

**THE
GORILLA
GUIDE TO...**®



Hybrid Cloud With Hitachi Vantara and VMware

Katherine Gorham

INSIDE THE GUIDE:

- Why Organizations Are Moving to the Hybrid Cloud Model
- The Ways Hybrid Clouds Introduce More Complexity
- How Hitachi and VMware Combine to Solve These Challenges

**HELPING YOU NAVIGATE
THE TECHNOLOGY JUNGLE!**



ActualTech Media
www.actualtechmedia.com

In Partnership With

HITACHI
Inspire the Next

vmware®

THE GORILLA GUIDE TO...

Hybrid Cloud With Hitachi Vantara and VMware

AUTHOR

Katherine Gorham

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

- Chapter 1: Forecast: Cloudy.....7**
 - The Current State of Play for Enterprise IT.....9
 - The Case for Private, Hybrid, and Multicloud In a
Supposedly Public Cloud World.....11
 - Emerging Technologies.....13
 - Hitachi + VMware = Superpowers.....16
- Chapter 2: Data-Driven Digital Transformation.....19**
 - The Churn.....21
 - Virtualization.....22
 - Cloudy, Cloud, Cloud, Cloud.....25
 - Containers.....26
 - Application Design.....30
 - Infrastructure Meets Application Design.....32
 - The Future of IT.....34
 - Automation, Anywhere and Everywhere.....36
- Chapter 3: Your Data in a Hybrid Cloud.....39**
 - Data Locality.....41
 - Data Sovereignty.....42
 - Privacy.....44
 - GDPR.....44
 - Muddy Data Regulations.....45
- Chapter 4: Better Together: The Hitachi Vantara and
VMware Advantage.....47**
 - Complexity Is About More Than Just Technology.....49

Multiple Moving Parts From Multiple Vendors on Multiple Infra-structures = Problems!.....	50
The Value of Vendors.....	51
The Value of Ecosystem Continuity.....	53
VMware and Hitachi: Better Together.....	53
Chapter 5: Technology Considerations.....	55
Raw Speed.....	56
VMware Cloud Foundation and vRealize.....	56
Hitachi Enterprise Cloud.....	58
Addressing Real-World Problems.....	59
The Big Takeaway.....	61


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!

CHAPTER 1

Forecast: Cloudy

The IT industry changes at such a breakneck pace that it's hard for anyone to keep up. Cloud computing is the new hotness, but with so many competing technologies falling under that umbrella, how can one transition one's skills—let alone workloads—from what exists today to what will be needed tomorrow? Hitachi Vantara¹ and VMware offer a mix of product and services that can help those who find themselves in any part of their cloud journey.



Extracting signal from noise is the role established enterprise vendors play. Established enterprise vendors have focus, experience, and a good idea about when to pounce and when to leave an idea to gestate within the startup that initiated it.

Life changes at such a breakneck pace that now it's possible to have a successful small business without a business phone number—because Instagram. If you're feeling a bit too old for all these changing paradigms, (as your author surely is) remember that the whippersnappers number in the billions and both social and technological change are intertwined.

¹ HITACHI is a trademark or registered trademark of Hitachi, Ltd. All other trademarks, service marks, and company names are properties of their respective owners.

As frustrating as the pace of change is, it's by many accounts accelerating, not slowing down. In life, in business, and certainly in IT, it's easy to become discouraged. The bleeding edge of shifting paradigms get all the press and attention, and start-ups seem to have limitless energy to chase after every sparkly object.

Some start-ups will be bought up. Others will grow to join the ranks of the world's tech titans, redefining IT for at least an entire generation. It's hard to know which of the myriad start-ups working hard to redefine the future will be successful and, as a result, which bets to make in training your staff or which technologies of tomorrow to pick.

Extracting signal from noise is the role established enterprise vendors play. Established enterprise vendors have focus, experience, and a good idea about when to pounce and when to leave an idea to gestate within the startup that initiated it.

IT may be an industry in which constantly accelerating change is an everyday reality, but it's important to remember that the final product delivered to customers must be engineered to last. Those customers may be internal to your organization, and the timeframe of their service may be only a few years, but there must exist a disconnect between the rapid churn of the IT industry and what is delivered to end users.

End users need IT that actually helps them get things done. IT practitioners need to be able to find a path through the churn that allows them to select relevant tools. No individual—and no organization—can know everything. By working together, however, leading enterprise vendors can combine cutting edge research with generations of expertise in picking winners to provider organizations, their IT practitioners, and their customers with IT solutions that meet today's needs, all while preparing them for tomorrow.

The Current State of Play for Enterprise IT

The re-election campaign of Barack Obama famously used² DevOps, cloud computing, and open source software. One oft-quoted tweet exclaimed that the campaign's IT achieved "4Gb/s, 10k requests per second, 2,000 nodes, 3 data centers, 180TB and 8.5 billion requests. Design, deploy, dismantle in 583 days to elect the President." This was almost a decade ago, yet the very same technologies and approaches used in this campaign are still slowly being integrated into enterprise IT.

For a host of reasons, few established organizations choose to do without on-premises data centers. Personal experience from anyone who has worked in an enterprise will show that today's enterprises use an opportunistic mix of public and private IT, and that in most organizations the overriding philosophy is finding the mix of products and services that works for that organization. Ideological philosophies, such as "all in on cloud," take a back seat to pragmatic reality.




Making things work and keeping them working are what matter. Which cloud is used, and whether that cloud is public or private, is emphatically not the important part.

A quick look at the results from the Rightscale 2019 State of the Cloud survey³ reinforces this: 94% of respondents use some form of cloud; 91% use some level of public cloud services, whereas 72% use some level of private clouds. Public cloud use and adoption, however, are a top priority for only 31% of enterprises. (see **Figure 1-1**).

² <https://arstechnica.com/information-technology/2012/11/built-to-win-deep-inside-obamas-campaign-tech/>

³ <https://www.rightscale.com/blog/cloud-industry-insights/cloud-computing-trends-2019-state-cloud-survey>



Source: RightScale 2019 State of the Cloud Report

Figure 1-1: Companies see the value in cloud computing, as most use some form of it.

Making things work and keeping them working are what matter. Which cloud is used, and whether that cloud is public or private, is emphatically not the important part.

The results of the survey, while informative, should not be shocking to anyone who works in IT. Public cloud services range from the simple—such as Dropbox—to the complex. It makes perfect sense to see nearly universal adoption of at least some public cloud services.

If there is a surprise, it lies in a 72% adoption rate for private clouds. Private clouds have been, until relatively recently, something of a pain to deploy and manage. Such a high adoption rate, despite the historic challenges, speaks highly of the utility of private clouds to the world's enterprises.

What is not revealed by the data in this survey is how extensive the adoption of either public or private clouds is, not which workloads are tending to live on which clouds. General industry wisdom has data warehousing, big data analytics, and similar “need to be able to scale quickly” workloads living in the public cloud. This accounts for much of the consistent public cloud growth as well as projected future growth.

