

NUTANIX PRESENTS

**SILVERBACK
EDITION**

THE GORILLA GUIDE TO...[®]



Hyperconverged Infrastructure Foundations

Scott Lowe, Joep Piscaer, James Green, Max Mortillaro,
Stephen Catanzano, Arjan Timmerman

INSIDE THE GUIDE:

- An Overview of HCI Architecture
- Top Tips for Managing HCI
- HCI Best Practices for Performance and Sizing

**HELPING YOU NAVIGATE
THE TECHNOLOGY JUNGLE!**

In Partnership With

NUTANIX[™]



ActualTech Media
www.actualtechmedia.com

THE GORILLA GUIDE TO...

Hyperconverged Infrastructure Foundations

AUTHORS

Scott Lowe, Joep Piscaer, James Green, Max Mortillaro,
Stephen Catanzano, Arjan Timmerman

EDITOR

Keith Ward, ActualTech Media

LAYOUT AND DESIGN

Olivia Thomson, ActualTech Media

Copyright © 2019 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

Introduction	7
Chapter 1: The Basics of Hyperconvergence	8
Hyperconverged Infrastructure 101	8
Common—and Not-So-Common—HCI Use Cases	13
Securing HCI-Led Enterprises	20
Architecting Security Beyond the Core	25
The Role of HCI and the Cloud	27
Chapter 2: Managing Hyperconverged Infrastructure	30
The 5 Pillars of Infrastructure Management	30
Data Center Infrastructure Layers	31
Moving Toward Invisible Infrastructure	33
Transparent Resource Management	38
Data Protection	45
Managing Workloads	47
Managing Global Infrastructure	49
Machine Learning-Based Capacity and Performance Planning	50
Chapter 3: The Role of the Hypervisor in a Modern Data Center	54
From Product to Feature	56
Legacy Hypervisors Weren't Built for the Modern Data Center	57
Requirements for a Next-Generation Hypervisor	59

Chapter 4: How To Assess Hyperconverged Infrastructure Performance	70
Storage Performance Measurement Has Evolved	71
Different Architectures Require Different Approaches	72
Methods to Assess Hyper-Converged Infrastructure Performance	74
Testing Prerequisites	78
Building the Test Environment	79
Identifying the Proper Data	79
Carefully Weigh the Pros and Cons	82
 Chapter 5: Migrating Virtual Machines to a Hyperconverged Environment	 85
The Need to Simplify	85
Migrations vs. Replacing or Rebuilding	88
Benefits of Automated Migrations	89
Migration Planning	90
 Chapter 6: Sizing Your Hyperconverged Infrastructure Environment	 96
Getting it Right(sized)	96
Why Rightsizing Your HCI Environment Is Crucial	97
The Importance of Monitoring Your Current Environment	104
Tying It All Together	105
Your HCI Blueprint	106


CALLOUTS USED IN THIS BOOK



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



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



When we have a great thought, we express them through a series of grunts in the Bright Idea section.



Takes you into the deep, dark depths of a particular topic.



Discusses items of strategic interest to business leaders.

ICONS USED IN THIS BOOK



DEFINITION

Defines a word, phrase, or concept.



KNOWLEDGE CHECK

Tests your knowledge of what you've read.



PAY ATTENTION

We want to make sure you see this!



GPS

We'll help you navigate your knowledge to the right place.



WATCH OUT!

Make sure you read this so you don't make a critical error!

INTRODUCTION

Welcome to this Gorilla Guide to hyperconverged infrastructure! This book is your Sherpa; it will guide you unerringly through the foundations of this exploding technology. By the time you're done, you'll have a solid grasp of hyperconverged infrastructure (HCI) fundamentals, including its architecture, management, and underlying technologies.

You'll also understand why HCI has become so popular. Organizations are moving to this new computing paradigm at lightspeed, because the advantages it brings in terms of efficiency, cost savings and ease of management can be significant.

Once you finish this short book, we think you'll agree that HCI represents a revolution in data center infrastructure. Whether you know nothing about HCI, or you know a lot, there will be information presented that will help you on your journey through the world of HCI, including lots of practical advice on how to get started.

So, if you're ready, strap on your virtual pith helmet, and let's head into Chapter 1. We'll begin at the beginning: with an overview of what HCI actually is.

CHAPTER 1

The Basics of Hyperconvergence

A decade or so ago, hyperconverged infrastructure, or HCI, entered the modern data center lexicon and turned legacy data center architecture on its head. Since the very first computers were invented, technological progression has been the norm, with organizations seeking to harness the raw power of these advancements to fuel their businesses, improve their bottom lines, and race ahead of their competitors. Unfortunately, even as technology became increasingly accessible for end users, it also became more complex, forcing companies to hire legions of expensive personnel to manage it.

Throughout the rest of this article, you will discover the challenges and trends that led to a set of conditions that made HCI the next logical step in the evolution of computing in the modern era. During this journey of discovery, you will also learn about the high-level drivers that led to the development of HCI and use cases for hyperconvergence, as well as the ways in which hyperconvergence intersects with security and the public cloud.

Hyperconverged Infrastructure 101

In its most basic original form, HCI is a conglomeration of servers, storage, and a hypervisor. All of the workloads in an HCI environment run inside the hypervisor, although recent iterations of the technology have added native container support as well.

In the Beginning

At its core, hyperconvergence began life as a solution to storage challenges. More recent iterations of the technology have sought to break out of this mold and embrace networking, security, and cloud services, all of which have become essential elements in IT organizations. The original focus on simplifying storage has not suffered, though, as hyperconverged solutions slurp up more aspects of enterprise IT.



One area that's been particularly interesting in hyperconvergence is the impact that it's had on the hypervisor market. The hypervisor formerly held the lofty distinction as a strategic differentiator in enterprise IT. Today, that formerly strategic piece of software has been commoditized and democratized as a number of HCI players in the market have taken to building their own hypervisors, often starting with the open source KVM.

The result for these vendors is that they are able to more easily build solutions that leverage these purpose-built hypervisors and design new services around them, without needing to worry about the development cycle for a proprietary hypervisor.

Replacing Traditional Infrastructure

The sad truth is that technology has often been the cause of business inaction, with costly upgrades and onerous complexity as the norm. Businesses have hired expensive staff to try to wrangle these complexities into submission, but there have often been underlying factors that prevented success in these endeavors. The technology itself was built in such a way that expecting simplicity was considered a fool's errand.

Perhaps the most notorious data center element in this reality was storage. Plagued by arcane constructs that hampered growth, and anchored down by a lack of advancement in spinning disk performance characteristics, organizations yearned for a day when storage was no longer chaining the company to the past, but was reimagined in a way that would lead to a better future.

This is where HCI has its original roots. Hyperconvergence began life as an answer to the common storage and scaling complexities of the day in the data center, but it didn't come to life by itself. In fact, there were a number of trends merging at the time that HCI hit the scene, including:

- The beginnings of the rise of affordable flash-based storage
- CPUs with cores and cycles to spare. This meant that commodity CPUs could begin to replace what once required dedicated hardware engineering. Between this and the rise of virtualization, the broader software-defined or software-led ecosystem was born
- The rise of the public cloud, which shifted the perception of how IT *should* run, and unveiled to the business untold possibilities when no longer shackled by outdated technologies and practices

HCI Architecture

There are a couple of different methods by which HCI works its magic. Since storage was the original driver behind the development of hyperconvergence, the focus here is on how that storage is managed: either with or without a controller virtual machine (VM). See **Figure 1**.

Both approaches have associated pros and cons. The most visible benefit of using a controller VM is flexibility. The controller VM doesn't have to care all that much about the underlying hypervisor; it independently manages the storage functions on the host, and is eminently portable between hypervisors.




Figure 1: Hyperconverged infrastructure with a controller virtual machine architecture.

Perhaps the biggest downside of these constructs is that they consume resources and add an additional element in the data path that isn't always present in other architectures. But the impact now is largely negligible, particularly as CPUs have gotten beefier and more capable.

It's true that the controller VM approach is RAM-heavy and ties up CPU cores, but typically the benefits outweigh the costs. These VMs are often handling data deduplication tasks: this requires CPU cycles, and RAM is needed to hold a lot of metadata. The positive is the data reduction—if you can substantially reduce the amount of data that has to be stored on disk, using more RAM is worth it.

It's also true that hardware cost associated with these tasks is far less than just a few years ago, so the overhead as a part of the architecture is reduced. Also keep in mind that today's CPUs have dozens of cores, so carving a few off to handle storage functions that replace complex dedicated hardware is more and more a no-brainer.

The second approach to storage management in hyperconvergence, shown in **Figure 2**, is to allow some kernel module or hypervisor kernel module to handle storage functions. While this method provides a small bit of extra performance, it also shackles you to a single hypervisor. You

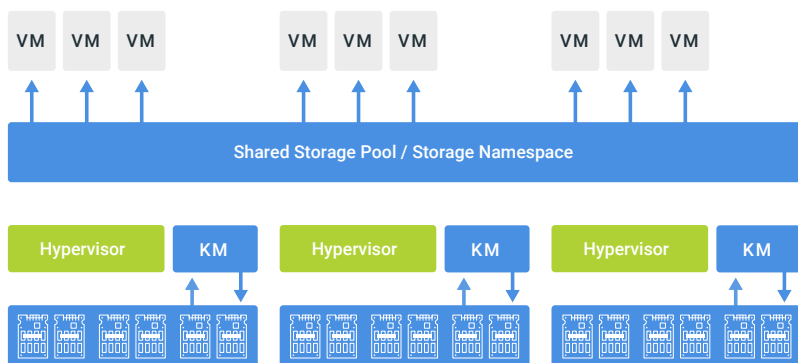


Figure 2: Hyperconverged infrastructure built on a kernel module architecture.

may end up stuck with an expensive piece of software within an ecosystem that you'd prefer to leave, but can't.

Moreover, some of these solutions still carry with them the need for RAM overhead for metadata tables related to data reduction; so, while you may save a few CPU cores, it won't have a major impact on the amount of RAM needed.

Under the hood, both solutions have the same outcome. The storage made available via the HCI cluster can replace what used to be a completely dedicated resource, which often required dedicated personnel and skills to manage. The heavy lifting of features such as encryption and data reduction shifts from what used to be dedicated hardware constructs to software-centric ones that can leverage much less expensive commodity hardware and achieve the same, or, in many cases, even improved outcomes.

You make a decision as to whether you want to buy an appliance with the hyperconvergence software preinstalled, or a supported hardware platform on which you install the software. From there, you create a HCI cluster that provides redundancy and availability in the case of a failure.

When the time comes to grow, you simply add more nodes and join them to the cluster. The management network behind the scenes handles the really hard work of making sure that data is always available, and to the appropriate VMs. Redundancy is provided through a combination of local hardware features as well as software capabilities that make sure data exists on at least two nodes in the cluster.

Common—and Not-So-Common—HCI Use Cases

In the very early days, HCI was pushed pretty hard for “edge” use cases. One such use case was virtual desktop infrastructure (VDI), but it wasn’t uncommon to see hyperconvergence touted for other, smaller uses cases back then.

Over time, the perception of HCI has shifted from edge cases to an architecture that can support even the most demanding mission-critical workloads. There was skepticism at the beginning of the HCI era, but that skepticism has given way to excitement as organizations seek to enjoy the benefits of HCI adoption, including reduced operational overhead and faster time to value for new initiatives, among many others.

General Purpose Workloads

It should come as no surprise that HCI has become a staple for general purpose workloads. These are those functions that every organization has to have; they can include infrastructure servers (DNS, DHCP, Active Directory, print servers, and so on), file servers, application servers, database servers, and anything else that the company needs to operate.

In the context of “general purpose,” there are a lot of ways that workloads can be defined. The easiest way is to say that, in essence, everything can be included. Prior to flash storage, this might have been more difficult to achieve, but with the kinds of performance benefits

flash offers, even hybrid HCI environments can support a wide array of workloads all vying for storage, RAM, and compute resources.

Databases

Databases are the workhorses for most businesses, powering everything from ecommerce sites to point-of-sale systems to customer relationship management tools to enterprise resource planning systems.

Regardless of the actual application, databases all have one key fact in common: they have to perform, since they're usually linked directly to the bottom line. Poor performance can impact revenue in different ways: by increasing expenses due to slow applications holding back employee productivity, or reducing revenue by driving away customers frustrated that the checkout process, for example, is taking too long.

Logging and Analytics

In recent years, logging and analytics have emerged as key workloads requiring significant enterprise IT support. These tools carry with them some relatively unique characteristics. For logging, the underlying

Just Add Nodes

HCI solutions can support even the most intense database applications, thanks to their inclusion of flash storage and the efficiency of the storage stack on each cluster node. Moreover, as databases grow, HCI makes it far easier for companies to expand their storage footprint. The recipe: just add more nodes. That's it! HCI solutions are purpose-built to enable easy scale, which was one of the most significant shortcomings of legacy environments.



infrastructure needs to support a significant level of data velocity: if, for example, the platform has to support fast writes because there's so much data coming in. Moreover, the platform has to support quick and easy capacity expansion, since logging can consume vast swaths of storage capacity as it's being ingested into the system.

Analytics can have similar characteristics as logging; the type of workload generated depends on your usage. Are you gathering data to analyze? That's a lot like logging. Are you mining data for insight? If that's your goal, the environment needs to support fast reads since the analytics platform will need to consume all the underlying data to yield results.

Regardless, HCI can provide support for both write- and read-intensive applications. Of course, there still needs to be some thought given to how to architect the hyperconverged deployment. You can't just deploy a bunch of spinning disk nodes and expect high levels of performance. You'll need hybrid or all-flash nodes to achieve appropriate levels of performance.

Data Protection

Data protection means different things to different people. It may mean just maintaining a high level of availability, which HCI solutions typically do by default; it might mean providing comprehensive disaster recovery services, or enabling a strong data protection and disaster recovery partner ecosystem.

Data protection is a strength of HCI environments. With easily scalable storage capabilities, it's not hard to make sure that there is sufficient capacity to enable data protection services.



In recent years, the term “secondary storage” has come into use as a way to describe non-primary storage needs; it includes backup and recovery as an included use case. HCI is a key enabler of secondary storage offerings, due not just to the easily-harnessed and expanded storage footprint, but also because it’s so central to all workloads operating in the environment.

File Storage

File servers are often the “dumping ground” for everything that doesn’t fit somewhere else; but they’re also chock full of corporate data and secrets, and IT needs to ensure that these resources are well supported and protected. Unlike other workloads, file servers don’t generally demand high levels of performance.

Capacity, on the other hand, is a different story. File servers can demand a lot of capacity, since they might store everything from small text files to corporate board reports to entire libraries of video content from the marketing department.

The sheer scalability of HCI makes it a great match for file services. If you need more capacity, just add nodes. Even better, companies such as Nutanix have taken steps to add native file services¹ to their hyper-converged platforms. Native file services allow customers to deploy powerful and highly-scalable file storage structures in their HCI environments, without having to build out separate Windows File Servers. This simplifies the overall architecture. Of course, these services still integrate with Active Directory to enable secure authentication to what are often sensitive company resources.

¹ <https://www.nutanix.com/products/files>

Edge Computing

Edge computing describes computing activities that take place outside an organization's data centers and cloud environments. Edge computing locations can include remote office and branch office locations, but can also include other locales, including inside self-driving vehicles, which require tremendous computing power that's also immediately accessible. See **Figure 3**.

In the traditional remote/branch office (ROBO) sense, HCI is a perfect fit, since hyperconverged clusters can often start very small. Even the biggest enterprises have very small needs at the edge. The ability for an HCI deployment to scale down to support these environments is critical.

