else), they want a compelling narrative that explains how the initiative will help to contain cost, reduce risk and improve productivity or profitability.

With these design criteria in mind, we're ready to look at CI models, beginning with a leading candidate: FlashStack from Cisco Systems and Pure Storage.

Case in Point: FlashStack

FlashStack is the tradename of a CI solution developed jointly by Cisco Systems and Pure Storage. It offers best-of-breed components, pre-integrated and modular in design for scale-asyou-grow deployment to host a broad set of data center workloads. It's delivered, implemented and supported by FlashStack resellers and can enjoy single number support from Cisco Solutions Support for Critical Infrastructure.

The collaborators, Cisco and Pure, have developed reference architectures for deploying and supporting their CI in conjunction with Oracle, SAP, Microsoft SQL Server, VMware, Citrix and Exchange. They're working with additional technology partners to become certified for a broad range of additional mission-critical workloads.

The FlashStack CI can scale from two server hosts to more than 160 without disruption or downtime, and its storage can scale independently of the servers from five terabytes to more than a petabyte and a half. The flash array technology recently announced by Pure enables tremendous density, achieving the high end of the storage range in a 3u footprint.

The resulting infrastructure meets the "cloud-enabled" criteria readily. It's agile: easy and non-disruptively deployed. It's elastic: scaling is simple and servers and storage can scale independent of each other, in response to application demand. And it's resilient:



Figure 4-1: A rendition of FlashStack's integral components

workloads remain online as infrastructure is scaled or maintained. Typically, capacity upgrades require less than one hour, during which the infrastructure can continue to process workload IO.

As for business criteria, FlashStack again makes the grade. From a deployment and administrative cost-containment perspective, FlashStack is demonstrated to provide a 38% lower TCO than competing stacks, including private and public cloud offerings such as AWS, based on models created by Cisco. The primary sources of savings are three-fold:

- FlashStack delivers lower OPEX due to simplicity of design and outof-the-box automation of common storage management tasks.
- FlashStack delivers lower CAPEX because of efficient equipment featuring high density, low software costs and built-in data

reduction technology that enables more data to be stored in a fixed amount of capacity.

 FlashStack further limits CAPEX with its "upgrade in place" design and evergreen storage model. When refreshes of the technology occur (generally, every 3 – 4 years in enterprise infrastructure), the modular replacement strategy of FlashStack dramatically reduces the cost of upgrades.

The certainty of cost is a major selling point of FlashStack, according to evangelists and customers. So is ease of management, which contributes to the risk reduction value of the CI.

With FlashStack, there are 20 times fewer elements to manage than competing infrastructure alternatives or virtually any hyperconverged infrastructure solution with the same scale or capacity.

Tight Integration

That translates to fewer cables, server, storage and network devices, software packages, etc. Moreover, the infrastructure is tightly integrated with common management tools, including VMware vCenter, vRealize, UCS-Director, OpenStack and others. FlashStack provides a comprehensive and clearly documented set of management APIs that enable further integration and customization with whatever management utilities the customer decides to use.

Management is the lynchpin of proactive downtime avoidance, and a critical component of risk reduction. As noted previously, FlashStack also avoids downtime by design, and through the use of stateless technologies. Server components from Cisco Systems come with UCS Manager service profiles.

This type of architecture allows the creation of workload deployments where identity is abstracted from the underlying physical hardware. Cisco Unified Computing System (Cisco UCS) hardware is defined within a service profile, while the Pure Storage FlashArray is likewise stateless and resilient.

Using simple tools, you can configure how, where, and when workload instances are deployed. These instances can be VDI, an Oracle database, an Exchange environment, SAP/HANA, or a range of others. Stateless technologies enable administrators to configure MAC, World Wide Name (WWN), Unique Universal ID (UUID), boot details, firmware, and even basic input/output system (BIOS) settings in software, through simple management interfaces. With a stateless architecture, you create the industry's most agile converged infrastructure.

Also related to uptime optimization and risk reduction is the design of the underlying storage layer: FlashArray technology from Pure Storage. Flash modules can be swapped for higher-capacity modules the same way blades can be swapped for new CPUs/greater memory dual in-line memory modules (DIMMs) (as new technology becomes available, older modules can be replaced by shelf evacuation). No downtime, no complete equipment upgrades; just pure efficiency around workload delivery and scale. This further insulates the business from the risk of early technology investment obsolescence.