That same general industry wisdom, however, still has many—if not most—of an enterprise's core workloads running on their on-premises

Building an Enterprise Cloud

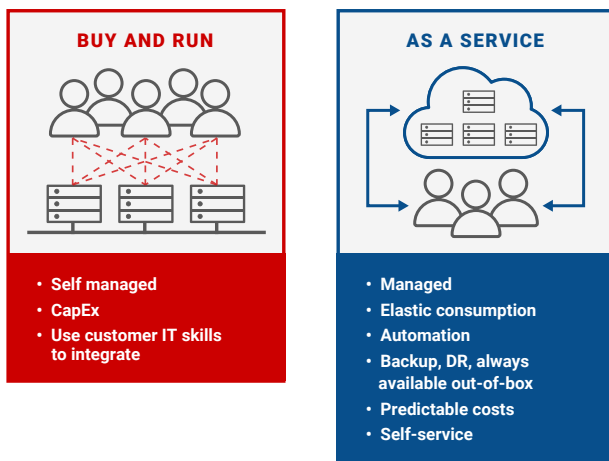


Figure 1-2: Pros and cons of private cloud development, implementation and management.

IT infrastructure. Ensuring that these workloads can meet the needs of all users, while remaining agile enough to respond to changes in demand, takes the right mix of technology and experience. **Figure 1-2** shows some of the considerations faced when deciding on which direction to go when using a private enterprise cloud.

The Case for Private, Hybrid, and Multicloud In a Supposedly Public Cloud World

Making on-premises workloads resilient, agile, and ready to meet the demands of today's increasingly distributed workforce means using a private cloud, even if that private cloud spans several data centers in different locations. In the real world, most enterprises don't own all of the data centers they use. Their private cloud stretches to include service provider managed services as well as nodes located in colocation facilities.

Increasingly, workloads need to be geographically close to workers in locations where it isn't feasible for an enterprise to actually maintain a full data center. In addition, the need to join on-premises workloads with those located in a public cloud is now an everyday requirement of IT.

The line between public clouds, private clouds, hybrid clouds, and multiclouds is blurring into insignificance. These terms are, for the most part, relevant to vendor marketing teams and venture capitalists, not to real-world IT teams.

Interconnecting workloads that live in multiple public clouds with on-premises workloads (whether they live on a private cloud or not) and then making them work with whatever other IT innovation happens to be next on the list is just IT. It's the buzzword-independent reality of keeping the lights on, and it has been since Oct. 29, 1969, when the first computer network was created.

Most enterprises are not, at this point, going all in on the public cloud, and they may never do so. Many IT teams feel they have good reasons not to. The cost would probably be prohibitive, even with volume discounts. And then there's the compliance elephant in the room. For many organizations, regulatory compliance—especially once international regulatory regimes are taken into account—means that keeping data in a private cloud remains the best way to have visibility and control of it.



The carrots of simplicity and ease of use are counterbalanced by the sticks of unpalatable fees as well as lock-in and regulatory compliance. This is the tightrope enterprise IT teams must walk, and somewhere in all of this, there must be the ability to continuously adapt to change.

All of this is changing and will continue to change, though not in the predictable and attractive upward curve of an analysts' graph. Public cloud adoption is more likely to look like the jagged peaks and valleys of an electrocardiogram.

Did one of the major public clouds develop great new features for data analysis that will drive increased adoption of public cloud services? Is there strict new data privacy or locality or sovereignty rules being rolled out as part of a new regulatory regime? This will trigger some backtracking into the private cloud.

The carrots of simplicity and ease of use are counterbalanced by the sticks of unpalatable fees as well as lock-in and regulatory compliance. This is the tightrope enterprise IT teams must walk, and somewhere in all of this, there must be the ability to continuously adapt to change.

Emerging Technologies

It isn't unreasonable to ask if a swift reaction to change is relevant to enterprises. After all, aren't enterprises famous for inertia? Aren't there entire libraries of books about whether or not enterprises can afford the agility and risk taking common in smaller organizations? Conversely, can enterprises afford stagnation or the status quo?

The answers won't be found here, but the fact that these questions continue to be debated, pontificated about, and the topic of umpteen career self-help books every year is itself the point: Opinions vary, and as a result, every organization's IT team needs to be able to do things differently from the next. Despite this, they must also be able to claim to be using industry-standard best practices. These things matter in a courtroom.

All of the above puts the enterprise IT decision makers in a place of unpleasantly conflicting messages. Go to the cloud! Be agile with extra buzzword sauce! Get every possible iota of insight from your data!

But don't get stuck with the wrong vendor when a data breach happens, a new regulatory regime drops, or the pricing model changes into something resembling a torture device. Make the wrong choice and you're the person trotted out for career sacrifice on the altar of public outrage when something goes wrong.

Do it now but do it right. The sharks are waiting for their next meal.

Which technologies will help you most? Are containers the one true future, or do you foresee decades of virtual machines (VMs)? Will serverless render all others irrelevant? What will you do about the need for your data to persist, even if workloads do not? How do you make your multifaceted next-generation architecture resilient, and how do you keep track of everything when workloads that used to be measured in the hundreds start to multiply into the hundreds of thousands?



Certain questions hang over everything modern IT teams do. What data will you keep? How will you secure it? What underlying infrastructure will you use? How are you going to store your lakes, oceans, or watery planets of data?

Containers are an example of a new enough space, and enterprises can definitely be forgiven for waiting to see what challenges the early adopters face before evaluating them as a possible tool. Despite this, they're clearly useful enough that—even if they aren't experiencing primary adoption today—IT teams will want to keep the container option open for the future.

While preparing for that future, it's worth thinking about redesigning workloads for composability now. This raises the question of whether or not your organization has the required skills internally to do that or whether time needs to be taken to build these skills now, ahead of the rush.

Even if you press pause on containerization to give the market a chance to mature, you still have a mess of options for getting into the cloud and for what you do once you're in there. Perhaps more importantly, what are the options available to get back out of the cloud should the need arise?

Coping with public cloud computing is not just a matter of reading contracts carefully to look for gotchas or getting some cost-monitoring tools running. There are technological considerations as well. Are you going to migrate whole VMs to the cloud, or break the workloads out of VMs and use cloud-native services?

Will you use your existing application architecture or break your application into microservices? Will the databases you use be traditional on-premises databases or those offered natively by a given public cloud platform? What are the consequences of all of these decisions?

Public cloud, private cloud, or Internet of Things (IoT) lightbulb as a service, and certain questions hang over everything modern IT teams do. What data will you keep? How will you secure it? What underlying infrastructure will you use? How are you going to store your lakes, oceans, or watery planets of data?

Will you rely entirely on cloud storage or turn to the costly but rock-solid traditional storage array? Will you adopt the less traditional but significantly easier to use hyperconverged infrastructure (HCI) option? And which networking options will you use to put it all together?