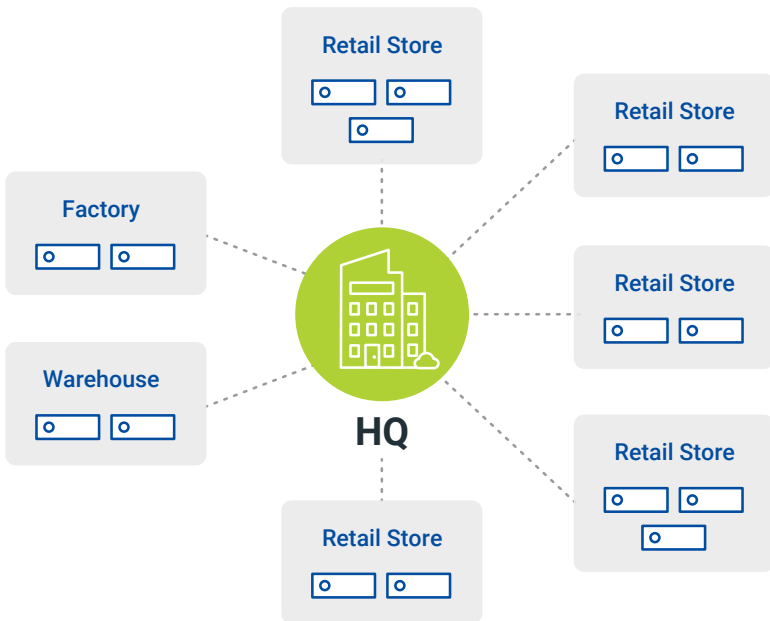


Figure 3: Edge computing pushes processing out to the front lines of the business where IT staff may not even exist.

Moreover, this is one use case in which ease of use and simple scalability are truly key. Many edge environments don't have dedicated IT staff, so the infrastructure deployed into those locations needs to be rock-solid and easily administered. It also has to be scalable. If a store grows and needs more workload capacity, it should be easy to expand.

The needs around the edge are simplicity, scalability, and cost-effectiveness. HCI makes it possible for organizations to design a standard edge architecture and then deploy it as many times as necessary to support the needs of the business while retaining the ability to scale as needed.

VDI and DaaS

VDI and Desktops-as-a-Service (DaaS) are two methods by which organizations seek to bring order to what can be desktop chaos. Around the time that HCI was originally hitting the market, CIOs and desktop architects were struggling with VDI deployments, and often giving up on the promise of the technology. Often, VDI failure was due to underperforming storage, as well as sheer architectural complexity.

HCI collapsed the hard parts of VDI into a single appliance, often imbued with just enough flash storage to help organizations ride out the boot storms and login storms that plagued previous efforts. This is why VDI was paraded as one of the top HCI use cases at its outset (**Figure 4**).

VDI's cousin, DaaS, is everything you'd expect from a software-as-a-service (SaaS) offering. Initial iterations of DaaS, such as Nutanix's Xi Frame,² are fully managed, allowing customers to simply consume their desktops from the cloud. In the future, expect to see on-premises DaaS offerings—based on HCI—that will provide more flexibility and allow organizations to maintain an on-premises desktop environment while enjoying the consumption-centric benefits DaaS has to offer.

² <https://www.nutanix.com/products/frame>

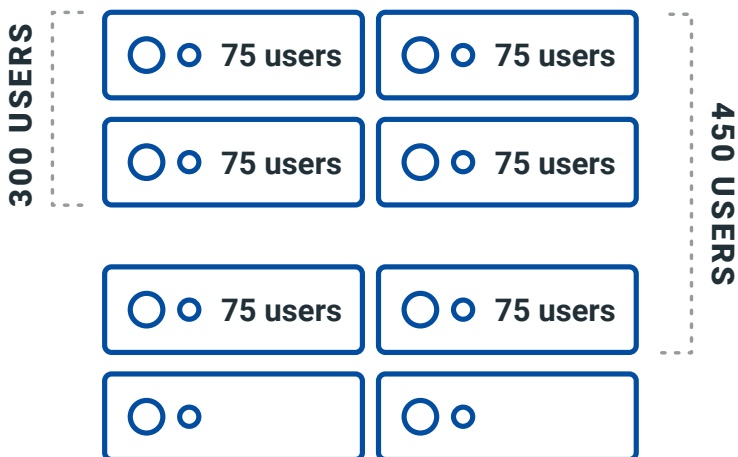


Figure 4: Hyperconverged infrastructure scales easily as your needs grow, turning virtual desktop infrastructure architecture into a simple building-block approach.

Test and Development

Test and development workloads are sometimes relegated to the IT castoff gear that's no longer in use. This is a mistake. In an age in which those with the best software and processes win, having the best gear for developers is paramount. It's critically important for developers to have access to gear that operates similarly to production, but is still cost-effective and can scale as development workloads increase.

HCI provides developers with a programmable infrastructure environment they can include right into their development workflows. They can create and destroy VMs on command, and run it all on infrastructure that performs well, speeding their efforts.

Securing HCI-Led Enterprises

All those use cases have one thing in common: they need to be secure. Without reasonable security precautions being taken, organizations can no longer safely operate on the world stage (or even the local stage). The interconnectivity provided by the internet has made it trivial for bad actors around the world to attempt to infiltrate companies large and small.



Security today has to be baked into every level of the organization, from the infrastructure to the software to the human beings.

For decades, security has been a focus at the architecture level, with patching and updating paramount; this is a basic part of IT security. But there's a whole lot more that needs to be considered as you're kicking the tires of new data center infrastructure, including HCI solutions.

Role-based Access Control

It starts with *who* can do what to *what*. If you're buying infrastructure today and it doesn't provide robust and granular role-based access control (RBAC) to manage who can do what with the hardware, you need to look for a better solution. See **Figure 5**.

RBAC needs to be a consideration in everything you buy. Certain users need broad rights to manage the environment, but others need only enough access to create a VM. This isn't necessarily an issue of whether or not someone is trusted—although it can be—but about what kind of damage can be done by someone with too many rights when their account is compromised, or when there's a falling out between employer and employee.

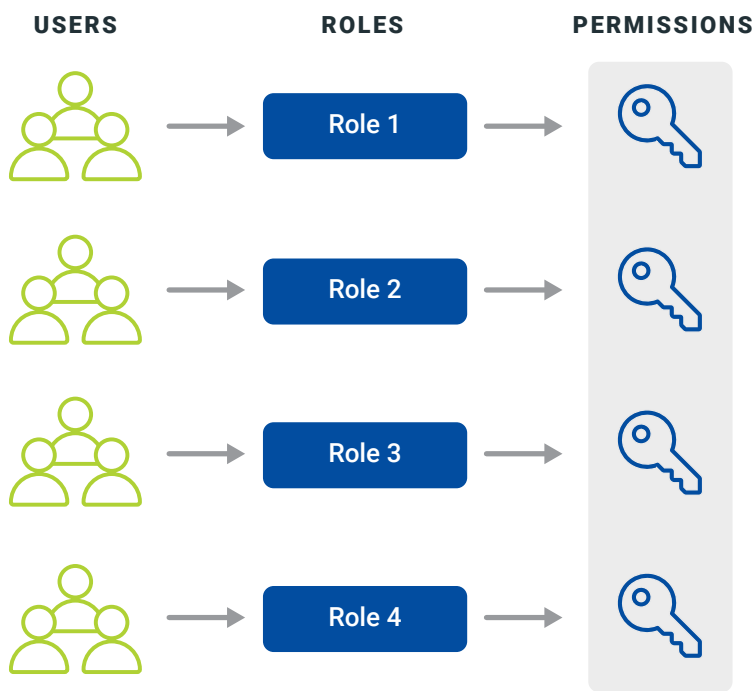


Figure 5: Role-based access control allows security professionals to grant only the access necessary to perform a given job function—and no more.

The software used to manage an HCI environment needs to support this kind of delegation and security. More importantly, the levels of access need to be configurable by the customer. Not everyone needs or wants a bunch of preconfigured roles that might not match local requirements. Very granular custom RBAC permissions enable customers to specify exactly what they need. Tool such as Nutanix’s Prism³ provide the ability for the customer to granularly create a security framework that can be used to design an appropriate access control configuration.

³ <https://www.nutanix.com/products/prism>

Data At Rest Encryption

Physical security is no longer sufficient for organizations that want to maximize their security posture. Every aspect of the environment needs to be secure, whether a particular component will leave the confines of the data center or not.

Let's take storage, for example. It's clear that accessing storage resources from across the world is pretty common, at least by authorized users.

But what about unauthorized ones? What happens if they gain a foothold in your environment and start snooping around? In an ideal world, they still can't see anything, because it's sitting on disks in your data center in an encrypted state.

There was a time when encrypting data at rest might have been optional. Not anymore. Today, your HCI solution must support this ability. Whether the vendor uses some proprietary method or uses disks that support encryption natively is less important than the kinds of security capabilities the vendor provides.

This blog post⁴ goes into greater detail about how Nutanix implements data at rest encryption, and how customers can securely manage the cryptography keys associated with this security feature in a Nutanix environment.

It's important to note that you don't necessarily need self-encrypting drives to enable data at rest encryption. The goal for every environment should be to support highly secure computing practices without regard for what the underlying hardware can do. If the hardware supports data at rest encryption natively, that's great. If not, the software part of the hyperconverged solution should spring into action and provide those services.

⁴ <https://next.nutanix.com/blog-40/security-with-simplicity-encryption-for-your-data-with-1-click-28225>

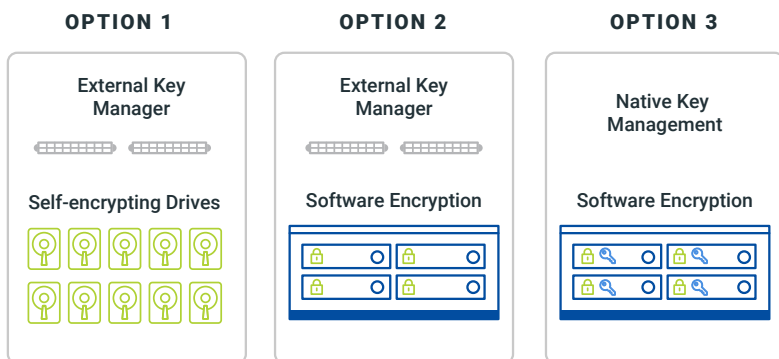


Figure 6: There are a range of options for protecting data at rest: from dedicated, external key management and self-encrypting drives to native key management with software-based encryption.

Single Sign-on

Scattered logins are a severe security threat, in a number of different ways. First, they force users to create different passwords for different resources, which can lead to people keeping written password lists to track everything.

Second, when someone leaves or changes roles, there needs to be an accounting done to determine which systems that user had credentials for; those credentials need to be shut down or changed. It can get messy, particularly if a key system is missed and that defunct user's account lingers on for months or years, just waiting for someone to exploit it.

Single sign-on (SSO) services were born to address the need for centralized authentication mechanisms. These services provide critical authentication capabilities in a centralized way, with the SSO service having hooks into all an organization's systems. SSO communicates securely with these other systems and eliminates the need for separate credentials for every system.

When a new user's provisioned via SSO, they log into an SSO portal; then they're able to immediately access all allowed resources for which their role is configured. They don't need to memorize 57 different passwords for different services, or manage different logins and a maze of password complexity requirements. See **Figure 7**.

HCI components should support SSO for both administrators and end users. Administrators need to access centralized management portals, and users need to access certain services that may be provided directly by the HCI environment. Moreover, any ancillary services offered by the solution need to support SSO. Fortunately, most enterprise-grade hyperconverged solutions provide this support.

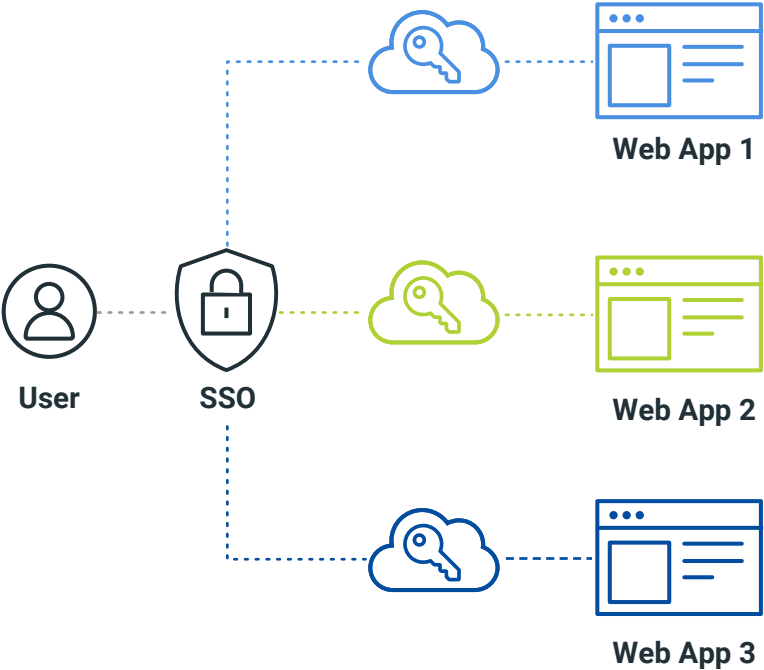


Figure 7: Single sign-on configurations ease the management burden on both administrators and users by allowing access to myriad resources with a single login.

Architecting Security Beyond the Core

Security inside the walls of the hyperconverged solution is satisfied by most of the major players in the market. A strong security solution, however, smashes through those walls and permeates the rest of the enterprise computing environment.

Let's look at how HCI can form the basis for an enterprise platform that enables distributed security, both on-premises and in the multi-cloud world.

Network Microsegmentation

In a legacy computing environment, once a bad guy gains a foothold behind the corporate firewall, jumping from system to system and application to application is trivially easy. The firewall, while playing an essential role, is a single point of security failure in many organizations. If breached, there are no more backstops to keep hackers from spreading like the plague and infiltrating other systems.

This is where microsegmentation comes in. In the context of HCI, microsegmentation is a security service that operates on each hyperconverged host and provides what is essentially a distributed firewall. Microsegmentation learns about your applications' communications patterns before it begins to enforce deviations from expected behavior. This enables the microsegmentation service to fully understand exactly how your applications interact at the network and port level. See **Figure 8**.

Prior to network segmentation, you had to sift through reams of documentation and attempt to piece together the methods by which various parts of an application communicated with one another. No more: By watching your network, microsegmentation can do most of that hard work for you, providing actionable information you can use to help lock down your network.

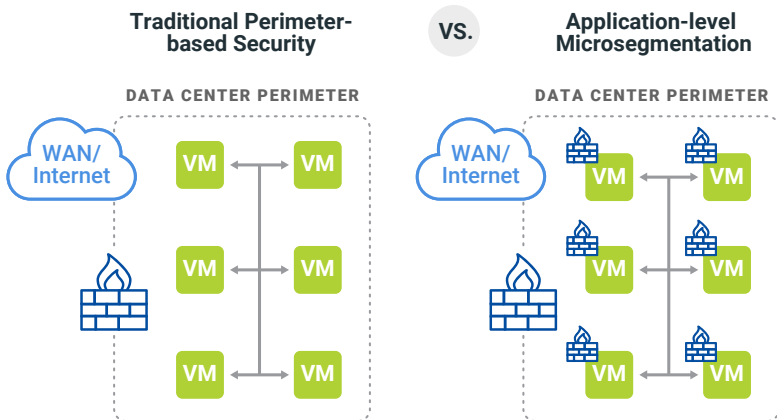


Figure 8: Microsegmentation enables a zero-trust environment, unlike traditional security implementations.

Not all HCI vendors provide this kind of service. Nutanix, with a combination of its Flow network security service⁵ and AHV hypervisor,⁶ is up to the job. Flow is Nutanix’s answer to helping organizations solve critical network visibility challenges that can leave you in a vulnerable state.

Enabling Multi-Cloud Security

HCI is most often associated with the private cloud portion of the hybrid cloud function. By providing local, scalable services, HCI brings to the enterprise some of the benefits of the public cloud—and at the very least, it gets much closer than is possible with legacy architecture.

Some HCI companies take things to a whole new level, though, by embracing the public cloud and helping organizations ensure high levels of security for their new multi-cloud environments. This is an area in which leveraging HCI as the core of the computing environment can pay big security dividends, as that environment morphs into a broader

⁵ <https://www.nutanix.com/products/flow>

⁶ <https://www.nutanix.com/products/acropolis/virtualization>

platform. Nutanix, for example, makes its Beam⁷ service available for both private and public cloud environments.