Improved productivity—the third business value criteria—is enabled by the simplicity and performance of FlashStack. Numerous customer testimonials available on the Pure Storage website provide snapshots of the specific performance improvements yielded by the implementation of FlashStack.

Generally speaking, FlashStack architecture creates a highperformance infrastructure for the most demanding users and workload profiles. More work is accomplished by reducing latency and introducing powerful metrics around workload performance with the Cisco Unified Fabric coupled with a Pure Storage all-flash array. These workloads can range from VDI to high-performance database systems. Performance also means resiliency. Users can have confidence that the app they need will be available when they need it. With FlashStack, enterprise deployments can enjoy a highly-resilient CI environment, which enables non-disruptive upgrades and 100% performance, even if a single component or path experiences a failure.

Smooth upgradability without downtime also allows FlashStackbased data centers to take advantage of technology changes (such as higher-performance blade servers) without any penalty and without downtime. Furthermore, in a virtualized environment, the process doesn't require a reboot of the application or VM. In fact, the virtualization layer (VMware) can reposition and replace the server during the transition using a vMotion operation.

Finally, the power of an all-flash stateless ecosystem will allow your applications to use sub-millisecond performance, for realworld efficiency gains and optimal user experiences. This is driven by reliable, award-winning systems: Cisco UCS servers are consistently the highest-performing in their class, across a broad spectrum of workloads.

UCS has set more than 100 world record benchmarks with their UCS architecture. Furthermore, every design has been comprehensively tested and documented by Cisco engineers to provide a deployment guide and best practices to help ensure faster, more reliable, and more predictable deployments.

FlashStack vs. Competitors

FlashStack isn't the only game in town when it comes to CI, of course. When it was originally announced in 2014, the first hurdle was to answer those who wanted to know why we needed another kind of infrastructure, rather than a new reference architecture.

Reference architectures are primarily whitepaper exercises that describe configurations of hardware and software that worked in a given implementation. Often, a reference architecture is simply a set of documentation that shows how to cobble together technology components that delivered required capabilities in a laboratory, proof of concept, or application testing environment.

Reference architectures generally aren't supported by committed cadres of product vendors with either integrated product or onethroat-to-choke warranty and maintenance agreements. There's also no guarantee that the reference architecture will show the same value in a different implementation context, nor that it will scale consistently.

Perhaps the biggest competition to FlashStack comes from HCI. To be certain, HCI is a worthy competitor in certain situations, but there is something of a Venn diagram scenario that takes place as you compare the two technologies.

FlashStack works in most organizations given a baseline level of financial and staffing resources. HCI shines when budget is very limited or staffing is very lean.

On the scaling front, modern HCI solutions can go far, but cannot reasonably achieve the same kinds of capacities as FlashStack without introducing some potentially significant internode networking challenges. That said, HCI can still scale well, but not to the levels of FlashStack.

At first, you may think that the single-vendor nature of HCI would lend itself to a more streamlined support experience, but with the way that Cisco and Pure have implemented a single support structure, that perceived advantage is erased, resulting in a tie between FlashStack and HCI on this front.

As always, though, the market is constantly changing. New, selfstyled HCI offerings are beginning to appear in the market that represent themselves as platforms, rather than appliances. They are, arguably, doing a better job of creating managed scale-out capabilities, mostly by mimicking Google's server node scaling architecture (to scale, just add more nodes). Some new HCI vendors are also limiting their use cases to "secondary storage" (backups, file and object storage, etc.) and excluding SLA-driven, performancehogging applications.

But these kinds of applications are exactly where CI solutions such as FlashStack truly excel. CI is looking more and more like a trustworthy method for building future infrastructure. Perhaps the greatest contribution of the CI stack vendors is the attention that is—at last—being paid to management and administration up and down the hardware stack. This has always been the Achilles' Heel of IT infrastructure, and is even more important to resolve as business leaders begin to believe the marketing of cloud companies and expect meaningful improvements in agility, elasticity and resiliency from "IT as a Service."

The Future: Going HyperScale?