There are, of course, an unlimited number of IT vendors who will tell you that they are more than happy to help. But the hard reality is that each of them can only help with a little piece of the puzzle. Today's IT requires multiple vendors working in concert, and that presents its own set of challenges for IT practitioners.

Hitachi + VMware = Superpowers

There's a limit to how many management interfaces, monitoring tools, APIs, and “single panes of glass” any sane individual or IT department can work with. The human brain can only contain so much knowledge, yet specialization brings its own problems to the table.

When IT teams hyper-specialize, knowledge and experience can become stranded. Silos emerge, and accomplishing any task starts to rely on rigid business processes to establish and maintain communication across teams. This creates a yo-yo effect where organizations alternately attempt to break down walls and force sub-teams to communicate, and then engage in proliferating specializations to cope with the flood of emerging technologies and their consequences.

In walking the narrow path between stagnation and recklessness, you need allies. You need integrated solution stacks, carefully curated by experienced vendors, whose purpose is to minimize technological translation problems and keep management manageable.



You need hybrid IT. You're not alone. You probably already have hybrid IT; you just need it to be easier to use, simpler to deploy, and provided with a level of vendor support that means you can focus on advancing the needs of your organization, and not simply on keeping the lights on.

You need to focus not on technology for the sake of novelty, but on technology that gets you the specific results you need within a budget you can sustain. You need public-cloud-like simplicity and private-cloud-like control of sensitive data.

Is this asking for a unicorn? Not really. You need hybrid IT. You're not alone. You probably already have hybrid IT; you just need it to be easier to use, simpler to deploy, and provided with a level of vendor support that means you can focus on advancing the needs of your organization, and not simply on keeping the lights on.

It isn't unreasonable to be looking for pre-engineered, end-to-end managed solutions with Day 1 guaranteed service-level agreements (SLAs), self-service catalogs, and outcomes tying costs to business usage. As we head into the 2020s, these should be table stakes for any enterprise vendor's offerings.

The Partner Advantage

Hitachi Vantara, together with VMware, delivers solutions that allow organizations to efficiently manage, govern, and mobilize key data services across data centers and clouds, both public and private. This reduces the complexity and cost of IT, while offering increased agility.



The experience, flexibility, and ease of use provided by Hitachi Vantara and VMware drive accelerated innovation and enable successful digital transformation. Hitachi and VMware are committed to accelerating customers' digital transformation journey with deep integration that speeds deployments. This integration spans three levels:

1. **Modern infrastructure.** Automated, software-defined, flash accelerated, hyperconverged, private/hybrid cloud; with data and professional services that are consistent with other Hitachi solutions.
2. **Simplified IT Delivery.** Full integration with vSphere, vRealize, NSX and more, allowing for a quicker, secure implementation of ITaaS, private/hybrid cloud and IT automation.
3. **Diversified consumption.** On-premise, managed, or "as a service."

Technologies such as HCI have simplified physical infrastructure to the point that traditional infrastructure specialists are empowered to refocus their considerable talents on emerging technologies. Meanwhile, VMware—by far the dominant provider of virtual infrastructure to today’s enterprises—has been building out a diverse and capable portfolio of enterprise-ready multicloud technologies.

Hitachi Vantara and VMware bring decades of enterprise IT experience to the table and offer customers efficiency, integration, and choice. The sum of this partnership is greater than either of its parts: Hitachi Vantara and VMware have a partnership of long standing and have had time to fine tune the bits that matter. This applies not only to the technologies involved but to the messy business processes, red tape, and human communication that must be sorted out for any vendors to work effectively together.

Hitachi Vantara and VMware both have broad product portfolios, and customers have many options for tackling their unique challenges. But VMware and Hitachi have quite a bit of experience assembling these products into stacks, allowing for real-world integration and simplicity. Together, Hitachi Vantara and VMware build upon the skills and experience already present in today’s enterprises to prepare them for the future.

CHAPTER 2

Data-Driven Digital Transformation

A long time ago, in a galaxy far away, computers were people. “Computer” was a job description, meaning a mathematically competent minion who did calculations for you. Your computer may or may not have used that hot new device, the slide rule, invented in 1622. Alternately, they may have used an abacus, counting board, or exchequer.

The word “computer” did not evolve to mean a calculating machine—infrastructure as opposed to staff—until the late 1890s, and it did not acquire the modern meaning of a programmable digital device until about 1945. Computers as we think of them today—in which application software is so abstracted from the hardware that it could be developed independently of the infrastructure upon which it would execute—didn’t emerge until the late 1970s.



Early mainframes had very little storage by modern standards. They could hold about 1,000 words or 2,500 numbers. Storage media consisted of mercury, metal tape, or magnetic drums. Common complaints about computers of the day were that they were bulky, complex, power hungry, and chronically short of storage. So, very much like the computers of today.

Computers in the 1940s and 1950s were unwieldy and complex machines; the smallest among them could be the size of a bull, although they were more often the size of a bus or even a building. The cabinetry that contained and supported these monstrosities was called the “main frame,” and this eventually became the name for the device itself: mainframe.

Some mainframe computers had a terminal that allowed a user to make inputs and receive outputs, but that’s all that the terminal could do. The actual calculations took place inside the hulking heap of metal, or “big iron.” The inputs to these machines might be made by flipping switches, manipulating a plug board, or inserting a punch card.

Ensuring that these enormous computers received the correct inputs was a specialized job, so if you needed something computed (ballistics trajectories, for example, since there was a war on), you went to the computer, made sure no one else’s calculation was currently running, and asked the input worker to put your problem into the machine. You then waited for a result.

Early mainframes had very little storage by modern standards. They could hold about 1,000 words or 2,500 numbers. Storage media consisted of mercury, metal tape, or magnetic drums. Common complaints about computers of the day were that they were bulky, complex, power hungry, and chronically short of storage. So, very much like the computers of today.

Transistors and integrated circuits really started the ball rolling on miniaturization and eventually gave rise to Moore’s law: that the number of transistors we’re capable of building into an integrated circuit doubles approximately every two years. The number of transistors in an integrated circuit is a rough proxy for speed and inversely for cost.

Integrated circuits made possible the microprocessor and, thus, the personal computer. Personal computers were a revolution: they could do calculations on their own and could fit on a desk. Ease of use was

a careful consideration, resulting in the ability to input data into the computer without requiring a dedicated specialist.

Applications could be developed for personal computers in high-level languages and then compiled, meaning that the development of applications could occur without needing to know every quirk of the underlying hardware. This was the DevOps of the day—bleeding edge, emerging technologies type stuff, even if those particular terms wouldn't themselves be in use for decades.

The Churn

The development of the personal computer firmly established a pattern of technological progression in IT: a cycle that has repeated multiple times in the decades that followed. Computer use became less of a specialized profession and was eventually made more widely. IT infrastructure went from being bulky and complex to being miniaturized, simplified, and under the user's direct control.