It's worth noting that the multi-cloud, if it isn't already, soon will be the norm, so onboarding an HCI solution that supports this future is generally going to be your safest option.

The Role of HCI and the Cloud

Of course, the cloud isn't always about just security. There's a lot more reasons why organizations seek to adopt the public cloud or multi-cloud for workloads.

Operationally, the public cloud carries with it a lot of benefits, including on-demand provisioning, a consumption-based expense model, no practical scaling limits, and much more, which enables faster time to value for new initiatives. It's also accessible from pretty much anywhere on the planet.

Simply put, hyperconvergence can be the core of your multi-cloud environment. Your HCI deployment might be tricking you into thinking that you're still living squarely in the confines of the private cloud. Nothing could be further from the truth. More than likely, you're operating workloads in public clouds too. This means that you have, at a minimum, a hybrid cloud architecture; and more than likely, a multi-cloud mashup.

The problem here is that the more discrete services you operate, the more complex things become and the harder it is to secure it all. With the right HCI solution sitting at the center of all this, you might run into a wildly different situation.

That's what on-premises hyperconvergence should look like; it should resemble the core of an expanded platform that provides hooks into the multi-cloud and acts as the conductor in a multi-cloud orchestra.

⁷ <https://www.nutanix.com/products/beam>

Nutanix Beam

Beam is a multi-cloud security compliance and cost optimization tool that keeps a watchful eye across all of your infrastructure silos, from your on-premises ones to the ones stored in various cloud provider environments.



By providing deep visibility into the costs of these services as well as your security posture as aligned to the best practices recommended by each cloud provider, Beam brings your HCI environment out of its silo and makes it a full-fledged player in the race to multi-cloud.

You might wonder what this really has to do with hyperconvergence, though. Beyond Beam's ability to help you keep tabs on your local environment, there's no real direct relationship, but bear in mind this one critical fact: most organizations are adopting cloud services in some way. What if you could adopt cloud services and retain some of the overarching management tools that you've become used to with your local deployments?

For example, your HCI solution might watch for diversion from security best practices across both the on-premises and cloud components of the environment. It might have the ability to keep an eye on web-centric database platforms or provide container services that span the public and private cloud.

It may make it possible to deploy DaaS, both in the cloud and as a part of an on-premises offering. Regardless of how you deploy hyperconvergence, the right solution becomes the very core of your architecture, with support for additional layers that can expand your on-premises capabilities or seamlessly extend into the cloud.

The services provided via this expanded set of capabilities makes it far easier to support the quickly-growing needs of the modern enterprise, and make it far easier to maintain these services, regardless of where they reside.

This brief overview of HCI has introduced you to many concepts, but these are just the basics. As you use and explore more of what it can do, you'll find ways to expand HCI's capabilities into a more comprehensive computing platform that can handle containers, multi-cloud security, file services (block, files, and object), database management, data protection, and a whole lot more.

Now that you have an understanding of HCI basics, we'll move into the practical arena: managing this infrastructure on a day-to-day basis. That's the focus of Chapter 2.

CHAPTER 2

Managing Hyperconverged Infrastructure

In Chapter 1, you learned the basics about HCI, including the ways it's transformed the industry. Now we'll dig into the operational side, and how to manage this new infrastructure. It's easier—much easier—in general to manage HCI; in fact, that's one of its main selling points. But that doesn't mean it's without its challenges. Anyone that has ever managed infrastructure, especially on-premises physical or virtual environments, knows how much work goes into the day-to-day operations, as well as the continuous integration of all the components.

The 5 Pillars of Infrastructure Management

Infrastructure management is the planning, design, delivery and control of the foundational back-end support structure for applications and data. There are five major pillars of management, as shown in **Figure 9**:

- Applications
- Database management
- Servers and operating systems
- Network and storage
- Backup and disaster recovery

Infrastructure management includes its own set of software for simplification and automation of day-to-day tasks and workflows, including monitoring, lifecycle management (software upgrades and the like),




Figure 9: The 5 pillars of infrastructure management.

capacity planning, and configuration management. It focuses on the qualitative aspects of availability, scalability, security, performance, recoverability, and manageability.

Data Center Infrastructure Layers

A modern-day virtualization infrastructure includes a number of components, split up into a number of layers:

The Hardware Layer

The hardware layer consists of compute, storage, and networking. Depending on the type of infrastructure, these can be separate silos of hardware entities or converged into simpler appliances that combine them in different ways. Newer infrastructure architectures push some of the storage and networking logic into decoupled software layers to provide new levels of flexibility and to reduce complexity and cost.

The Infrastructure Software Layer

The infrastructure software layer is the software foundation of the data center. Software components include the hypervisor for running VMs, storage software, and networking software. These three categories consume and abstract the physical resources efficiently to create and run VMs and containers.

Additional infrastructure management software adds capabilities for automation, easy consumption of applications from a catalog, file serving, networking, security features, and hardware and software lifecycle management.



The infrastructure software layer is the software foundation of the data center. Software components include the hypervisor for running VMs, storage software, and networking software.

This infrastructure management is the go-to place to perform operational tasks, like creating a new VM or managing the storage infrastructure, as well as seeing operational telemetry and fixing issues. While many different infrastructure management platforms have similar features, their usability and simplicity differ significantly.

Traditional platforms require you to integrate existing—often open source—products into a coherent management platform, while others eliminate that complexity in favor of a seamless, highly integrated platform. Hyperconverged infrastructure does not only collapse the physical layers of compute, storage and networking, but it also eliminates the need for separate management solutions, each needing separate licensing, deployment and operations.

The Application Layer

Different applications need different kinds of middleware—like databases, web servers and application servers—in order to function. Depending on the application’s architecture, these all run inside a single VM, a group of VMs or in multiple containers.

Older infrastructure solutions do not usually include solutions for easy consumption of operating systems, middleware, and applications. HCI not only saves IT ops teams time by eliminating the daunting task of infrastructure deployment, it also extends the turnkey element to the application layer by providing easy consumption of operating systems, middleware, and applications through app store-like features.

Moving Toward Invisible Infrastructure

The reality of many infrastructures is that most time is spent on the infrastructure itself, not on the applications running on top. Infrastructures are complex beasts, with many sensitive integration points between the different silos; these points often break, especially during upgrades and configuration changes.

Unfortunately, infrastructure on its own does not bring inherent value to an organization. While important, it’s not something of value in and of itself. In reality, it’s more of a necessary evil, a cost center rather than a revenue generator.

This is why simplification of the day-to-day operations and management tasks via automation, smart workflows, and intelligent operations makes sense. If infrastructure holds little intrinsic value, why not make it simple—even invisible?

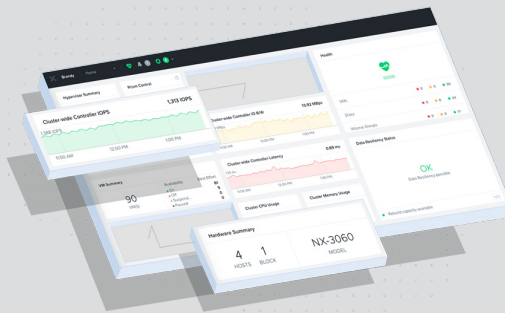
Making infrastructure invisible and operations intelligent is the goal of many public cloud and hyperconverged solutions in the market; it’s what sets the current generation of IT infrastructure apart from the virtualization platforms of yore.

Nutanix Prism: Integrated, Simplified Management

An example of highly integrated, simplified management can be found in Nutanix Prism.

Prism manages all of infrastructure environments across different hypervisors, from storage and compute up to and including VMs and containers on a global, multi-cluster, multi-data center level. Prism simplifies infrastructure management using a one-click approach and streamlines common workflows.

With less time spent on chores, IT admin time can once again be spent on the work that does add value: deploying, managing, securing and scaling applications in use by employees or customers.



Day-to-Day Operations

Infrastructure management in a hyperconverged world isn't fundamentally different than the traditional model—the basic components that make up the infrastructure are the same. Many of the responsibilities and tasks are identical, too.

The big differences lie in the level of integration between the components and the management software that manages the whole. In older siloed approaches, there would be a management tool or suite for each of the pillars: servers, virtualization, storage, networking, back-up, disaster recovery, databases, and operating systems.

These disparate tools lacked integration, with little to no contextual information flowing between the systems; each tool just gave their own limited view of the infrastructure's health, performance, and other issues. This made the life of an infrastructure admin more complicated than necessary.

With the infrastructure and virtualization space maturing, the HCI model provided integration of both the hardware and software components into a single-vendor solution, which in turn opened up the opportunity to tightly integrate all the infrastructure management tools into a single solution.

Many HCI players jumped on this opportunity to create a single, unified view of the infrastructure world. They built management software that can manage the entire infrastructure from a single interface, simplifying workflows and adding richer contextual information from different parts of the stack, allowing for comprehensive monitoring and alerting.

Manage Resources in Bulk, Rather Than as Individual Entities

With highly integrated systems like HCI, cluster management is highly integrated as well. Resources are managed as a whole, instead of per-system or per-silo. This vertical (across silos) and horizontal (across many instances) integration allows for much simpler management. See **Figure 10**.




Figure 10: The management of hyperconverged infrastructure unifies both vertical and horizontal management silos.

With this level of control over hardware, hypervisor, storage, and networking, many of the day-to-day workflows can be automated out of the box, leaving as little manual work as possible for operators.

HCI management is often focused on higher-level workflow management, where system admins follow an automated workflow instead of manually going through various steps. Workflow automation allows admins do more work that adds business value, instead of bogging them down with technical minutiae.

Another HCI advantage is the integration level of operational telemetry, which provides admins global system-at-a-glance information, instead of having to correlate monitoring data from different systems manually. This allows for much richer monitoring and alerting, which in turn means quicker troubleshooting of issues. Tracking all

environmental mutations by encapsulating most of the changes in an audit log is a major part of this.

In addition, by not having to go through multiple (sometimes up to a dozen or so) monitoring consoles to get a glance at the environment's health, it also takes much of the guesswork and manual correlation out of the equation; the system only surfaces up admin alerts that require manual action.

This is especially important for correlating performance metrics throughout the infrastructure. Being able to drive granularity down into the VM, virtual disk or container level is becoming more important, since many of these workloads are becoming more ephemeral in nature. Understanding the resource usage patterns of your applications, and knowing how and what to scale, is a crucial capability of infrastructure management.



There are a number of areas where a holistic view of the hyperconverged environment really stands out. One is data resiliency. Just like any storage environment, data consistency and resiliency are important to monitor. In scale-out architectures like HCI, this is doubly true, as too many node failures lead to data consistency problems.

Monitoring the health states of physical nodes and reporting the number of failures that can occur until the critical point is reached gives admins an immediate overview of the status of their cluster.

Finally, at-a-glance health and capacity planning information is vital for delivering a cloud-like experience to the enterprise. Scaling the infrastructure on a just-in-time basis is an important capability that will make your CFO happy. We'll dive into this later.

Transparent Resource Management

Centralizing physical infrastructure as pooled and shared resources is the core functionality of any HCI environment. Being able to manage all aspects of the pooled resources means managing both the hardware that contribute resources to the pool, and assets like virtual disks, VMs and containers that consume those resources.

Let's look at the core areas of managing a virtual and container infrastructure.

Hardware Is Foundational

The cornerstone of every infrastructure is its hardware. In a hyperconverged world, the hardware layer is simplified into a single appliance. Making sure that this physical layer is available, performant, and scalable is a priority.


Deploying new physical appliance is a breeze due to the simplified hardware model inherent in HCI. It's a matter of racking and stacking the system, then adding it to the appropriate cluster. These workflows are simple and accessible within the management plane of the hyperconverged system, removing the complexity of cluster management. Much of the work to keep the cluster healthy and balanced is done by automated workflows within the system, making the data center self-healing without administrative interaction.

Storage Made Simple

While there are a number of different hyperconverged architectures, general design principles include separate physical servers, each containing storage devices. The storage can range widely, from hybrid flash and magnetic disk to all-NVMe systems.

Infrastructure Lifecycle Management Via HCI

Infrastructure lifecycle management is another advantage of HCI. These operations include keeping software versions up to date for the physical layer (UEFI and devices firmware), hypervisor, storage controller, and others.



HCI's high level of integration makes the management plane a one-stop-shop for all software version upkeep, removing the need for a specialized update tool for each layer and component in the infrastructure.

By pulling all updates into a single system, a high level of workflow automation is achieved, making software upgrades as easy as a single click.¹ This happens without the usual hassle of checking interoperability of new versions or the uncertain impact on the availability and performance of workloads of putting an infrastructure component into maintenance mode.

¹ <https://www.nutanix.com/blog/radically-simple-hypervisor-upgrades>

The physical storage from the appliances is merged into one or more pools or namespaces. From there, storage resources are assigned to VMs, file services or containers for consumption. See **Figure 11**.

The hyperconverged platform abstracts the physical components and performs the storage-related maintenance tasks invisibly. Physical storage across all systems in a cluster are available for consumption as a single resource, often called a *storage pool*.

From these pools, admins can create buckets (for example, containers or volumes) as a consumable resource. Admins have a number of controls to optimize the storage usage for each use case, VM or container. Luckily, most of the so-called “nerd knobs” are simplified

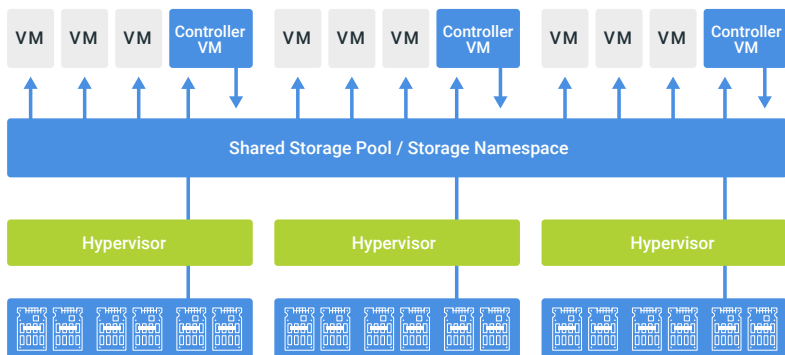


Figure 11: Hyperconverged infrastructure pools silos of storage to make it accessible across nodes and highly available.

into easy-to-use toggles for each storage volume. Each workload has a number of capacity and performance optimization options to configure, such as compression, deduplication and erasure coding.

Storage Capacity and Performance Optimization Explained

Compression is done both inline and post-process. Inline compression has large benefits for large or sequential I/O performance by reducing the amount of data to replicate or read from disk.

Inline compression skips random I/O. Those are picked up later by the post-process compression engine, which has the benefit of making those random I/Os relatively sequential by using spare processing capacity on the cluster.

Even though **deduplication** is a data reduction technique, when applied (inline) to the performance storage tier (read cache), it has a significant positive performance impact by minimizing duplication of data in the finite cache space.



In other words: more unique blocks are held in the cache, increasing the overall cache hit ratio. Post-process (capacity tier) deduplication is usually applied only in specific use cases to reduce on-disk overhead. These use cases include golden or master images for virtual desktop infrastructure, operating system virtual disks, and other datasets that have high duplication across VMs and containers.

Deduplication vs. No Deduplication

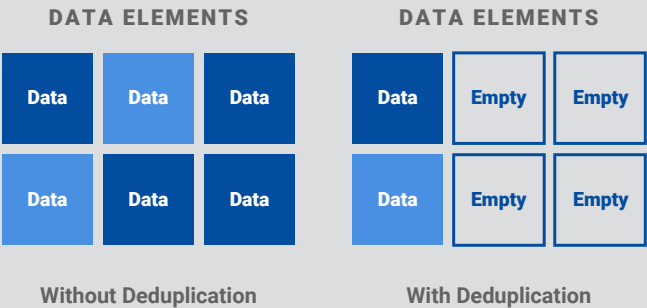


Figure 12: Deduplication can save a great deal of storage space.