Years ago, at the dawn of contemporary business computing, there was speculation that, eventually, there would only be a need for a handful of data centers. The idea of every business maintaining its own data center and IT infrastructure was about as far-fetched as everyone wearing a computer on their wrist or carrying one in their pocket.

The iPhone, and later, Android devices, proved the latter assertion incorrect. However, since 2016, public cloud vendors and their technology suppliers have been reconsidering the quick dismissal of the claim around data centers. Vendors and analysts projected that by 2020, the preponderance of enterprise computing would rely on just 485 data centers distributed around the world. Hyperscale data centers were on the rise and would soon account for 92% of workload processing.

For vendors of enterprise storage systems, the news did not depress earnings projections, but instead encouraged them. Those hyperscale data centers were, after all, going to have to store between 60 and 160 zettabytes of new data. So, even in a consolidated market, the future was bright.

Hyperscale isn't just a marketing term for "really big scaling." It refers to the ability of a system architecture to scale appropriately and independently as increased demand is added to the system. This is part process and part product. From a process standpoint, the human and machine processes must be in place to provision and add compute, networking and storage resources to the nodes of the system. From a hardware perspective, the infrastructure—conceived as integrated but independent layers—must be designed to accommodate this independent scaling to what might be considered today ridiculous or absurd levels.

Some vendors approach this challenge with a specific use case in mind, like multi-cloud integration. Other vendors, either individually or in groups, are puzzling out the requirements for building "hyperscaler storage" to keep storage alive as an industry and to enable the future of IT.

Growing Pains

Thus far, the resources available from the Storage Networking Industry Association (SNIA) are limited to case studies and white papers featuring do-it-yourself projects or proofs-of-concept undertaken by larger firms with hyperscale concerns. Latency is a huge issue, but so is the issue of management at scale, which has already beset a number of HCI users who have attempted to grow their complement of deployed systems.

Interestingly, a debate is resurfacing regarding the compatibility of HCI and web-scale. A debate of sorts on this topic appeared in an online forum at InfoWorld in 2015, but gained little traction at the time. In the discussion, Brandon Salmon made the case that web-scale architecture was not hyperconverged, and hyperconverged architectures are not web-scale. He cited the infrastructures of Google and Facebook for evidence of the incompatibility between the two concepts.

He noted that, traditionally, infrastructure that was designed for large enterprises was a poor fit for SMB environments and viceversa. He noted that hyperconverged systems are "marked by a few main architectural decisions" such as providing "a reliable hardware abstraction (the virtual disk) by replicating between machines." And they "design for an environment that is free of custom hardware." Simply put, these decisions are not the choices of engineers who design for hyperscale infrastructure—like Google, to name one.

Web-scale tends to rely on software abstractions (objects, file systems, etc.) "because hardware abstractions... are extremely difficult to scale, as they impose strong consistency requirements." Examples include the Google File System, Amazon S3, and HDFS: they're all software abstractions that help to create a distributed storage system that provides what workloads operating in a hyperscale environment need—eventual consistency or support for sequential compute frameworks like MapReduce.

The author goes further to explain that web-scale companies sometimes combine compute and storage, but often separate them into different services. This makes sense when you consider that multiple workloads don't all make the same demands of infrastructure. Some require customized combinations of kit compute, network and storage componentry. In the case of Google, they're actually requesting one disk manufacturer to make them a specialized drive with a different platter size to host data from one of their applications!

As we conclude this brief book on CI, it's useful to consider these points. CI sometimes gets negative press or the cold shoulder from storage hipsters (if there are such creatures) because it maintains a disaggregated model for compute, networking and storage. In some ways, it's perpetuating the design of 60 years' worth of IT infrastructure.

However, there are strong reasons for enabling these technology layers to scale independently of each other, especially as the world moves to hyperscale. Infrastructure needs to be simple and automated, but it also needs to be flexible to meet the demands of different types of workloads. At the center of any workable infrastructure option is the requirement for a coherent management scheme that can monitor, provision and optimize both infrastructure resources and services applied in hosting data.

Course Correction

The more automated the delivery of the right storage to the right data, and the delivery of the right protection, preservation and privacy services to the right data, the more appropriate the infrastructure option is. This was supposedly what the storage vendors were working on when the SAN world was disrupted by the virtual server crowd. Converged Infrastructure is intent on returning us to that course.