Centralized computing devices include not only mainframes but also servers and clouds. The persistence of each is important: there may be a lot more smartphones than PCs, for example, but PCs are not going away in the foreseeable future.

For all their attractions, compact and simple devices also had their limitations. Personal computing devices weren't—and today still aren't—exactly superstars in terms of memory or processing power. Centralized computing capabilities still had their uses, ranging from more number crunching power, to more memory, to the ability to share data between users more easily.

A dynamic emerged in which some tasks would be performed on an end user's device, and some tasks would be performed on a central computer platform. Both sides of this equation evolved over time, with each iteration never quite wiping out their predecessors.

End user computing devices today include not only personal computers but also smartphones and wearable devices. Centralized computing devices include not only mainframes but also servers and clouds. The persistence of each is important: there may be a lot more smartphones than PCs, for example, but PCs are not going away in the foreseeable future.

Similarly, there are quite a few more servers than there are mainframes, but mainframes are far from dead, and whereas mainframe market share may be largely stagnant, the volume of mainframes shipped keeps growing.

This says something about the scope and scale of computing as a market. By and large, mainframes are considered a relic among IT practitioners, despite being a vibrant, growing market. They are relegated to insignificance in the general public consciousness only because other implementations of “computer” have dramatically overtaken them in sales. The rate of change in IT makes it difficult to keep anything in context.

Virtualization

The incredible rate of change in computing has forced the development of technological solutions whose primary purpose has been to pretend that change isn't happening. Virtualization, which has been around in one form or another since the 1960s, is arguably the most successful of those technological approaches.

Depending on the implementation, virtualization has two primary aims: emulation and consolidation. Emulation allows applications to operate on newer hardware without modification. Consolidation

allows multiple applications to execute on a computer simultaneously. The details of how this is accomplished have varied over time, but the core goals of virtualization are as true today as they were decades ago.

Part of the churn in IT is that when new foundational technologies emerge, there is a burst of innovation. Dozens, sometimes even hundreds of startups appear, each trying to solve a problem using the new technology. Only a few survive, and even among these, one approach typically becomes dominant.



The incredible rate of change in computing has forced the development of technological solutions whose primary purpose has been to pretend that change isn't happening. Virtualization, which has been around in one form or another since the 1960s, is arguably the most successful of those technological approaches.

The emergence of personal computers was no different, and when the dust settled, the x86 central processing unit (CPU) ruled the roost. Inexpensive and reasonably powerful, it would become the dominant computer processor from the late 80s until today.

The emergence of the World Wide Web, however, caused a massive acceleration in demand for computers. Personal computer adoption was exploding, and everyone wanted to get on the internet to communicate, consume content, and exchange information. This created something of a slow-moving crisis, one which was overshadowed by the twin headline grabbers of Y2K and the dot-com boom and bust.

Demand for servers was not just outstripping supply, if we collectively kept proceeding with the old-school approach of putting one workload per physical server, very bad things would happen. We would run out of physical space to put servers. We wouldn't be able to generate

enough electricity to power them. That many servers would be impossible to manage with the management tools of the day, and all of this was ridiculous because most servers were running at single-digit utilization anyways.

Bringing Virtualization to the Masses

VMware brought x86 virtualization to the mass market, simultaneously solving most of the management issues through something that behaved a lot like emulation and the scale issues through consolidation. Multiple VMs could be run simultaneously on a single server, each with their own operating system (OS) and associated execution environment.

More importantly, VMware provided the ability for any x86 workload to be made highly available (HA), meaning that if a physical computer died, those workloads could be restarted on another physical computer. Clustering multiple physical computers to provide HA had been possible before VMware but only in extremely specialized cases where the individual application supported it.

VMware brought HA to all x86 workloads and followed it up with vMotion: the ability to migrate workloads between physical computers without taking the workload offline. This was, from a management point of view, nothing short of a revolution, and it dramatically changed end user expectations of application availability, service uptime, and more.



Cloudy, Cloud, Cloud, Cloud

There are always multiple factors that lead to foundational changes in IT. Virtualization technologies had been around for decades before VMware, but none of them seemed relevant until the emergence of the World Wide Web caused a crisis of scale that demanded action. Cloud computing has a similar history.

A cloud is, at its core, a collection of automation technologies married to an orchestration layer with a user-friendly front end. Computer automation has existed since the earliest days of computers, and even much of what we think of today as “cloud computing” existed in the form of timesharing on mainframes, or virtual private servers (VPSes) (though they weren’t called that yet) a decade before Amazon rolled out AWS.



Where things got messy—and where they remain messy for many organizations today—is that although VMware’s x86 virtualization dominated the enterprise data center, the major public clouds were not built on VMware’s hypervisor or upon their management products.

It was VMware’s development of x86 virtualization that kick-started cloud computing. VMware solved a number of critical management challenges that made responding to scale possible, and these approaches were eventually copied by competitors, including open source projects. Layering automation on top of a modern, VMware-like x86 virtualization allowed the development of modern cloud computing—resilient, highly available IT that could be consumed as a service.

Where things got messy—and where they remain messy for many organizations today—is that although VMware’s x86 virtualization

dominated the enterprise data center, the major public clouds were not built on VMware's hypervisor or upon their management products. In fact, each of the major public cloud providers used a different hypervisor and each has radically different approaches to automation, orchestration, and end user interface.

VMware responded by developing cloud capabilities of their own, and this gave rise to the modern enterprise private cloud. A decade after the emergence of Amazon Web Services (AWS), the promised landslide abandonment of the on-premises data center for the public cloud didn't happen, and so here enterprises sit today with a toe in multiple worlds. This is a problem that is being compounded by the development of useful but proprietary public cloud services and the rapid march of enterprise-focused management innovation from VMware.

Just as organizations were unwilling to simply rewrite all of their mainframe applications for PCs, abandoning what could be hundreds of millions of dollars' worth of investment, today's organizations have proven reluctant to abandon VMware's x86 virtualization and private cloud technologies—to say nothing of the myriad workloads operating on top of these platforms.

Containers

When discussing any technology in IT, context matters. How VMware turned x86 virtualization into an empire only makes sense if you understand the context in which VMware's ascendance occurred. Similarly, containers can only truly be understood in the context of their chief competitor: the x86 VM.

x86 virtualization works by creating a complete virtual computer into which an OS can be installed. There are virtual CPUs, virtual hard drives, virtual network cards, and so forth. For the most part, an OS doesn't know it's living inside a VM; it thinks it has complete control of a physical computer. This was purposeful: when x86 virtualization went mainstream, the OSs of the day were designed to take complete

ownership of a computer and behaved as though they were the only tenant of that hardware.

x86 applications were similarly designed: the overwhelming majority of them, at least up until the turn of the millennium, were created to behave as though they were the only application installed on a given computer. OS vendors of the day encouraged this practice. Even as their guidance on application development evolved, application development practices mostly didn't, leading to very real stability and manageability problems for x86 applications.