Resiliency for the Enterprise Cloud

In HCI, data is protected against any single (or double) node failure by replicating data across multiple systems. The simplest way of accomplishing this redundancy is writing each block of data to two or three systems. This is very similar to how RAID-1 works in traditional systems. As this requires two or three times the capacity, it’s heavy on the storage resources, but it’s the simplest method for achieving data resiliency. See **Figure 13**

Erasure Coding is a way of reducing this overhead. Instead of writing full copies to multiple systems, it provides a balance between resiliency and overhead. Erasure coding is somewhat similar to RAID-5 in that it calculates parity for each block and stripes it across different systems.




Figure 13: Losing a node from a cluster is no reason to lose sleep - data replicated or erasure coded across nodes ensures that your data is safe and available while you bring the failed node back online.

When restoring from a failure, there is some impact on performance, as the parity calculations are performed to restore full resiliency; but it has little impact on normal read or write performance.

When managing a distributed storage system, having a solid overview of the current resiliency status is an important factor in day-to-day operations. It lets the administrator know the current level of protection against physical system failures, and the available wiggle room to, for instance, put nodes in maintenance for software upgrades. Accidentally removing too many systems from the working system can have disastrous effects on data consistency, performance and availability.

Many hyperconverged systems offer an at-a-glance overview dashboard of the cluster's current resiliency status, informing admins about the number of failures the system can absorb. An example of one is shown in **Figure 14**.

In addition to data protection within a running cluster, creating resiliency against cluster and data center failure with backup and disaster recovery functionality is often included in HCI. We'll tackle data protection shortly.

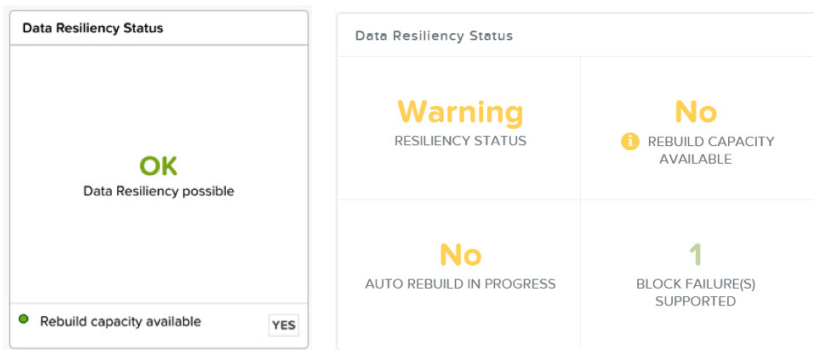


Figure 14: An example of the “Data Resiliency Status” dashboard on a Nutanix cluster.

Security by Integration

HCI brings physical and virtual networking into the realm of the system administrator. There are two distinct pieces of networking to monitor, manage, and operate: the physical *underlay* network and the virtual *overlay* network.

The first is the physical network, which provides connectivity between the hyperconverged appliances. This network, in its most simple form, consists of switches and cabling for interconnectivity between appliances within and across data center racks. Routers exist at the edge of the physical domain to provide connectivity outside the data center.

The physical network transports different traffic types, including the storage traffic across appliances for resiliency and availability, as well as management traffic. Workload traffic, for its part, uses the same physical transport layer, but is often separated into its own virtualized networking space.

In this virtual realm, containers and VMs are segregated into administrator-defined networking segments. Depending on the size and complexity of the infrastructure and workloads running on top, networking is done using traditional VLANs or overlay networking




Figure 15: An overlay network (B) connects endpoints in a virtual network segment. The underlay (A) is the transit network that carries traffic for various overlay networks.

technologies that encapsulate traffic to decouple it from the physical layer. See **Figure 15**.

One of the pertinent use cases for using overlay technologies is *microsegmentation*, or the ability to outfit each VM with its own, centrally-managed firewall. This is a more granular approach to workload security and is fully integrated into the stack, moving functionality from a traditional hardware appliance into the software stack.

Integrating networking into the stack increases the visibility of workload networking, as well. Instead of manually having to correlate technical networking configuration on physical network devices with workloads, application profiles and automation code, HCI stacks have all the relevant contextual information at the ready. This allows for application and workload-centric security, removing much of the complexity and nerd knobs from networking and security.

Having physical and virtual networking side-by-side not only simplifies configuration and automation, it visualizes application communication across workloads. Understanding the flow between different workloads that make up an application makes it easy to create secure, app-centric policies.

Data Protection

Data protection is a key element of any infrastructure. For each component in the layer cake of the infrastructure, a data protection strategy is required. With HCI, these functionalities have been built into the system natively.

Infrastructure Resiliency

Disks are the smallest entity in the layer cake. By using erasure coding, both individual disks within a system, as well as the appliance itself, are protected against failure by storing multiple copies or parity checksums on different disks across different systems. This allows any cluster to sustain one or more appliance failures without data loss.

Backup

Backing up data into a completely separated environment protects against failures in both the data itself (like accidental deletion) as well as in the primary software stack. No matter how battle-tested, there is always a risk of catastrophic bugs in software that makes data unreadable, corrupt or inaccessible. A completely separate stack of software and hardware mitigates that issue.

But the most common use case for backups is the ability to granularly restore items, like files, e-mails or individual virtual disks or VMs. This protects against accidental deletions, human errors and malicious behavior.

Restoring backups requires existing infrastructure to restore the data to. In some failures, the primary infrastructure is not available for restoring data. In those cases, the entire dataset is restored, rather than individual items. For this, we need available infrastructure.

Disaster Recovery

Disaster recovery uses different techniques to capture and transport the data to the secondary infrastructure. Replicating and shipping the primary environment's entire dataset protects against cluster-, rack-, or data center-level failures.



Clusters are usually bound to a physical location such as a single rack or data center. The exception to this rule is a *stretched* cluster, which spans physical locations.

Normal clusters, however, use data replication to another cluster in another physical location to protect against failure within a single location. There are two forms of replication: Synchronous and asynchronous replication.

Each has different characteristics for time-to-recovery (Recovery Time Objectives, or RTO) and amount of lost data (Recovery Point Objectives, or RPO) and has different associated costs and complexity. Stretched clusters are a special form of synchronous replication that allows for more flexibility and tighter integration between the physical locations.

Managing Data Protection

Managing data protection in traditional infrastructure requires a multitude of different tooling and solutions from multiple vendors. As it does for other aspects of management, HCI integrates all these tools into the stack, without compromising the level of protection.

From within a single user interface, admins can protect workloads based on their criticality to the business. Some workloads require just a weekly backup, while other, more business-critical

workloads require real-time protection using synchronous replication or stretched clusters.

The complexity of the underlying infrastructure is hidden by the hyperconverged management stack, freeing admins configure protection levels for business outcomes, instead of having to translate the business value into many deeply technical constructs, configurations, and automation platforms.



The complexity of the underlying infrastructure is hidden by the hyperconverged management stack, freeing admins configure protection levels for business outcomes, instead of having to translate the business value into many deeply technical constructs, configurations, and automation platforms.

Managing Workloads

Remember that there is no inherent value in the infrastructure alone. The infrastructure is important only in that the applications that businesses rely on need the infrastructure to operate. Time spent managing the infrastructure is time away from fine-tuning those business value-generating applications. A hyperconverged infrastructure adds the most value by minimizing the time infrastructure admins spend managing the infrastructure itself.

VMs and Containers


Let's take this idea a step further. VMs and containers by themselves are not valuable. They're a technical requirement for running the applications. Minimizing complexity and the time spent managing these should free up admin time to work on the applications themselves.

HCI hides most of the complexity behind the scenes in managing VMs. Storage, networking, and compute are abstracted away. The technical configuration of VMs and their operating systems are automated into easy-to-use workflows.

This paradigm is massively different than traditional virtualization solutions. Those put the VM at the center of the universe. While a major step up from the previous paradigm of physical servers, it was still not a value-centric approach.

Take the Easy Way Out

Nutanix Calm¹ is a great example of application lifecycle management built into HCI. With its App Store-like approach, Calm simplifies the setup and management of enterprise applications by helping admins incorporate VMs, configurations and binaries into easy-to-deploy blueprints that are consumable by business users in a self-service manner, using a marketplace. Calm provides a growing collection of pre-integrated, validated blueprints from software vendors.



Similar to the application marketplace, other assets such as containers, databases, and files are pulled into this ecosystem. For instance, Nutanix has solutions such as Files,² a fully integrated file server from within the stack, removing the need for separate file storage systems; Era,³ a database management tool; and Karbon,⁴ for managing container deployments at scale.

¹ <https://www.nutanix.com/products/calm>

² <https://www.nutanix.com/products/files>

³ <https://www.nutanix.com/products/era>

⁴ <https://www.nutanix.com/products/karbon>

On-Demand Self-Service for App Owners Speeds Up Software Delivery

Hyperconverged solutions are moving up the stack in recognition of the true value of applications. Not unlike Google's Play Store or Apple's App Store, the simplicity brought by HCI minimizes the effort admins have to put into managing the lower-level constructs in the infrastructure. But that does leave unanswered the question of how to efficiently get applications running and up-to-date in the data center.

Managing Global Infrastructure

Managing a global infrastructure is a daunting task. Managing a single location is difficult enough, but managing multiple, geographically dispersed clusters multiplies an admin's work. Tasks such as upgrading software to a newer version must be repeated for each cluster. Managing applications that circle the globe is nearly impossible without the ability to manage from a single interface.

The generally accepted approach to this is to have a management interface for each cluster, data center, and so forth in smaller environments, but have those rolled up into a global "manager of managers." The global interface monitors all individual clusters for additional insight across clusters, collating and summarizing information into dashboards as well as providing some distinct features for multi-cluster management.

This approach unifies all assets under management into the proverbial "single pane of glass," simplifying the day-to-day operations for multi-cluster management.

Automation

In addition to a graphical interface, having command-line interfaces and programmatically accessible interfaces is crucial in further simplifying management and automation of the hyperconverged infrastructure. These are especially important in the infrastructure-as-code paradigm, where the state of the infrastructure is stored as configuration code on a version control repository.

Machine Learning-Based Capacity and Performance Planning

A highly-integrated stack is great, but it's not a major leap forward unless it can drive itself. Everything discussed in this chapter is the foundation required to create the self-healing, self-driving data center that operates itself, detects anomalies, and creates actionable insights into capacity and performance planning.

Capacity Planning

Capacity planning in traditional virtualization is a nightmare. For each resource type, admins have to dive deeply into each silo to gather capacity telemetry. Trend reports then have to be created manually. Finally, the admin has to manually right-size everything. This is, to understate the case, inefficient. Each task is time-consuming, and overlapping and/or missing data can lead to errors in analysis. Admins need specialized knowledge and experience with multiple storage and virtualization vendors.

Automatically surfacing this information in a visualized “runway” allows plenty of time for the admin to add capacity to the cluster before time runs out. As mentioned before, this adds to your bottom line (and makes your CFO happy) because there's no need to buy capacity that sits idle. You'll be able to bring the cloud's pay-as-you-grow capacity strategy to your data center.

To further streamline resource utilization, right-sizing tools can reduce unneeded resource claims by workloads to free up resources that can then be used in new or constrained workloads. Right-sizing is a key element of optimizing resource consumption.

Automatic capacity forecasting, planning, and optimization makes HCI a true self-driving, self-healing data center that optimizes cost, utilization, and infrastructure health.



HCI simplifies capacity planning. All the telemetry the software needs is in a single, highly-correlated collection. Add in machine learning capabilities and the system will tell you, without any manual effort, when the currently available resources run out.

Anomaly Detection

Traditionally, monitoring systems sent out alerts based on fixed triggers or thresholds. This creates more noise than value. Critical alerts can be missed or even discarded when too many basic alerts must be weeded through.

Machine learning-based algorithms, however, can sift through the noise to find signals that indicate a real issue. Because the system knows expected behavior, it can detect anomalies and serve those up to the admin to address.

Operational Insights


Relying on behavioral analysis and predictive monitoring to detect anomalies removes the manual interpretation from troubleshooting and enables finding and remediating bottlenecks automatically.

Every organization is different. Dashboards should be created based on what’s important to your organization. After all, monitoring data just because someone else does isn’t the point. Monitor data you care about with custom dashboards that provide quick, at-a-glance summaries of application and infrastructure status.

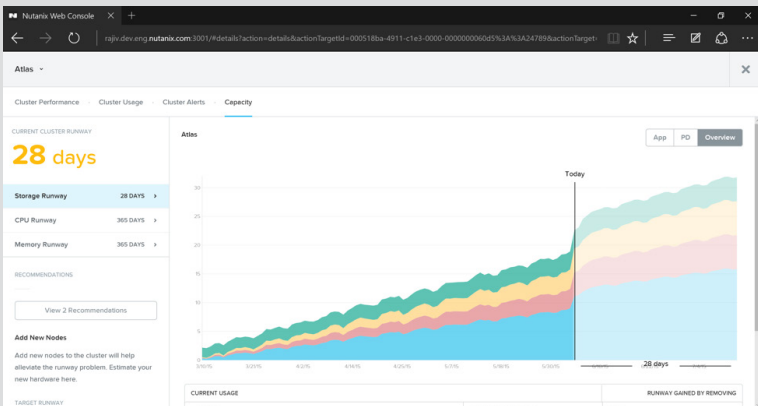
The days of silos and static infrastructure that doesn’t leverage automation are gone—or at least, they should be. Doing things the way they’ve always been done because you know how to do it is over. It’s

See Capacity Problems Before They Arise

A great example of using machine learning to forecast the runway until resources in the cluster run out is found in Nutanix Prism Pro.



Prism Pro identifies the most constrained resource and indicates what the top consumers of those resources are. For storage, usually the biggest constraint, Prism Pro identifies potential options to clean up capacity and suggests ideal node types for future cluster expansion.



impossible to keep up anymore without streamlining operations and automating as much repetitive work as you can. If that sounds like the promise of HCI, it shows you've been paying attention!

Managing HCI properly is, in one sense, about managing the hypervisor, which is at the core of HCI—one might call it the most basic building block of the infrastructure. But not all hypervisors are created equal; is your hypervisor up to the task of HCI? That's what we'll explore in Chapter 3.

Add More Bottom-Line Value

As you've seen, infrastructure management is no picnic, no matter what strategy is used. But applying the advantages of HCI can make it more efficient, saving your admins tons of time they can use to work on projects that deliver more bottom-line value to the business. The old ways are hopelessly out-of-date now, especially in the era of public cloud, the Internet of Things, edge computing, artificial intelligence and machine learning, and so on.



CHAPTER 3

The Role of the Hypervisor in a Modern Data Center

The first two chapters of this book dealt with foundational principles of HCI. Now we'll move onto the most foundational piece of all: the hypervisor, which is the straw that stirs the HCI drink. There is no HCI without the hypervisor; in fact, there's no virtualization at all. That's why the hypervisor deserves its own chapter, starting with a brief bit of history.

The server consolidation trend fueled by the 2001 release of VMware GSX Server fundamentally changed the way enterprise computing would take place over the next decade and beyond. Rather than a simple one-to-one mapping of physical servers to operating system environments (OSEs), data center architects leveraging x86 virtual machine (VM) technology gained the ability to both maximize efficiency and




Figure 16: Consolidation of a large number of servers onto a much smaller number of virtualization hosts changed the game.

decrease costs by consolidating multiple OSEs in isolated virtual servers onto a single physical server. See **Figure 16**.

In the ensuing re-platforming of the data center, the likes of Microsoft, Citrix, and Oracle showed up to claim their piece of the pie, and the race was on to build the hypervisor that would be the foundation of enterprise data centers the world over.

Early hypervisor platforms were fairly rudimentary, but they got the job done. The real catalyst to widespread adoption, however, was the additional features that came with future iterations of the commercial hypervisors. Capabilities like the online migration of a VM from one physical host to another without downtime changed the game for IT organizations.