For the most part, an OS doesn't know it's living inside a VM; it thinks it has complete control of a physical computer.

The short version of the x86 application crisis is that, one way or another, it is safest if x86 applications “believe” that they are the only application installed, while being isolated from any other applications that need to execute on the same physical hardware. This allows the application in question to behave in whichever unfriendly—and often insecure—fashion that its creators feel is required, without actually breaking anything.

x86 virtualization solves this problem by giving each application its own OS inside its own virtual computer (see **Figure 2-1**). It's a nice, neat solution to a very messy problem, and it's one that prevented a massive scale crisis in the early 2000s when broadband Internet connectivity changed the world.

Unfortunately, OSs are resource-hungry beasts. A modern OS can consume 2GiB, or even 4GiB of RAM while sitting around doing absolutely nothing, with no applications even installed. This is problematic for a number of reasons.

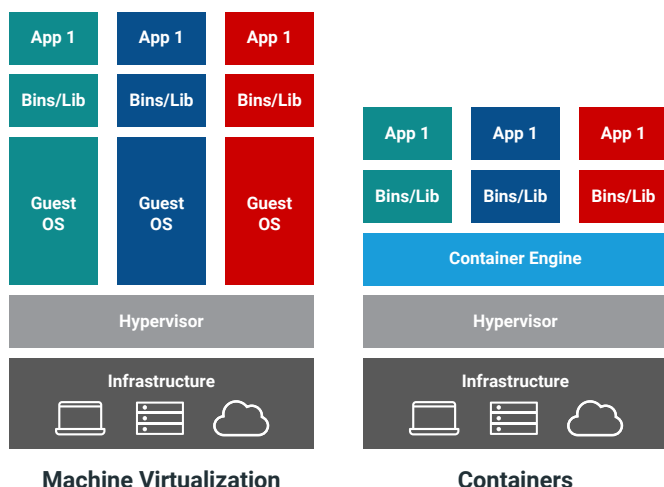


Figure 2-1: Comparing virtual machines and containers, from an architectural standpoint.

RAM is manufactured in the same type of fabrication facilities as solid state drives (SSDs), and both have been experiencing perpetually increasing demand for at least the past decade. Furthermore, technological progress in cramming more and more transistors into ever smaller spaces—“process shrinks” that enable both increased density and lower prices—is very quickly going to be physically impossible.

Simply put, we’ve reached the point where individual transistors are now so small that we’re measuring their size in individual atoms. There isn’t any smaller we can go. As it is, manufacturers already have to deal with quantum effects with transistors that small, meaning that the days of cheap server RAMs doubling in size every few years are, for the most part, behind us.

Containers are one possible answer to this problem. Containers isolate applications, meaning that they provide a means to run multiple applications in a single OS without letting any one application step on another. Applications can still behave badly, be poorly designed, and attempt to do silly things like put files where they shouldn’t, but the

containerization technology lies to these applications about the environment they occupy, which sandboxes them and prevents them from harming anyone but themselves.

The end result with containers is that multiple applications can be run on a single physical host without needing to spin an entire VM for each application. This is great for resource usage efficiency, but it comes with a pair of rather large caveats.

All OSs need to be patched on a regular basis, and patching an OS requires rebooting it. Cram 100 applications into a single OS, and when you reboot that OS for patching, you take down 100 applications. Currently, no version of “vMotion for containers”—where individual containers migrate to another container OS without taking down the application—has achieved any sort of widespread adoption, in part because of the second caveat to containers.



Containers isolate applications, meaning that they provide a means to run multiple applications in a single OS without letting any one application step on another.

The second caveat to containers is that outside of some edge case configurations, all containers deployed onto a single OS share that OS’s kernel. This means that if, for example, you have an application that needs a newer kernel, you must either patch and reboot the underlying host OS or tear down that container and move it to a host with a newer kernel.

This has some related problems, especially for live migration. Moving applications between host OSes with different kernels is extremely difficult, if not borderline impossible. Kernels differ not only because of the version of the kernel but also because of individual kernel modules loaded for that kernel; hardware drivers are one example. Two

computers running identical kernels but ever so slightly different hardware can and do count as different kernels for the purposes of live application migration.

The end result of this is that containers are unlikely to evolve vMotion-like live application migration capabilities with the stability and ease of use required for enterprise adoption anytime soon. This relegates containers to use cases where workloads can be easily created and destroyed, especially those use cases in which the creation and destruction of individual workloads do not affect the availability of individual workloads.

Application Design

In true IT fashion, this process is cyclical. Before VMware, highly available applications were few and far between because clustering could only be easily achieved if it was directly built into the application. VMware made HA mainstream by ensuring that any x86 application running on any supported x86 OS could achieve HA by using VMware's virtualization platform in combination with a centralized storage product.

Today, in order to achieve HA when using containers, applications are broken into smaller pieces—called microservices—and each individual microservice is made resilient by using application-level clustering or by front-ending the microservice with some form of load balancer.

Traditionally designed applications—today called monolithic applications—were generally either a single executable or a small number of executables. These applications were (and still are) designed to operate together in a single environment. This design philosophy goes all the way back to the era when each application was the only one to execute on a given physical system, and any form of HA was typically an afterthought.

Where and when HA did occur in monolithic applications it would typically occur in one of two ways. The first would be the use of shared storage: the application would execute on more than one physical system, but all of them would share a single, centralized storage system. If one physical system were to die, the others would pick up the slack. This is traditional clustering.

The other approach is replication. Multiple instances of an application exist on multiple systems, each with their own independent storage system. Every transaction that occurs on one system is replicated to the other(s). Both cluster- and replication-based HA comes in an active-active and active-passive failover, depending on the application.

In a monolithic application world, there is one application, and it either has native HA capabilities or it doesn't. If it does, then there is usually only one approach employed, and that's that. The microservices approach to HA is different.

A microservices-based application is actually a whole bunch of much smaller workloads working together pretending they're a single application. Each individual microservice can have its own approach to HA (see **Figure 2-2**).

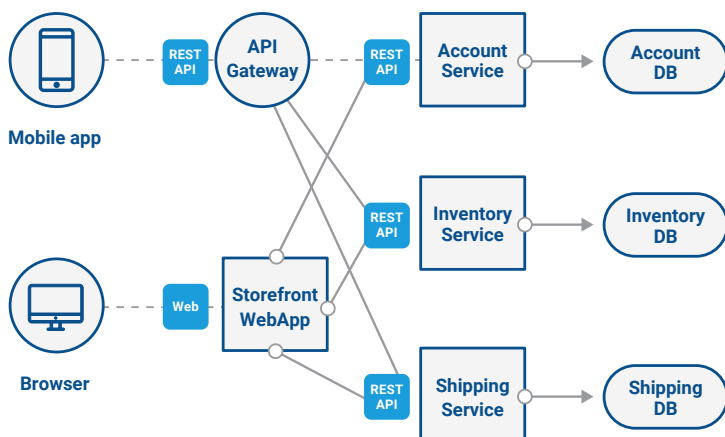


Figure 2-2: A sample microservices architecture.

Code that executes in a web server, for example, might rely on the ability of web servers to be front-ended by load balancers to handle HA. Load balancers themselves have application-native clustering capabilities and so do databases. Other executables that make up a microservice may have no HA capabilities at all—meaning that they are either reliant on operating inside a VM (instead of a container) if HA is desired or that taking down that microservice results in an outage for the entire application.

Infrastructure Meets Application Design

An application's microservices need to communicate in order to interoperate. Most microservices-based applications do this using a message bus. There are a number of different message busses. Without getting too into the weeds on them, the purpose of a message bus is to enable a standardized means of information exchange between different application components.

How message busses are deployed is up to the individual organization. Many organizations deploy a message bus architecture for each application. Others stand up an organization-wide message bus infrastructure and subscribe all their microservice-based applications to it. Still other organizations use the message busses provided by public cloud providers.

This adds an additional layer of complexity—and planning—to modern application design. In addition to having to think about the many and varied ways in which individual microservices might achieve HA, attention must be paid to how those microservices communicate and to ensure that both the message bus and the underlying network infrastructure are available to allow that communication.

This leads to a requirement for increased testing of novel failure modes. What happens to your application if all the microservices are up, but some fraction of them becomes isolated and unable to communicate with the rest?



In addition to having to think about the many and varied ways in which individual microservices might achieve HA, attention must be paid to how those microservices communicate and to ensure that both the message bus and the underlying network infrastructure are available to allow that communication.

Care needs to be paid both to where individual microservices instances are placed—there's little point in putting both database container instances on the same host, for example—as well as to what the communications pathways between the various microservices groupings look like. In the real world, organizations frequently use a combination of VMs and containers to solve this problem, often as much for ease of management as any attempt to provide HA.

Placing a handful of containers in a VM and then running multiple VMs on top of an HA x86 virtualization cluster can, for example, make it logically simple to group containers based on failure domains as well as to provide the ability to live migrate and provide HA to those groups of containers via the underlying x86 virtualization platform.

Here, infrastructure and application design feed into one another, again bringing us to the beginning of yet another IT cycle. This intermingling of application and infrastructure design, while providing significant efficiencies and enabling unprecedented scale, is on the other end of the spectrum from the decoupling of application design from hardware, which helped catapult the personal computer to mainstream success.

The Future of IT

All of the above is further complicated by the blurring of the lines between emerging technologies and traditional IT practice. An individual application can have microservices in multiple public clouds, spread across multiple on-premises data centers, with ingest handlers located in edge computing facilities provided by service providers (such as at the bottom of a cell tower).

This application could have some components built much like traditional monolithic applications—some components that use standardized stacks of third-party applications to execute and some components that use emerging runtime environments, such as serverless. Although such massively distributed applications are still rare today, they are becoming more common, and enterprises frequently state that it's a lack of hybrid/multicloud infrastructure expertise that's holding them back from developing the applications of this scale that they wish to build.

Microservices-based applications that exist across multiple infrastructures (multiple public clouds, on-premises, as well as in edge DCs) are often used to ingest vast quantities of data and then to make sense of them. These applications frequently take in data from IoT devices at the edge and attempt to apply ingest filtering to reduce the firehose of data into something a little bit more manageable. The most famous example is probably CERN throwing away 99.9%-plus of data collected from the large hadron collider (LHC).

CERN throws away most of the data obtained from the LHC, in part because despite having one of the largest IT budgets in the world, it doesn't have anywhere near enough storage space—or network capacity—to handle the total sensor output generated by the LHC, even though all that equipment is located on the same campus. In an IoT context, this sort of distributed application may be handling sensor inputs from sensors monitoring an oil pipeline or from a shipping port tracking the millions of containers that pass through it.



One application, though massively distributed, with different components of different sizes and different requirements running across multiple infrastructures: this is the future of IT. And there are a lot of moving parts to get here.

This data would have to be filtered and then sent somewhere else—likely the public cloud—for in-depth analysis, as edge data centers don't have the compute lift to do more than basic discarding of inputs. The results of that analysis may be then sent to an on-premises data center for secure storage or integration into line-of-business applications within the organization.

One application, though massively distributed, with different components of different sizes and different requirements running across multiple infrastructures: this is the future of IT. And there are a lot of moving parts to get here.

From an infrastructure operations standpoint, even once we get past which infrastructures individual microservices will operate on and what their individual availability requirements will be, these applications will have microservices that are updated independently. Using modern developers like continuous integration/continuous delivery (CI/CD) means that some or all of the application's individual microservices may be updated multiple times per day.


Individual enterprises can have thousands of applications. Although very few of these applications—at least over the next ten years—will fall into the massively distributed, multicloud nightmare territory described above, many enterprise applications are already at least hybrid applications, with components that operate on-premises as well as in at least one public cloud.

The Great 'Pets Vs. Cattle' Debate

The automation versus manual administrator discussion is known in the IT industry as the “pets vs. cattle” debate.¹ This is based on a presentation frequently given by Randy Bias, VP Technology—Open Source Software at Juniper Networks.



In the old way of doing things, we treat our servers like pets, for example Bob the mail server. If Bob goes down, it's all hands on deck. The CEO can't get his email, and it's the end of the world. In the new way, servers are numbered, like cattle in a herd. For example, www001 to www100. When one server goes down, it's taken out back, shot, and replaced on the line.



¹ <http://cloudscaling.com/blog/cloud-computing/the-history-of-pets-vs-cattle/>

Automation, Anywhere and Everywhere

If the above seems daunting, bear in mind the real underpinning of clouds: automation and orchestration. If your goal is to continue to operate IT in the traditional manner—where administrators install every OS, manually configure every application, and spend their time individually configuring every switch, router, virtual host and so forth—it's time to retire. The future won't be kind.

Enterprise IT in the 2020s will not be possible by taking an entirely manual approach. It arguably hasn't been possible for quite some time. Automation is, unquestionably, an absolute requisite for IT at scale.

The applications of automation are both subtle and gross. They extend from automating individual microservices within an application to the creation of fully autonomous data centers.

To be clear: full automation of IT is likely impossible. There will always need to be a human in the loop somewhere, but the purpose of most modern pushes for IT automation and digital transformation is to automate as much as possible so that expensive and difficult-to-source talent can be refocused on the tasks that we can't yet offload to computers.

In the context of the monster distributed application discussed above, where would you rather spend your IT practitioner's hours: manually rebooting VMs and instantiating containers or contemplating the application and infrastructure design as a whole in order to identify points of failure and design solutions to them?