The increased availability that was realized when a virtual server was no longer bound to a discrete physical server was, in some cases, even more valuable to organizations than the increased per-server utilization and reduced costs for power, cooling, certain software licensing, and more.

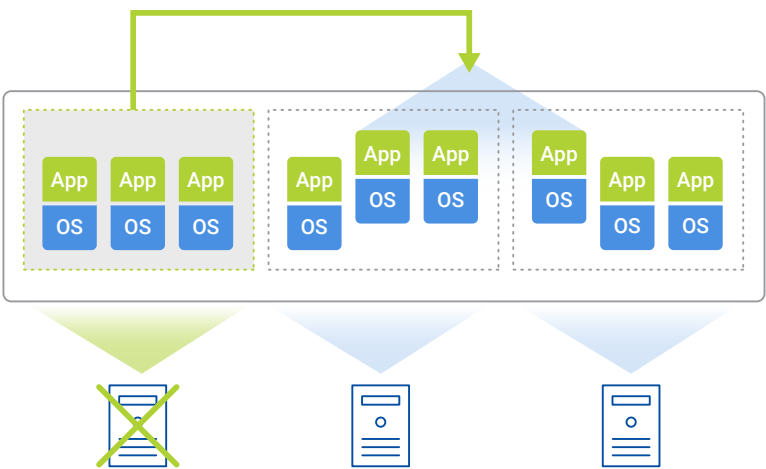


Figure 17: Virtual machine high availability features restart virtual machines on a surviving host in the event of a node failure.

Moreover, VM high availability (HA) features gave enterprises the opportunity to increase resilience for workloads that didn't have native, application-level HA. Across the board, server virtualization in the enterprise was a hit.

From Product to Feature

Nearly a decade down the road from when server virtualization technology became mainstream, a lot has changed in the data center. One of the most critical changes is that the capabilities that come with an enterprise hypervisor—like ESXi or Hyper-V—cemented their place in the contemporary data center are all but expected. Things like HA and dynamic workload placement are now the norm. The new and shiny look of server virtualization has dulled a bit, and sysadmins that were completely floored when they saw their first vMotion aren't terribly enthused when the topic of hypervisor technology comes up today.

Now, as the industry continues to build more and more modern applications that handle things like availability at an application level rather than at an infrastructure level, some of the most critical features of incumbent hypervisors become less important. Cloud-native, microservices-based applications don't need VM HA to be resilient; they have a distributed nature that allows individual, ephemeral nodes to be terminated and spawned without any disruption to the application.

The fact is that when it comes to hypervisor technology, we got what we were looking for. Hypervisor technology has matured to the point where we can easily virtualize any mainstream operating system and business-critical application, run it at the highest levels of performance, and keep it available 24/7. The major competitive advantages of one commercial hypervisor over another have mostly evened out.

The hypervisor used to be a product, exciting and revolutionary in and of itself. The value of a hypervisor in the coming decade will be less determined by the capabilities of the hypervisor itself and more by the

level of integration that hypervisor software has with the rest of the infrastructure stack.

Because the capabilities of an enterprise hypervisor have become the status quo and there isn't much more that IT architects are clamoring for from the hypervisor itself, the focus of IT leaders is shifting away from things like which hypervisor is best and toward higher-level considerations such as how to most effectively deliver IT services to the organization.

As we look at the future, the hypervisor should really just be a feature of a greater platform. When the hypervisor is a product, the focus is on the underlying hardware, the hypervisor layer, and the management of that stack (including network and storage and so on). When the hypervisor becomes invisible, the focus becomes the application layer—which is what we really care about anyway.

Legacy Hypervisors Weren't Built for the Modern Data Center

Statistics show that hyperconvergence is spreading quickly. A key reason is that the infrastructure design paradigm is attractive on all sorts of levels; one of the most prominent is the way HCI deals with storage. Necessarily, hypervisors are in the critical path of storage operations, so how a hypervisor handles storage is incredibly important.

The same is true for networking; how a hypervisor deals with virtual networking is a principal concern as it relates to the performance, ease of use, and overall efficiency of the platform.

When it comes to hypervisors for hyperconverged infrastructure, there are two types. To visualize these two breeds of hypervisor, consider the analogy of a couple of electric cars at the drag strip.

In one lane, you've got a Toyota Camry; a tested and tried vehicle which has been refined over more than three decades of contact with the market. It's a *really* solid vehicle. But in an effort to become more energy

efficient and environmentally responsible racer, you could take the trusty Camry and retrofit an electric motor. Et voilà—an electric car!

In the lane next to the Camry with a bolted-on electric motor sits a sleek Tesla Model 3. This comparable sedan has been engineered from the ground up around the electric motor. For example, the battery is elegantly hidden in the chassis. A far cry from the crude array of batteries and heavy-gauge wires that take up the entire trunk of the retrofitted Camry.

While the *label* of both cars is the same (they’re both “electric sedans”), the experience to the consumer is clearly different. There are a variety of justifications for why you intuitively prefer a Model 3 to a retrofitted Camry if you get to pick which car you hop in and race with; perhaps the most essential reason is that from Day 1, the Tesla was developed with its electric nature in mind.

Now return to the enterprise data center, but keep the electric car drag strip in mind. In one rack, you have hyperconverged infrastructure built on a legacy hypervisor. The entire storage subsystem was developed with a SAN and NAS data center design at the center. All of the storage constructs IT admins manipulate come from the storage array world.



What you’re looking at is the Tesla of the modern data center. It’s a platform (hypervisor, storage, networking, and all) built from the ground up on the principle of hyperconverged infrastructure.

The distributed networking is there too, but it’s clunky and you had to pay extra to get it. The management plane is technically highly available, but only because you deployed two copies of the same monolithic management tool for redundancy. It’s clear that the hyperconvergence experience is bolted on, like the Camry’s electric motor.

But in the rack next to it, you have hyperconverged infrastructure built on a hypervisor that was engineered from the very beginning to support virtualized workloads *on hyperconverged infrastructure*. Nothing is bolted on. Pooled storage maps directly to virtual disks without any sort of odd intermediate constructs. Virtual networking is distributed by default, and management scales out proportionally as the cluster scales out.

What you're looking at is the Tesla of the modern data center. It's a platform (hypervisor, storage, networking, and all) built from the ground up on the principle of hyperconverged infrastructure.

It's pretty clear that a legacy hypervisor is the second choice for building a hyperconverged data center. Let's get practical about what a hypervisor properly equipped for the future looks like.

Requirements for a Next-Generation Hypervisor

It takes a special kind of hypervisor to adequately serve the modern data center, built on the premise of hyperconvergence. There are a few key characteristics you're likely to find in a hypervisor that's up to the task. Consider the following qualities as you evaluate your hypervisor options.

It's a Feature of a Greater Platform

It was once the case that the hypervisor *was* the platform. We've come a long way since then, and the number one feature of a next-generation hypervisor is that while it's rock solid and feature rich, it's merely a part of a larger, deeply integrated platform. That means that the hypervisor is seamlessly integrated alongside the hardware, the storage stack, the networking components, and the management layer.

When a hypervisor is one part of the whole, as opposed to a mix-and-match ingredient in a recipe, the infrastructure becomes easier to manage, is more secure, and has the potential to perform better.

From an operations standpoint, an integrated hypervisor that's a part of an appliance-like stack (vs. a roll-your-own stack) has the distinct advantage of being easy to procure, since it's just included with the platform.



When it comes to security, a hypervisor that's baked into the platform has a smaller attack surface because of a lack of external, third-party dependencies. For example, many hypervisors and their accompanying management tools rely on a database like Microsoft SQL Server or Oracle Database to house their critical data.

Moreover, as ongoing maintenance and upgrades take place, interoperability issues are much less likely to arise since the whole platform is developed together. It's not at all uncommon for a storage array firmware version and a hypervisor code version to be incompatible. While not out of the realm of possibility, this is less likely to be an issue when the hypervisor is developed in tandem with the rest of the stack.

While no one questions the ability of such databases to support business-critical applications, this external dependency can be a weak link in the security of the virtualization infrastructure if not carefully secured. Compromised database credentials and misconfigured roles can lead to an attacker having the ability to tamper with one of your most critical platforms.

A final advantage of deep integration of the hypervisor into the rest of the platform is that if a single vendor develops the hardware compatibility list—the hypervisor, the management layer, the storage stack,

and the networking components—they inherently have the ability to take advantage of new synergies between those different components and develop special features or increase security thanks to the fact that they have control over all components.

While many enterprise infrastructure vendors have strong partnerships that allow for healthy communication across company lines as products develop, each company is still looking out for No. 1, and they're going to serve themselves first. Developing all the components within a

single organization offers the most possibility.

It Provides Reliable High Availability Out of The Box

One of the most exciting capabilities of hypervisor technology in the last decade has been the potential increase in availability created by abstracting physical hardware from the operating system with a virtual hardware layer.



If high availability configurations aren't tuned properly, uptime is at risk. Business leaders *think* they're protected and make decisions based on that belief, but due to the improper configuration, it only takes the right hardware failure or combination of failures for promises made to customers to be broken. Service-Level Agreements (SLAs) could be violated. In some cases, there are significant fees associated with failure to meet SLAs.

Simultaneously, one of the more disappointing problems that's lingered over mainstream virtualization is the fact that high availability (HA) configurations can be tricky to get right. Often, the default

configuration of the HA mechanism is acceptable, but nothing more. Even if you get it tuned properly, it can start off configured correctly but become outdated and invalid over time as the infrastructure grows and changes.

A modern hypervisor should come configured with VM HA out of the box that will reliably protect workloads in the event of a host failure; and it shouldn't take a certified virtualization expert to properly tune the HA settings. Further, it should continue to be tuned properly over time without constant revisiting from an IT admin.

Dynamic Workload Placement Comes Standard

Another exciting feature that the last decade of virtualization gave to IT architects was the ability to balance VM workloads across a cluster for optimal resource utilization and remediation of contention. While this feature is almost always helpful and desirable, it doesn't come standard on the most commonly-deployed hypervisors. It's a feature that must be attained by upgrading to a higher tier license.

Additionally, resource optimization tools on less-modern hypervisors (in conjunction with their hypervisor management software) takes some tuning, just like high availability does.

Since effective resource utilization and elimination of hot spots that could degrade performance is critical in the modern data center, the hypervisor platform of the future should democratize access to this valuable tool and make it available universally. Of course, it should also work optimally the first time you turn it on, and continue to work correctly over time as the infrastructure matures. See **Figure 18**.




Figure 18: Workload optimization within a cluster should be a standard feature and be configured optimally out of the box.

Cluster Upgrades Are as Simple as Upgrading Your iPhone

Upgrades are kind of like taxes. Nobody really wants to do them, but you have to do it. Upgrades come with security improvements, performance enhancements, new features, and new compatibility. All of that is desirable, but it's not always easy to get there. Due to interoperability, upgrades can sometimes be a pain.

Furthermore, it's not uncommon for infrastructure like hypervisors to need a somewhat manual approach to completing the upgrade. Very few (if any) upgrades are performed in a legacy data center without some very attentive babysitting.

It should be expected that a next-generation hypervisor platform makes upgrading this part of your infrastructure simple, straightforward, and painless. Because the hypervisor is integrated with the larger infrastructure platform and has fewer external dependencies, interoperability isn't much of an issue. And due to the unified nature of the platform, the upgrade can babysit itself to some extent while the humans focus on higher-value tasks.

It Ships from The Factory Already Secure


Security is an ever-growing concern in business today, and a hypervisor would be a sweet entry point for an attacker looking to have maximum access to your infrastructure.

Beyond that, a next-generation hypervisor should be just as usable and easy to manage at its most secure. One of the downfalls of some hypervisors is that while you can significantly ramp up security and reduce attack surface, it comes at the expense of usability. Data center operators won't be willing to make that compromise moving forward; they're going to want both a factory-hardened platform *and* optimal usability.

The Delicate Security Balancing Act

Securing infrastructure is tricky because there are conflicting requirements. On one hand, the infrastructure should be usable and operationally efficient; no one wants to manage clunky, cumbersome infrastructure and users don't want to wait for slow, inefficient implementation due to security at the expense of usability.

On the other hand, new vulnerabilities in software are discovered daily and it's important to maintain the smallest possible attack surface and the most secure security posture attainable. To meet both of these requirements, it's undesirable to deploy a hypervisor that's somewhat secure but needs to be modified from the default configuration to be truly hardened.



Management and Analytics Services Are Distributed and Scale-Out in Nature

There are numerous areas in the contemporary data center where the management components of a technology exist independently from the parts actually doing the work. Often, this is a control plane—data plane sort of arrangement.

This model is especially prevalent in networking. Virtualization has been no exception, and the prevailing model for the past decade has been to manage hypervisors and clusters of hypervisors from a control plane that exists as a standalone service. Common examples are VMware ESXi managed by VMware vCenter, and Microsoft Hyper-V managed by Microsoft System Center Virtual Machine Manager.

There are, however, problems with a monolithic, scale-up management model for hyperconvergence. First of all—and most painfully obvious in older versions of vSphere—is that the management server becomes a single point of failure. If (or when) the management service becomes unavailable, whether that's due to a failure or as a planned maintenance outage, the hosts become manageable only on an individual basis. Cluster-wide settings such as high availability and load balancing configurations become unavailable, as do cluster-wide analytics and historical performance data.

The monolithic management model also doesn't scale very nicely. As you increase the number of virtualization hosts and VMs in the environment, resource requirements for the management processes grow. In the case of a fixed and independent management server, it's easy to end up in a scenario where the management server used to be sized appropriately, but as the infrastructure has grown, the originally allocated resources are insufficient. The result is poor performance for administrators and instability across management functions.

A better model for the hypervisor of the future would be to use a distributed control plane where each virtualization host participating in the

cluster or group of clusters runs a management service; as a group, they share the responsibility of providing resilience and performance to the management and analytics services.

Adopting this model ensures that there's no single point of failure that can knock the management capabilities offline, and that as new nodes enter the management domain, the resources scale up proportionately without oversight and input from humans.

It Provides a First-Class Experience to Every User

Obviously, licensing software is a bit different than assigning prices to a tangible good like, say, a cheeseburger. Software costs can be more arbitrary, and tend to be influenced as much by the value they provide as by the Cost of Goods Sold.

That said, the tolerance end users will have for wringing more money out of their budgets to access features now seen as table stakes is likely to decline. Certain features that used to be worth paying a premium for are starting to be seen as necessary, and users may expect them to be a part of the base package in the future.



Some of the hypervisor technologies that have previously only been available with expensive upgraded licenses that are becoming commodities include:

- Automated resource scheduling and work-load balancing
- Distributed virtual networking
- Data protection services such as backup and replication

In a next-generation hypervisor, features like these come standard, and users aren't surprised and dismayed to discover that although the software can technically do these amazing things, they're artificially restricted from those capabilities due to their licensing budget.

Consumption Is Application-Oriented

There comes a point at which porting forward constructs from decades past stops making sense. In many ways, virtualization has reached that point with regard to how physical infrastructure resources are consumed by virtual ones. For example, the notion of LUNs and volumes may still be a necessary construct from a storage protocol standpoint, but administrators have grown weary of dealing with it.

A next-generation hypervisor that's a part of a comprehensive platform should enable administrators to provision resources in a logical way that isn't limited by legacy constructs. In the case of storage, administrators should be able to provision VM storage from a global pool of physical storage straight to the VM without having to worry about arbitrary boundaries like NFS exports, for example. Then, at the virtual disk level, the system will manage things like performance guarantees and capacity optimization techniques.

It's Free From Artificial Limits

It's standard practice in the software world to have at least three numbers in mind:

1. How high do we *think* this will scale in theory?
2. How high have we actually *tested* in the lab?
3. How high are we willing to support?

It's only the answer to the last question that becomes relevant to customers. Cluster size is a perfect example in virtualization. Mainstay

hypervisors have always had a maximum number of nodes per cluster (8, 16, sometimes up to 64).