To be clear: full automation of IT is likely impossible. There will always need to be a human in the loop somewhere, but the purpose of most modern pushes for IT automation and digital transformation is to automate as much as possible so that expensive and difficult-to-source talent can be refocused on the tasks that we can't yet offload to computers.

If monster distributed applications are few and far between, there are far more mundane, but no less important, architect-level tasks that need addressing. Extending today's software defined data center (SDDC) beyond the on-premises data center and into the public cloud is an important precursor toward enabling the sort of IT infrastructure

that application developers will be demanding in the next decade and that will occupy many organizations for the next several years.

Establishing a common operating framework across private data centers and public clouds as well as creating a uniform underlying infrastructure to enable seamless workload mobility between infrastructures both fall into this category as well. Both involve lots of planning, lots of testing, and lots of business processes and rules, all with the end goal of scalability and automation and preferably accomplished in such a manner as to make future development easier to get a handle on.

This then is the real future of IT: planning. Playing the long game. Choosing frameworks and automation platforms, cloud providers and even message busses that are designed to provide stability over the entire lifecycle of applications that might contain thousands of individual components, stretched across multiple infrastructures, only some of which your organization controls.

CHAPTER 3

Your Data in a Hybrid Cloud

Control and trust are inextricably intertwined in any discussion about the future of IT. This is only becoming more apparent as organizations move beyond infrastructure as a service (IaaS) and the point application of software as a service (SaaS) and start adopting platform as a service (PaaS), serverless, and other riskier products and services.

When organizations first moved beyond the on-premises data center, they did so in a purely transactional manner. Adopting Office 365, for example, was not much of a stretch. Microsoft® on-premises to Microsoft in the public cloud made sense; it was a like-for-like replacement.

When it first launched, Office 365® was largely self-contained. E-mail and IM as a service, some cloud storage, and Office. It could integrate with our on-premises active directory. It was largely benign. If you really wanted, you could walk away from Office 365. Email and IM as a service are a dime a dozen, and Office can still be had in the non-cloud version for those who want it.

This conversation completely changes if you're building core capabilities around a public cloud vendor's proprietary service. Let's consider a simple facial recognition application.

It takes a few minutes to write an AWS Lambda function with a trigger based on the addition of a new file to an S3 storage bucket. When a new file is added, the serverless function is triggered, and that function runs the image through Rekognition, Amazon's facial recognition



When organizations first moved beyond the on-premises data center, they did so in a purely transactional manner. Adopting Office 365, for example, was not much of a stretch. Microsoft on-premises to Microsoft in the public cloud made sense; it was a like-for-like replacement.

as a service offering. The output of Rekognition could then be stored in Amazon's Aurora database and analyzed using yet more proprietary tools.

At any point, Amazon can change how any part of that stack works. The changes can be breaking, or not, as Amazon deems fit. With a standalone cloud application or service, it's annoying, but not critical, if the vendor changes something. As soon as anything "as a service" becomes one part of a much larger integrated whole, however, even seemingly minor changes can have a dramatic effect.

With traditional IT, if you didn't like something the vendor had done with the latest update, you could simply choose not to update. If a vendor made changes to an API, or a framework, you could delay implementing the new API or framework until your developers had a chance to catch up. This isn't always the case anymore today, and it places trust in the vendors involved at the very center of our relationships with them.

Our relationships with our vendors must be true partnerships. They can no longer merely be transactional. Any vendor that demonstrates contempt for the importance of change management—or for their customers—is difficult to trust and, thus, becomes a risk.

Nowhere is the importance of trust more apparent than when dealing with your organization's data.

Data Locality

Data locality is a term that is often associated with Hadoop clusters, but it applies to a wide variety of different technologies, including (but by no means limited to) HCI. The focus of data locality is to ensure that the data being processed is either logically or physically proximate to whatever is processing said data.

From a technical point of view, this can be solved in a few different ways. Using the example of large object stores, it can mean having multiple access nodes, such as a name node or a shim presentation node. Multiple access nodes allow compute nodes accessing the object storage to use the access node closest to them, preventing unnecessary network traffic. Smart storage clusters will track which systems access which data and will store that data on nodes closest to where they are used.



The focus of data locality is to ensure that the data being processed is either logically or physically proximate to whatever is processing said data.

Most HCI systems work in a similar manner. With HCI clusters, there are typically at least two copies of any VM kept on the cluster—one copy on the node where that VM is executing and one copy somewhere else on the cluster.

Although it is certainly possible for a HCI node to execute a VM whose storage resides on another node, the underlying storage fabric would eventually move the blocks for that VM to the node where it was executing, ensuring that the VM in question was able to operate with the lowest possible latency.

Data locality has legal implications as well. Different jurisdictions have different laws about data storage as well as data processing. This can have serious implications, not only in the design of large, distributed

applications but also in traditional IT where it's operated in urban locations that straddle a jurisdictional border.

Consider, for example, the Kansas City metropolitan area. This metropolitan area straddles two jurisdictions: the state of Kansas and the state of Missouri, both in the US. Google famously laid fiber throughout the metropolitan area, offering inexpensive fiber to the premises. This predictably led to a boom of IT-based innovation, including the creation of data centers, both big and small.

An organization operating in this area may reasonably place one data center on the Missouri side of the border and another on the Kansas side of the border. This would allow them to have metro area redundancy and take advantage of technologies, such as VMware's stretched clustering, to provide an extra layer of resilience to traditional applications.

Doing so, however, would mean that the organization in question would have to abide by state-level regulations regarding privacy, data security, and so forth for both states, even for a single workload. As a result, data locality can have both technical and legal implications, even though it is primarily discussed as a technical consideration.

Data Sovereignty

Data sovereignty is the data locality discussion writ large, and it is typically a discussion about national or supra-national jurisdictional issues regarding data ownership and access.

The concerns of data sovereignty typically include complex matters, such as when two jurisdictions (or even a single jurisdiction!) demand that an organization comply with conflicting laws. Perhaps the most famous such case is *Microsoft Corp. v. United States*.⁴

The goal was to ensure that the U.S. government could get access to any of the data, regardless of the laws of the nation in which the data

⁴ <https://www.lawfareblog.com/microsoft-ireland-case-supreme-court-preface-congressional-debate>

Microsoft vs. the United States

The short version of this extremely long case is that an individual suspected of a crime, which was being prosecuted in the U.S., had data stored on a Microsoft Ireland server. Microsoft Ireland is a separate legal entity from Microsoft, which is incorporated in the U.S., and must obey EU laws. The U.S. government wanted the data on the Microsoft Ireland server and found a judge that was willing to demand Microsoft (U.S.) to turn it over.



This was widely viewed as an attempt to set legal precedent. If the U.S. government really cared about the data, there were established legal channels for it to request the data via extant law enforcement cooperation agreements with Ireland. This wasn't the point; the point was to establish precedent that would allow any court in the U.S. to demand that any U.S. company must turn over any data to which it had access (Microsoft U.S. has access to Microsoft Ireland's servers), if requested.

resides, or the legal relationship between the different entities of which the data was requested. The legal theory went that if you were a U.S. organization and you could access data elsewhere that the U.S. government wanted, you must provide that information, regardless of the cost or risk to your organization.