A maximum cluster size was necessary to prevent features you expect to work in a timely manner from slowing down due to the number for participants in a cluster, and due to the increased resource need to support intra-cluster communication as the number of nodes grows. The problem is that this limit creates inefficiencies in operations and in utilization. Deploying two 8-node clusters with high availability features enabled will reserve (and essentially waste) more resources for failover than a single 16-node cluster will, so avoiding such artificial limits is especially desirable as the number of total nodes in play grows.

A next-generation hypervisor, however, will be free from such limits. Due to the scale-out nature of the platform, the scale is effectively limitless, and the cluster is just as efficient and in control at 256 nodes as it is at 4.

Why Nutanix AHV Is a Next-Generation Hypervisor

The incumbent hypervisor platforms in the enterprise data center are still improving, every day. They're tried and true, and enterprises trust them. But in an era where the most prized virtualization features are also commoditized, many organizations are starting to feel constrained by their hypervisor, and ready to try something different.

They want to be free from the license fees they pay for the privilege of using the established hypervisor; they're ready to break away from platforms with excessive external dependencies and single points of failure. They want something that looks more like Nutanix AHV. Consider the following characteristics in light of everything you've just read.



- Nutanix users pay no additional license fee for AHV. It's included as a part of their adoption of the greater Acropolis platform.
- It comes standard with enterprise-class virtualization features like high availability, dynamic VM placement and load balancing; these features are enabled by default and properly tuned right out of the box.
- Nutanix Prism, which manages AHV (and other hypervisors), runs on each node and is completely distributed. This means there's no single point of failure for management capabilities, and the resources allocated to management and analytics facilities scale inherently as the cluster grows.
- One-click upgrades from the Prism UI make maintenance on an AHV cluster a snap.
- AHV ships with a security posture that's already prepared to comply with common compliance requirements like PCI, SOX, HIPAA, etc.
- Thanks to the tight integration with the rest of the platform, AHV allows one-click deployment of advanced features such as microsegmentation (via Nutanix Flow) and production-ready Kubernetes (via Karbon).

It's clear that modern infrastructure requires a modern hypervisor, one that isn't restricted to old-school thinking. One that addresses the changing nature of the data center, and how the old boundaries of the DMZ have been obliterated. If you haven't examined your hypervisors lately, it may be time to open that door and let some fresh air in.

As you've seen, the hypervisor has a major impact on HCI performance. But how do you determine that performance? It's not as easy as you may think, so Chapter 4 helps you get a handle on it.

CHAPTER 4

How To Assess Hyperconverged Infrastructure Performance

You've been introduced to the building blocks of HCI in the previous chapters, including its major use cases, management issues, and the hypervisor that makes it all happen. It's now time to move up a level and talk about how to squeeze the most performance out of this infrastructure.

Performance assessments via various benchmarking methods have been used by IT professionals to help evaluate and select the right technology to support business goals for decades. Used correctly, benchmarking can help identify the strengths and weaknesses of products and if they're a good fit (or not) for an organization. Every architect, systems engineer, or infrastructure manager has individual preferences, opinions, and assumptions about which vendor, solution, or architecture will best fit their requirements.

These opinions and preferences are often based on years of first-hand experience—complete with IT horror stories—in managing IT environments. While experience provides the needed knowledge to evaluate “what's next” in IT infrastructure, that experience can also lead to biases that can hold back progress.

It's a fact that IT people are opinionated and willing to die on a hill to defend their opinions. What happens when experiences and horror stories lead to opinions that aren't entirely correct?

Benchmarking helps to prioritize facts over opinions. But what kind of benchmarking is best? Let's take a look at various benchmarking methods and formulate a new option: customized benchmarking.

Storage Performance Measurement Has Evolved

The broadest performance differences are seen in storage. It was easier to assess storage performance years ago: Most of the solutions leveraged only HDDs. RAID level was one of the few criteria to determine the underlying system architecture, amounts of reads/writes, and the eventual time needed for rebuilds.

Infrastructure Performance Assessments: A Cure Against Personal Biases

Technologists love objective and quantifiable data. Benchmarking tests are the primary way to gather performance metrics.

Benchmarks provide a standardized, repeatable, and measurable way of assessing the performance characteristics of a given technology solution as part of the evaluation process. The results of these benchmarks will help decision-makers and stakeholders determine whether the selected platform is the right fit for their business and technical requirements.

There's at least one industry benchmark available for every category of IT system. For instance, CPU and GPU performance is standardized: A given Intel x86 processor model will always deliver the same performance as long as external factors (including RAM) aren't considered.



Because the speed and throughput of drives wasn't fundamentally different, it was possible to perform some of the performance evaluation calculations via an elaborate set of formulas on an Excel sheet.

Flash devices have forced organizations into facing complex choices. Now different media types have different physical and performance characteristics. Manual efforts are hopeless; there are too many variables to account for today.

Different Architectures Require Different Approaches

To add to the complexity, monolithic SAN arrays are no longer the only players in the field. New distributed storage architectures have emerged, such as hyperconverged infrastructure (HCI). These architectures have different attributes, and each architecture has its own way of handling how data is written and read.

Most HCI architectures scale linearly, which has long been a sales point in favor of HCI. The addition of a new HCI node or of new capacity should provide a predictable performance increase, because the physical limitations of typical dual-controller storage arrays are eliminated. This means as the HCI system scales, each node adds one controller to the solution, which equates to increased I/O parallelization and therefore increased performance.



This explanation is true in theory; in practice, it will also depend on the load a production environment may be subject to, a common theme throughout this discussion.

Data Locality: A Key Attribute of HCI Systems

Data locality refers to how close the data is to the compute. In the case of an HCI cluster running on top of a distributed file system, the compute portion of a workload (for example a virtual machine (VM)) will be running on one of the cluster nodes (hypervisor hosts).

Because it's distributed, the file system's data is spread across the storage on each of the participating cluster nodes, and mounted as a network share to each of the cluster nodes.

Data locality ensures that the data relevant to a workload should be local to the cluster node where the workload is executing. In other words, if a VM is running on host No. 1 of a given cluster, data locality ensures that the VM's data (such as its virtual disks) is also available on host No. 1.

If a VM has been migrated to another host (host No. 2), either because of a live migration or a high-availability (HA) event, data locality-aware

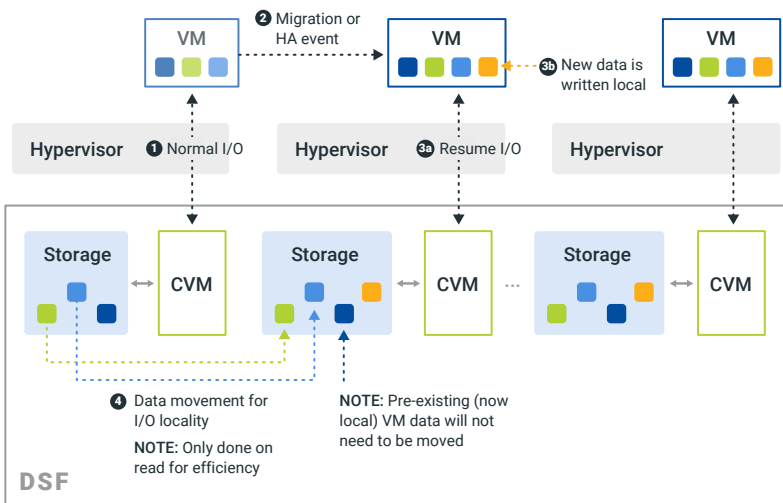


Figure 19: An example of data locality on the Nutanix Distributed Filesystem.
(Source: The Nutanix Bible, www.nutanixbible.com)

HCI systems will begin processing all the new I/O on host No. 2, while also initiating a background data motion to ensure the data is available locally on host No. 2 if it wasn't present there already. See **Figure 19**.

Methods to Assess Hyper-Converged Infrastructure Performance

There are several traditional methods currently available to assess HCI performance. These methods are ordered from the easiest to the hardest to execute. Keep in mind that the more complex the performance assessment, the more accurate the results are supposed to be.

Manual attempts at deriving performance attributes from individual components will not be considered here as a valid method to assess performance. IOPS ratings for individual components (such as SSDs and HDDs) are available. However, they don't take into account the architecture of the system (network protocol and requirements, data locality, caching mechanisms, and so on).

Vendor-Provided Performance Metrics

This is the most basic method to assess infrastructure performance. Vendors go to great lengths to perform performance assessments, whether internally or with external partners, such as industry research organizations.

In most cases, vendors will provide some context about the assessment:

- Which workloads have been evaluated
- Criteria used to evaluate performance (for example, the number of queries per second for a database system)
- Hardware configuration used and how many hardware components were part of the testing cluster

Organizations sometimes decide to rely solely on vendor-provided performance metrics. When they do so, they must scrutinize the testing details. One vendor may end up using fewer nodes but with beefed up configurations (higher CPU class, more RAM, a very fast flash tier, and so on) to beat up the competition. Other vendors may take a different route by proving that they can deliver high IOPS with “average” configurations; in other words, realistic configurations that match what their customers generally order.

Vendor-provided performance metrics are a good starting point. They are a rough idea of the expected performance of a given solution. The downside of vendor-provided performance metrics is that the evaluation team is placing its credibility in the hands of a third party, and some vendors have a tendency to publish “hero numbers,” which are a distortion of the reality.

Synthetic Benchmarks

Synthetic benchmarks will perform one or a sequence of tests, under a specific condition set. These benchmarks, such as IOMeter, help assess the performance of a storage subsystem when using a certain block size and a certain mix of operations.

This requires the evaluation team to understand the characteristics of their workloads, to properly simulate the workload. They need to ask questions like: What block size needs to be used? Is the application profile predominantly write-oriented or read-oriented? Is the application behavior changing over time? For example, an application may perform a lot of writes at a given time, but more reads in another period.

Synthetic benchmarks come with a warning: They’re often used to measure “hero numbers”. Vendors often go with a block size of 4KB, which seldom reflects the reality of block sizes used in real-world use cases.

Hero Numbers: Read the Fine Print!



In the IT industry, “hero numbers” is a marketing term for performance values intended to have a lasting impact on the reader. For the storage industry and the HCI world, these will be values such as millions of IOPS or ultra-low latencies.

While some claims can be legitimate, especially when the testing has been done following an established industry framework, these numbers should always be taken with a grain of salt.

It’s important to read the fine print and understand what configurations have been used to perform such tests, because not all “multi-million IOPS” benchmarks results are equal. One vendor may use just a handful of appliances, while another may need to populate an entire rack to achieve the same result.

Testing methodology can also be used to provide a misleading impression. Often, the test supposes a hypothetical workload that is constant in both duration and throughput, with no external influencing factors. In the real world, workloads behave differently and rarely the same at the same time.

Even with predictable workloads, larger systems are likely to have a certain level of entropy. This means that throughput will not necessarily be constant. Usage will vary in the real world, from periods of peak usage to quieter periods.

Think of hero numbers as extreme performance testing, not an indicator of everyday reality. It takes time and specific conditions to fine tune the fastest car in the world to run at top speed; the same applies for hero numbers. Look for the faithful, reliable and predictable car, not for the mind-blowing 250 mph sports car requiring extensive maintenance after each top-speed run.

Because of this, synthetic benchmarks are useful only as an indication of the potential peak performance of a system under ideal conditions; they shouldn't be trusted to assess real-world performance.

Specialized Benchmarks

Specialized benchmarks attempt to mimic real-world use cases.

One example is the TPC-C framework, which aims to provide a standardized method to assess the performance of OLTP systems (On-Line Transaction Processing). OLTP systems are transactional databases such as Microsoft SQL Server, Oracle DB, and so on.

The advantage of these standardized methods is that since the test characteristics are the same, the outcome of the test can be used to establish a clear comparison between different hardware platforms.

The downside of these benchmarks is that they're often related to a specific workload type. As the complexity of a business system increases, it becomes more difficult to assess performance adequately unless each of the components is tested individually.



A challenge with specialized benchmarks is that most of the mainstream ones are showing their age. The TPC-C website shows that their last release (version 5.11.0 at the time of writing) is from February 2010. The obvious concern is relevance: Is it meaningful to use a nearly 10-year-old benchmark to assess today's workloads? The simulation tools may simply no longer reflect reality, so even if the correct tools are used, applying the results to today's workloads can be a hit or miss exercise.

When components are individually tested, the results of the benchmarks must then be assembled into a coherent, unified performance measurement for the entire system. After all, an OLTP system is usually just one of multiple system components.

Application-Tailored Benchmarking

Application-tailored benchmarking takes a deeper dive into assessing HCI performance. Specific effort is made to reproduce a real-world application.

This benchmarking type provides an ideal insight into what performance can be expected when running in live production conditions. To do it correctly, several steps must be taken.

Testing Prerequisites

First, define the scope of testing. How large is the system to be tested? Can a subset of the system be used? Based on the workload characteristics, a subset of the envisioned infrastructure can be used to perform the test. (Note that this isn't always feasible, as a minimum set of components may be needed, e.g., database servers or application servers).

Second, there should be agreement on the minimum requirements that must be met. It's not uncommon for issues to arise in live production systems despite excellent performance results during the application-tailored benchmarking phase. These can be due to key components having been omitted during the testing phase or an underestimation of the user/data loads.

It often happens that business systems—which were originally designed to support a fraction of an organization's user base—are massively adopted within the entire organization without scaling the application and its underlying infrastructure.

An often-forgotten step here is designing test scenarios. Not all applications have the same behavior at all times. Certain activities take place periodically during the data lifecycle of an application, and those activities may have an impact on the workload characteristics; this can cause a cascading impact on the infrastructure performance requirements.

Building the Test Environment

Once the prerequisites are agreed upon, it's time to prepare the test system. Operating systems and applications still need to be installed and configured on the physical infrastructure.

An on-premises test deployment can take days (if not weeks) to prepare because of dependencies on technical requirements: Racking and cabling, network configuration, availability of key personnel, and so on.

Identifying the Proper Data

The next step is loading the test system with data, which is often trickier than it sounds. This is where the messiness of business intrudes into the benchmarking environment. Identifying the appropriate data to use is critical. Unfortunately, not all organizations have a strict separation of production data vs. test data.

A responsible organization will usually take precautions to avoid using a copy of production data for test purposes. The testing team should at least use anonymized data/non-production data, and always consult with the compliance and data protection teams within the organization prior to any data transmission. (While this step isn't technically related to performance testing, careless use of data can have dire consequences for the organization in case of leak, theft, or data loss, especially for regulated organizations like finance, healthcare, and so on.)

Vendor Certification

Not all workloads are born equal. Certain software vendors (such as SAP, Oracle, Meditech, and others) build specialized software used by organizations to run their critical business or research and development systems.

These business systems are very expensive to purchase, customize, deploy, operate, and maintain. Their vendors usually have very stringent technical requirements that customers must adhere to in order to benefit from technical support. Because of this, software vendors have created certification programs that infrastructure vendors can leverage to validate their solutions.

In their early days, HCI systems were initially not considered adequate for business-critical systems. Experience has proved that HCI systems can perfectly support business-critical workloads.

Vendors invest a lot of resources into software vendor certification programs to support the needs of their most demanding customers. For example, Nutanix AHV (Acropolis Hypervisor) is the first hypervisor to be certified to run production SAP HANA on converged infrastructure. Similarly, Meditech HCIS (Healthcare Information System) is also supported on AHV.

This is important because vendor certification is a guarantee to customers that the solution will not only deliver the necessary performance, but also that both vendors—HCI and software—will support the entire scope of the infrastructure stack.



If the testing environment has been carefully designed, it should be possible to obtain performance metrics that reflect the reality of a production workload.

While application-tailored benchmarking is the most complete way to assess performance, it’s not all sunshine and rainbows. The details that go into this type of benchmarking also make it difficult for businesses to commit to. It’s an intrusive and disruptive way of assessing infrastructure performance, and that costs time and money.

Another concern to take into account is that as technical resources are diverted from supporting the business to enabling application-tailored benchmarking, an organization could miss deadlines and business opportunities. There’s a delicate balance between planning for the future and keeping the organization’s lights on.

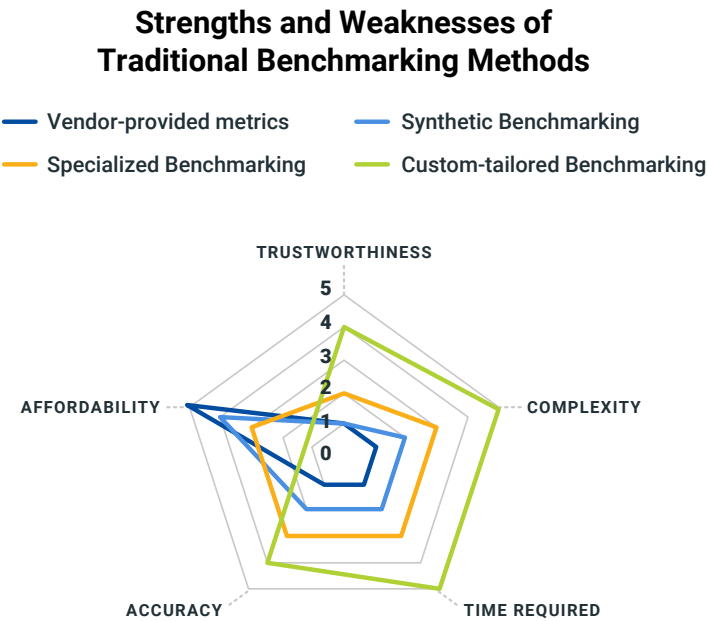


Figure 20: Strengths and weaknesses of traditional performance assessment methods, rated from 1 to 5, where 5 represents the highest value.

Carefully Weigh the Pros and Cons

Traditional testing methods can be useful for organizations to use as they choose a new IT system. However, they all have drawbacks that organizations must be aware of to avoid making a poor decision based on inadequate or misunderstood data. Consider these general strengths and weaknesses of the various approaches and make an informed decision.


To help with that, see **Figure 20**, which ranks the four types of performance assessment methodologies discussed. Some notes related to the strengths and weaknesses of each type of benchmarking:

- Vendor-provided values provide peak performance data based on an ideal environment. But ideal environments don't exist in the real world.
- Synthetic benchmarking just gives an indication of raw I/O performance.
- Specialized benchmarking gives a rough indication of workload performance, but doesn't necessarily simulate a complex enterprise production workload under real usage conditions.
- Custom-tailored benchmarking gives a good indication of performance, but requires extensive planning, timing and availability of resources, and may not provide sufficient data about operating under contention, i.e., working with limited data center resources.

If all this has convinced you that it's time to move to HCI, there's one more important consideration: how to move your existing VMs to the new infrastructure. It can be a tricky, complicated, and manual process. Or, it can be a much more simple, automated, less error-prone process. The choice is yours, and those choices are outlined in Chapter 5.

Customizable Benchmarking: A New Paradigm to Assess HCI Performance

There is one major drawback to benchmarking. Benchmarks are conducted on a brand-new system, as bare as possible, to assess the maximum performance of the selected infrastructure.



The issue, however, is that data center workloads are rarely separated in some pristine environment. While it's useful to know a system's capabilities in an ideal situation, evaluating how a workload will cope under resource pressure is even more useful. Resource pressures such as an environment impacted either by “noisy neighbors” (very large VMs) or by competition for resources must be accounted for.

There is a way, however, to test these specific real-world use cases on HCI-based data center infrastructures. Nutanix has introduced X-Ray, a vendor-neutral, open source testing platform. Organizations can test their HCI infrastructure using reality-based scenarios, not just idealized versions of reality.

X-Ray goes beyond specialized benchmarking tests and takes into account not only contention scenarios, but also the impact of common data center activities such as snapshotting and cloning, as well as rolling upgrades.

Beyond that, X-Ray can also assess application performance during failure scenarios, such as the loss of a node. It comes with a comprehensive reporting module.

Without the need to set up a complex testing environment, X-Ray automates infrastructure testing with meaningful scenarios and delivers detailed reports of the results.

You could think of Nutanix X-Ray as the “Goldilocks Zone” of HCI performance testing: neither too close to the sun, nor too far away, just at the right place, like our planet earth.

That doesn't mean it's always the best tool: for instance, if the testing organization needs to validate extreme workloads or very specific applications, X-Ray might not be sufficient. The most demanding use cases may require application-tailored testing.

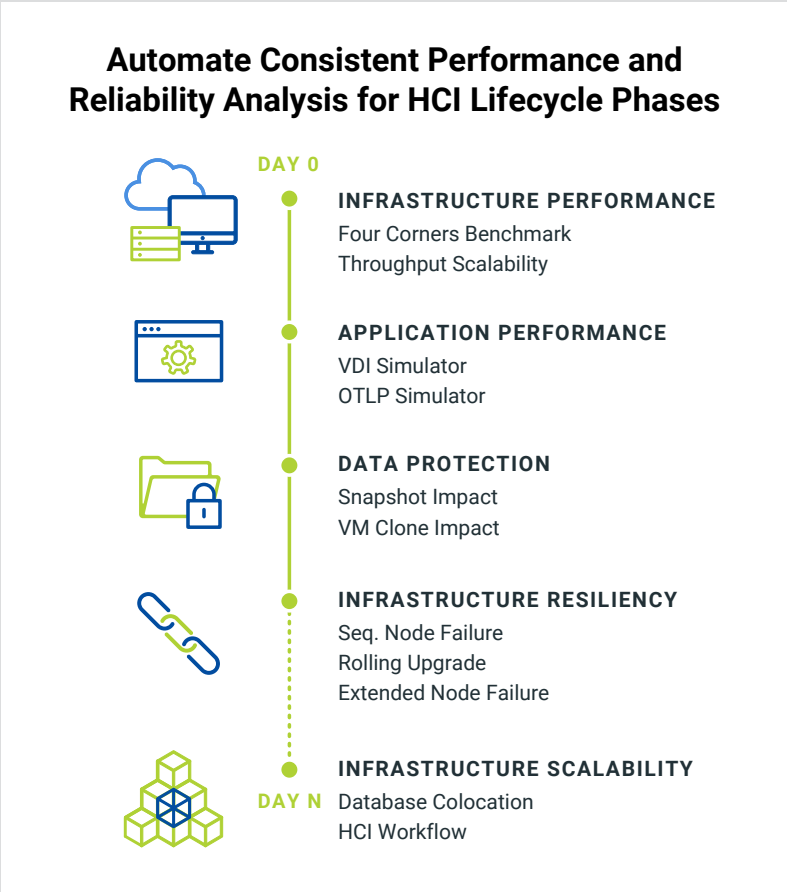


Figure 21: Nutanix X-Ray capabilities

CHAPTER 5

Migrating Virtual Machines to a Hyperconverged Environment

If you're like most companies, your VMs are the heart of your infrastructure. That's why moving them, even to a cutting-edge platform like HCI, can be a scary proposition. Putting your VMs into this environment is the most important part of the move to HCI, so getting this step right is crucial.

The Need to Simplify

To get there, though, let's first set the table by examining the rationale for HCI. As previously discussed, HCI technology is a complete, 100% software-defined stack that integrates compute, virtualization, storage, networking, and security to power any application, at any scale and across multiple cloud environments.

Organizations are rapidly evolving to a hyperconverged environment to reduce costs, achieve greater operational efficiencies, increase business agility, and consolidate resource management to a single strategic platform.

Benefits of Moving VMs to Hyperconverged Infrastructure

Enterprise IT teams today are looking for ways to deliver on-premises IT services with the speed and operational efficiency of public cloud services such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud.

They are increasingly being asked to spend less time on infrastructure and more time and budget on application services that add business value. Despite a continuous stream of IT hardware and software enhancements, the infrastructure challenges faced by IT teams continue to rise.

The IT infrastructure and virtualization software required to meet the needs of a business are complex and expensive, and data center management has become painful. Far too much time and effort are focused on just keeping the lights on.

Taking cues from web giants, HCI combines common data center server hardware using locally attached storage devices (spinning disk or flash) with intelligent software. This eliminates common pain points associated with legacy infrastructure. Migrating to HCI has numerous benefits for IT and the wider organization.

It starts with the bottom line. HCI lowers costs by combining commodity hardware with a simple operational model. Storage networks and storage arrays can be replaced with a single HCI solution to create an agile data center that easily scales.

All elements of conventional, hardware-defined systems are virtualized; these include, at a minimum, virtualized computing, a virtualized SAN, and virtualized networking. These changes represent lower infrastructure costs for an organization by taking advantage of commodity server economics with no dependency on proprietary hardware; that's coupled with the ability to run the latest hardware at higher utilization

rates, while avoiding over-provisioning and ultimately eliminating tech silos to simplify workload deployments.

In addition to hard cost savings, an organization that moves to a hyperconverged model brings many other benefits, including:

- **Operational Efficiency:** HCI consolidates multiple administration functions into a simple, unified management framework, which empowers IT with simplified management and the ability to scale up while avoiding complex maintenance and integration challenges.
- **Business Agility:** HCI enables rapid application deployment and workload-centric, policy-based management. IT can provision full stack infrastructure from one console in a few clicks, and easily tailor infrastructure to suit the needs of the workload.
- **Scalability:** HCI streamlines the deployment, management, and scaling of data center resources by combining server and storage resources with an intelligent, software-defined solution.
- **Strategic Platform:** HCI provides a flexible, strategic platform for next-generation application management.
- **Employee Productivity:** IT administration resources can be reduced and better aligned with the business goals with less focus on server, storage, and network maintenance of legacy systems.
- **Digital Transformation:** Moving virtual machines (VMs) to HCI provides IT organizations a path to digital transformation, delivering consistent infrastructure and consistent operations across data centers and public clouds to accelerate application speed and agility for business innovation and growth.

Migrations vs. Replacing or Rebuilding

After concluding that an organization needs to utilize cloud resources, IT and DevOps need to decide how to handle existing applications and VM workloads. Workloads may be either infrastructure-centric or application-centric; the move can be handled manually by replacing, rebuilding, or re-architecting them.

That's the hard road. It can also be automated using migration and application deployment tools.

Using automation takes advantage of cloud computing to switch from a capital expenditure (CapEx) to operational expenditure (OpEx) model,

Costs Vs. Benefits

When comparing the option to migrate workloads or replace existing infrastructure, an organization needs to evaluate the costs and benefits of each approach. Before HCI existed, the popular strategy was to start by evaluating new systems to replace those currently in use.



The benefits of this approach would typically be cost savings, performance improvements and additional headroom that resulted from taking advantage of new, faster processors and higher-capacity storage arrays. The overall cost of CPU, RAM, and storage declines over time.

The drawback was that the time to rebuild and manually replace existing workloads was often very long, with months of planning, possible down-time, and other risks. In the end, you had shiny new hardware, but hadn't achieved any operational efficiencies or organizational benefits: You were simply moving from one environment to a nearly identical one.

The use of modern, automated migration tools to move workloads to public or private cloud environments, on the other hand, unlocks many benefits. The process dramatically reduces planning and execution time, and provides greater infrastructure cost savings.

which has inherent cash flow advantages for an organization. Migration tools, such as Nutanix Move, allow you to automate the mundane and cumbersome steps required to manually migrate or rebuild VMs in a new environment. With automation tools, VMs can be easily migrated to a cloud and between clouds, providing greater flexibility and management capabilities.

Benefits of Automated Migrations

The process of automating a migration delivers clear technical and business benefits to an organization:

- **Lower Risk:** Business and technical risks are lower when compared to manual migration methods. Automation reduces human error and customer downtime, while creating a repeatable and consistent process for moving many VMs.
- **Time to Complete Migrations:** VM migration projects can be considerably quicker to complete than performing manual rebuilds or sourcing and testing new hardware solutions. The potential onboarding efforts are reduced from months to just a few days, and with minimal disruptions.
- **Licensing Costs:** Automated migrations greatly reduce hypervisor licensing fees, resulting in dramatic cost savings that can be passed on to the organization or customer.
- **Application Impact:** Automation results in virtually no impact on the customer and end-user applications, which remain accessible with minimal planned downtime.
- **Reduced Errors:** Automated migrations are not prone to human error, resulting in more consistent migration results. These migrations require fewer man-hours for migration tasks compared to manual migrations.




Figure 22: Automated Migrations reduce expense and risk exposure across the board

- **Lower Cost:** Both hard and soft costs are reduced through automation. This includes VM licenses, hardware, and software costs if moving from a traditional siloed environment to HCI. It also allows organizations to rapidly move workloads to one or more public clouds to reduce infrastructure costs. This is in addition to soft benefits, including the time to migrate, man-hours, and risk of downtime and disruptions, which can be extremely costly.
- **Employee Productivity:** Automating migrations frees up infrastructure admins' time to carry out other tasks, allowing them to focus their time and resources on projects that have greater business impact.

Automation produces a repeatable, consistent process, with fewer man-hours, less staff, no impact to applications, customers or end-users; and ultimately better overall results for the organization. See **Figure 22**

Migration Planning

When existing workloads require re-hosting or re-platforming, robust migration capabilities are necessary to streamline and automate the process. This can minimize the time needed, the amount of IT effort, and any application downtime.

If organizations plan infrastructure deployments without considering the best way to migrate existing workloads, it can lead to serious underestimation of the effort and time involved in the project. The goal of any IT change is ensuring minimal business disruption, which is

accomplished through good planning and execution. A solid migration plan consists of four stages as shown in **Figure 23**.

Planning

A good plan always starts with documenting and assessing the current legacy environment, then sizing the new environment. Some key planning steps include engaging business stakeholders, gathering all of the business and technical requirements, creating a project design, documenting a test plan, finalizing sizing, and engaging project/change management resources.

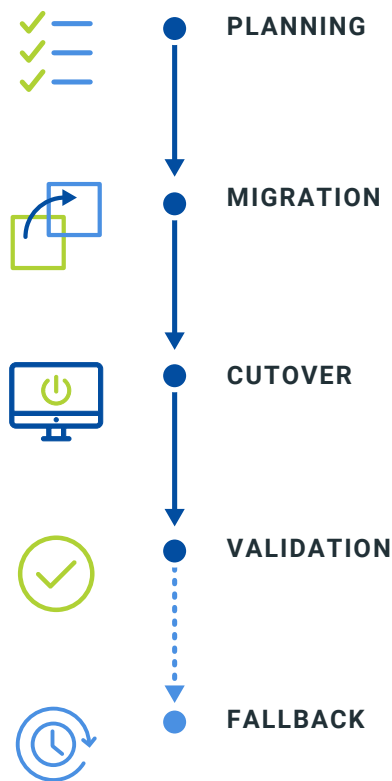


Figure 23: A four stage migration plan (and a fifth step that is ideally avoided!)



A good plan always starts with documenting and assessing the current legacy environment, then sizing the new environment.

Migration

Once you have established the new environment and have it ready, automated migration tools can be implemented to begin your planning and enablement steps. This starts by installing the migration automation software on the target HCI cluster and connecting it to the legacy host. The automation software will then pull all inventory from the hosts and allow you to create a migration plan of the selected VMs. You will need administrator/root credentials, which allow all guest OS preparations to be done automatically. Once completed, network mapping is set up as the final step.

After the migration plan is configured in the automation software, you can start seeding the data of selected VMs without impacting the users. Replication of VM data is at the vdisk level, via VMware VADP (a VMware vStorage API that backs up and restores vSphere VMs for VMware-centric environments), and not through in-guest agents.

This non-disruptive approach is completely agentless and enables VMs to be replicated, whether or not they're powered on or off. VMs are left synchronizing automatically in the background until ready for the cutover. Once the initial sync is done (which could take some time depending on VM size), the ongoing synchronization will happen every 10 minutes.

Cutover

The last step is to make the cutover, which will shut down the original VM, and then do one final offline synchronization before powering on the migrated VM.

Now the migrated VM boots for the first time on the hyperconverged environment. Since it detects new virtual hardware, including the NIC, it will initially start with DHCP. Then the exported IP configuration will automatically be applied so IP configuration will be preserved, including the MAC address if desired.

Validation

Validation is also a key part of the planning process. Once a test or production VM is migrated, a series of tests should be part of the plan. They may include: system integration testing, user acceptance testing, load testing, and workload tuning. All of those tests lead up to a final go/no go decision, with the ability to stop and readjust if needed.

Fallback

Automation also allows you to have a good fallback plan. If something doesn't work on the migrated VM, all you have to do is to enable the virtual NIC and power on the source VM; after that, you're back to the pre-shutdown state.

With the use of automated migration tools, managing the migration of a single VM or large-scale environment can be done with extremely limited risk and consistent results.

Migration Automation

By automating VM migrations to HCI, organizations can reduce costs and complexity, and streamline the challenging migration process.

DevOps and DataOps teams are faced with the increasing challenge of working in a fast-paced, agile environment where IT has become data-driven and more strategic than ever before. Organizations are going through a rapid digital transformation to stay competitive. This

Nutanix Move for Virtual Machines

Adopting new infrastructure platforms often require organizations to migrate existing workloads against tight timelines. One option is Nutanix Move,¹ which simplifies this process with infrastructure-level bulk virtual machine (VM) migrations.



Nutanix Move removes the friction associated with onboarding new infrastructure, and enables businesses to quickly leverage the full potential of Nutanix Enterprise Cloud, with near-zero VM or application service outages during migrations.

Nutanix Move simplifies and reduces the time taken to onboard workloads, providing customers with a short time-to-value. Nutanix Move automates “lift and shift” VM migrations and delivers best-practice configurations needed for optimal VM performance, saving significant time and cost.

Using Nutanix Move involves three simple steps:

- 1. Select the Source:** Specify the source cluster where the VMs need to be moved from. Nutanix Move supports both ESX and Hyper-V for private cloud, and AWS for public cloud.
- 2. Define the Migration Plan:** Once the source VMs have been identified and put into a migration plan, the VMs are replicated to the target AHV cluster through an efficient seeding mechanism.
- 3. Migrate with Minimal Downtime:** Non-disruptive replication of VMs via VMware’s VADP ensures consistent and efficient replication. Whether VMs are powered up or off, they’re kept up to date until ready to be cut over.

¹ <https://www.nutanix.com/products/move#inline5997471702001>

is forcing IT organizations to find ways to reduce costs, increase efficiencies, and move away from legacy data center models focused on infrastructure, to hybrid and cloud computing models.

To achieve this change, IT organizations are looking at their infrastructure challenges and how to seamlessly migrate their VM environments. The options are simple: to remain on a siloed infrastructure model, or shift to HCI with its advantages of consolidated virtual servers, storage, and network, with a single management console.

As you have seen, moving VMs into a new environment can be a laborious manual effort or a simple, streamlined process through the use of automated migration tools. Automation leads to reduce costs and complexity, and is far less error-prone. It also allows for fast access to hybrid and cloud deployments in a single or multi-cloud environment, an increasingly important consideration for businesses today.

When you decide to automate, you also decide to partner up. And choosing the right partner for a migration is a critical step. The right technology partner can make the ability to migrate virtual machines a systematic, error-free process—whether it's dozens, hundreds, or even thousands of VMs.

Before migrating your first VM, however, you need to make sure it's going into a properly-sized HCI environment; in other words, that it has the right amount of resources available to do its work efficiently. That's the subject of the next, and final chapter, of this book.

CHAPTER 6

Sizing Your Hyperconverged Infrastructure Environment

Getting it Right(sized)

We've talked about many aspects of HCI throughout this Gorilla Guide, but one thing we haven't done yet is talk about your current environment from a resource perspective. You'll be moving VMs to your new infrastructure (as detailed in Chapter 5), but before you do that, you'll need to know how much HCI to buy, right? HCI is brilliant, but it's not free; it's important to size such an environment correctly, so you buy enough capacity, but not too much capacity; the great news about HCI is that if you do underbuy, adding more is easy!

The differences between sizing an HCI environment and a traditional environment become clear once you understand the differences in the architectures; but there are also multiple factors that are the same for both.

The importance of rightsizing HCI gets even more important when you consider the multi-cloud, containerized and (micro)segmented world we live in these days, and the fact that more and more businesses are choosing hybrid cloud environments as their future-proof environment. Although sizing an HCI environment might look like challenging job, it's far easier to do than it is with legacy environments.

Why Rightsizing Your HCI Environment Is Crucial

There are a lot demands on the modern datacenter: it must be agile, efficient, effective, interoperable, connected, scalable and fallback capable inside the local data center as well as outside. No big deal, right?

Given these demands, it's crucial that the sizing is done comprehensively. This ensures that an environment will be built that can handle the workloads running in the existing environments, as well as having the capability to handle future workloads. A big advantage of HCI is the ability to scale fast once the capacity thresholds are reached and extra hardware needs to be added.

The Importance of Properly Sizing Your HCI Environment

HCI has gone from the new kid on the block to a solid citizen. The maturity of HCI products has grown impressively in the last couple of years, making it the first option for many companies modernizing their infrastructure. But before making that upgrade, it's important to size your HCI environment correctly. That means making sure you have a proper match of resources to application needs.

Let's take a look at the factors that go into rightsizing your environment for an HCI makeover.

CPU Ready Time

CPU Ready Time is the metric that records the amount of time a virtual machine (VM) is ready to use CPU but isn't able to schedule physical CPU time because all the physical CPU resources on the host are busy.

In virtual environments, creating a VM is easy. This ease of creating VMs leads to issues like VM sprawl and monster VMs that hog resources. In fact, it's not uncommon to see virtual CPUs exceeding the

physical CPUs by ratios of 5-to-1 or higher. The result is that VMs need to compete for CPU resources, often waiting for physical CPU resources to come available. The amount of time the VM needs to wait for enough physical CPU resources to complete the task is CPU Ready Time. The larger this number, the more trouble awaits.

This is why making sure the right number of vCPUs are allocated to a VM is important.



Adding more vCPUs to a virtual machine can actually slow down that workload, particularly in dense environments. More vCPUs means that VM has to await more physical cores becoming available, which may take longer, yielding more latency. Make sure you're sizing VMs for what they need and *only* what they need.

Memory Ballooning

Memory ballooning is a mechanism that provides the hypervisor the ability to claim a VM's memory resources when they're not needed by that VM. When the resources become needed, they're returned to the VM. To make this happen, the VM needs to be equipped with a device driver that's used by the host to (re)claim memory resources within that VM.

This becomes a problem when memory runs short. Allocating more virtual memory than the system has physical memory can cause slow VMs, because memory is paged to disk, which is much slower to respond. This causes the VM to be much more sluggish, leading to poor application performance—and complaints from users that they can't get their work done. Not good.

Noisy Neighbors

A noisy neighbor is a VM that's been given so many resources on a host that it starves other VMs on that host that need resources, too. Those VMs can take a significant performance hit.

To avoid this, it's important to know what resources a VM really needs, and compare that to available resources on an HCI node. You don't want to add a VM with 64 GB of virtual memory to an environment where the host has 64 GB of physical memory—this is bound to create a really noisy neighbor.

Storage Capacity

This is obviously a big one, and one reason HCI and its ability to pool storage has become so popular. Storage capacity constraints can cause VMs to slow down or even shut down once the ~100% threshold is reached.

You need to have enough storage capacity in your HCI environment, so monitoring your VMs for how much storage they're consuming is crucial.

With HCI, increasing capacity is usually as simple as adding more nodes, but it's important to know going in how much you'll need. Failing to have enough could lead to severe—and immediate—problems.

IOPS

Workloads require IOPS, and lots of them. Most applications need to read and write data as quickly as possible, and limited IOPS can become a limiting factor. This is improving, especially with advances in technologies such as faster flash storage and better storage protocols like NVMe. The claim of one million IOPS was rarely made 10 years ago, but a single NVMe drive can provide that today.

IOPS as a sole metric is the subject of some controversy these days, though. Much depends on latency, the size of the I/O in use, and more. But by failing to properly size your HCI environment for IOPS, you could end up with performance issues that will impact the entire environment.

HCI has an answer for these factors, especially in its ability to scale quickly. With HCI, it's easy to start small and grow as your needs expand. This allows you to create the perfect environment for your business needs at every stage: during the initial HCI rollout, and scaling out as necessary by the simple act of racking a new node and plugging it in. Capacity grows as you grow, but simply, predictably and cheaply.



IOPS is an important metric, but there's an overarching metric that is ever more important, as it can reveal whether or not you have storage performance problems in general: latency. If you'd like to learn more about how latency and IOPS are related, peruse this article.¹

¹ <https://www.actualtech.io/latency-the-king-of-storage-performance-metrics/>

Rightsizing Requires Answering the Right Questions

Before you add your first HCI node, you must understand, down to a very granular level, what your environmental needs are. This means getting the right information at the outset. Here are the major factors you need to determine.

How Many VMs Will I Run?

Applications are often operated as a part of a 3-tier architecture: a web tier connects to a middle tier, which handles the business and application logic; this tier connects to the data tier and things like databases and storage. In many cases, applications use at least two VMs; that often depends on whether the database can be hosted on an existing database cluster or has a dedicated VM due to its resource needs.

Finding out how many VMs are needed in an HCI environment means talking to the application developers, database administrators and other front-liners to determine the number of VMs needed and the resources they require. The resources part will be further described below.

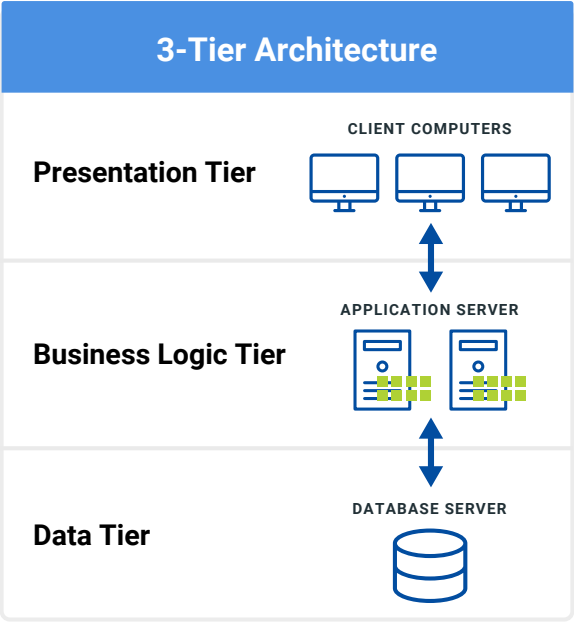


Figure 24: A traditional three-tier setup.

How Much CPU Do I Need?

The number of virtual CPUs needed for a VM depends on many factors. Web servers need a minimal number, while a database server might need the maximum possible. Remember that CPU Ready Times are an important consideration here, making it important to understand the needs of the application, as well as the VMs that host the application.

It's easier and more efficient to add CPU resources to a VM than to investigate and determine noisy neighbors and oversized VMs. Start by making sure the VM has the minimum amount of necessary CPU resources, then monitor if this is sufficient. If needed, add more CPU resources to the VM—but always in minor steps, to avoid wasting CPU.

How Much Memory Will My Workloads Demand?

As with CPU resources, memory resources for the application are scarce and need to be allocated with care. Adding too much memory to a particular VM might cause other VMs to underperform, as discussed with memory ballooning.

Similar to the other factors, it's essential to talk to application stakeholders to determine an application's minimum memory needs. Monitor and investigate the application in the existing environment if you're upgrading to HCI, and allocate the right amount of memory to the application. Apply the same methodology here as discussed with CPU; start small, with the minimum necessary, and add resources as necessary. Remember that doing this for a VM is easy and can be done on the fly.

How Much Storage Capacity (And How Many IOPS) Do I Need?

Getting storage right is central to rightsizing HCI. It's usually not difficult to size basic storage capacity needs per application and can be done easily. But when it comes to IOPS and latency, correct sizing is much harder.

Fortunately, HCI helps here, too. With NVMe drives, solid state storage (SSDs), and having the data close to the CPU (previously discussed, and known as “data locality”), HCI brings the IOPS and latency game to a higher level.

How Much Bandwidth Do I Need on the Network?

Blaming everything on the network is an old and revered game in IT. In traditional environments, each server had its own Network Interface Card (NIC) or maybe even had more than one. With HCI, it's somewhat different; bandwidth can become a high-demand resource.



Blaming everything on the network is an old and revered game in IT.

Stakeholders might be able to provide insight in the bandwidth utilization of their application, but the network (or connectivity) team needs to be involved in this, too. Where 10 GB/s was the norm a couple of years ago, these days 25 GB/s, 40 GB/s and even 100 GB/s is available for HCI.

What's the Expected Application Growth Rate?

Is this an application for a company department with five workers, or something for all 5,000? It's important to know how many will use the application, and the future plans for it.

It might be that an application is implemented for a small subset of users within the company, but it ends up proving so useful that it's rolled out companywide; if that happens, it will impact the application's resource contention, which in turn affects the HCI environment. This is why it's important to know the potential growth rate of the application.

The Importance of Monitoring Your Current Environment

We've already talked about some of this, so we'll boil it down: To right-size your HCI environment, you need as much information as possible. Talking to stakeholders is great, but if you ask them what their application needs in terms of resources, they'll typically say that theirs is the most important and needs the lion's share of CPU, memory, storage, and, well, everything else.

Because of that tendency, monitoring and analyzing the current environment—not just handing a blank check to an app for resources—is what's needed to protect you from overprovisioning VMs.

Don't Forget Physical Data Center Needs

No IT environment is 100% virtual. Unless you outsource every bit of your infrastructure, you have hardware. That stuff is important too, so be aware of these factors.

- **Rack layout.** The rack layout of an HCI environment is very important. Its resistance to failure can be greatly expanded by locating

the HCI nodes across multiple racks or even multiple data centers so that when a rack fails, the environment can continue running. Building in resiliency can keep your business running.

- **Power.** The power consumption of hardware is increasing, even with the implementation of virtualization and cloud computing. And, although the HCI datacenter footprint can be minimized, power is still an important consideration. You'll have less hardware overall, true, but those nodes need plenty of juice.
- **Cooling.** This hasn't changed with HCI, although, again, it may be reduced. The cooling capacity of your data center determines how many nodes can be housed in there. If the cooling capacity is less than the heat generated by the HCI nodes, you've got some expensive choices ahead.

Tying It All Together

Rightsizing your environment *before* you add your first HCI node is crucial; you don't want to just start tossing boxes in the rack, wiring them up, and then hitting the On button. Without carefully determining your application's needs for resources, it would be easy to buy too much capacity that sits there year after year, unused. In fact, that was the original idea behind virtualization: to utilize CPU, storage and other hardware more efficiently, instead of running one email application on an entire server, while most of the hardware just sat there, staring into space.


But that's exactly what you'll be doing if you don't perform due diligence beforehand. One of two things are likely to happen: either your capital expenditure will be wasted on apps that don't need the available resources, or your nodes will be severely overtaxed because you've underestimated those requirements, causing you to spend more money on nodes—money you may not have in your budget.

Both of those outcomes are disastrous, so keep them from happening to you by figuring out what you need, before you need it.

Your HCI Blueprint

In a way, data center infrastructure is similar to building a house: it starts with a strong foundation. For a house, that means a good location, leveling, lots of cement. But first and foremost is the blueprint. The blueprint is your source of truth, and without it, a solid house won't be the result.

It's the same thing with your data center: you must have a blueprint. That's what this Gorilla Guide to Hyperconverged Infrastructure Foundations is—the plan for building your new data center. It's a data center that's ready for today's demands, that responds to the requirements of today's applications. A data center that's efficient and frees up your expensive IT staff to pursue tasks that have much more value to your organization.



HCI does all that, but without a plan in place of how to get there, you could end up with a structure that's rickety, and ready to collapse when the first big wind hits it.

This book is your plan. Keep this Guide handy and refer to it often. These six chapters covering the basics of HCI will help you construct a house on a rock-solid base that will weather the IT storms, now and in the future.

Happy Building!