Microsoft's legal victory in the appeals court resulted in the passage of the CLOUD act.⁵ The CLOUD act effectively codifies the U.S. government's position in *Microsoft Corp. v. United States* into law, bringing it into direct conflict with data security and privacy laws all around the world. The consequences⁶ of this have yet to be fully revealed.

⁵ https://en.wikipedia.org/wiki/CLOUD_Act

⁶ <https://legalict.com/2018/11/08/the-cloud-act-and-its-consequences/>

Privacy

Privacy, while often discussed and especially so when talking about data in a cloud context, is a somewhat nebulous concept. Legally, it is poorly defined in many jurisdictions and subject to conflicting and even outright contradictory requirements all over the world.

Depending on the jurisdiction, ensuring a user's privacy may involve preventing end users from seeing one another's information or preventing any employees from viewing a user's information unless viewing that information is essential to doing their job. Privacy may require preventing suppliers and partners from viewing that user's information unless absolutely essential, anonymizing end user information, and even keeping that information from foreign governments.

Because privacy requirements can change from jurisdiction to jurisdiction, complying with privacy regulation increasingly involves tying user data to a given location and may involve restricting processing of that data to a given location as well. Privacy issues are also deeply associated with data sovereignty issues as well as data security, encryption, and general IT security.

GDPR

The General Data Protection Regulation (GDPR) of the EU is the standard candle against which international privacy regulation is judged. Though several years old, the GDPR came into force on May 25, 2018, and the full ramifications are still being determined.



The General Data Protection Regulation (GDPR) of the EU is the standard candle against which international privacy regulation is judged.

By and large, the GDPR gives European citizens control over their data and over who can use it. It restricts where organizations can store data on EU citizens, what they can do with it, who can see it, and allows EU citizens to demand that their data be made accurate and/or be deleted.

On paper, there are a number of areas where the GDPR and the CLOUD act come into conflict. These issues are being sorted out at the political level but still represent hypothetical legal risk, especially for large enterprises and social media organizations, which are frequently the target of lawsuits regarding privacy.

Muddy Data Regulations

Although the U.S. and the EU get all the press when it comes to privacy, other jurisdictions also have privacy laws. On paper, PIPEDA in Canada provides strict privacy protections to Canadians; however, these protections were supposedly significantly weakened with the signing of the USMCA free trade agreement. The disparity between PIPEDA and the USMCA has not been challenged in court or addressed in ongoing efforts to strengthen PIPEDA within Canada, which will create uncertainty for several years to come.

Uncertainty is the watchword when it comes to data regulation. Even in the U.S., there is a movement by prominent lawmakers toward greater privacy laws. This⁷ draft bill is an example. It will never become law, but the fact that it was even proposed means that the U.S. is on the path to proper privacy regulations, especially in light of the growing number of state privacy laws emerging within the U.S.

The patchwork of privacy, data sovereignty, and data security regulations around the world presents a very real threat to the emerging multicloud technological paradigm of the 2020s. Organizations are not building distributed applications for fun, nor are they attempting

⁷ <https://www.wyden.senate.gov/news/press-releases/wyden-releases-discussion-draft-of-legislation-to-provide-real-protections-for-americans-privacy>

to create hybrid cloud setups or multicloud setups just because these happen to be latest buzzwords.

These sorts of applications and infrastructures are required to meet the needs of customers, suppliers, and partners in a cost-efficient manner. Above all, they're needed to cope with both scale and the ever-increasing pace of change that since the very beginning has been part of the fabric of IT.



Ensuring that data security is consistent, even across multiple infrastructures from multiple providers, is difficult, with many organizations turning to a single, trusted cloud provider with global reach in order to provide secure data storage.

Ensuring that data security is consistent, even across multiple infrastructures from multiple providers, is difficult, with many organizations turning to a single, trusted cloud provider with global reach in order to provide secure data storage. It is for this reason that Hitachi recently acquired REAN Cloud.⁸

Navigating all of this requires more than simply purchasing the right products or services from a vendor. It requires vendors with experience, talented and diverse legal teams, and a willingness to engage with both other vendors and their own customers as true partners.

⁸ <https://community.hitachivantara.com/community/innovation-center/hus-place/blog/2018/11/11/why-are-we-excited-about-the-rean-cloud-acquisition>

CHAPTER 4

Better Together: The Hitachi Vantara and VMware Advantage

It would be easy to read the challenges in the preceding chapters and want to give up in despair. In three chapters, we have engaged in only the highest level of discussions about the challenges that will be facing enterprise IT in the next decade, and even these can seem overwhelming, and disheartening.

Fortunately, Hitachi Vantara and VMware are here to provide real world solutions to these problems.

Although there are a lot of moving parts in today's hybrid multicloud world, the goal of the Hitachi Vantara-VMware partnership is to ensure that customers can focus on these high-level concerns without also having to get bogged down in the nitty-gritty technical details.

IT decision makers and practitioners alike should be focused on solving large-scale infrastructure design issues, not fretting about how many controllers or disk shelves need to be purchased for the next storage array. Virtual administrators should be focused on deciding which workloads will live on which infrastructures, not trying to build the plumbing to connect on-premises to a public cloud.

Hitachi Vantara and VMware aim to make decades of experience in automation available to customers, so that the focus shifts from keeping infrastructure going, to continually refining ease of use, and rooting out unnecessary complexity, as shown in **Figure 4-1**.

Automate IT with Infrastructure Aware Management

- **Business/app user** – I need xyz VM service today – and it may change
- **VMware admin** – what type of storage data services am I going to get?
- **Storage admin** – what type of application workload are you deploying?

Need policy framework to do this efficiently




Figure 4-1: How Hitachi and VMware collaborate to streamline IT operations.

Over time, complexity inevitably increases. It is a natural consequence of mergers, acquisitions, and the very human response to forgetting about a solved problem, even if the solution to that problem was inelegant. Over time, management applications increase ease of use—often by adding layers of abstraction; however, there is always a battle being fought between ease of use and complexity.



Hitachi Vantara and VMware aim to make decades of experience in automation available to customers, so that the focus shifts from keeping infrastructure going, to continually refining ease of use, and rooting out unnecessary complexity.

Putting aside any product or service details, the singular goal of the Hitachi Vantara-VMware partnership is to make the lives of customers simpler.

Complexity Is About More Than Just Technology

Cloud computing—and especially multicloud computing—brings IT administrators face-to-face with the reality that complexity in IT is about more than just technology. There have been innumerable books written on the topic of human complexity as it relates to IT. That IT teams need to think of their job as delivering a service rather than merely maintaining technologies is a discussion that has continued for decades.

Now, IT teams have to face a reality, in which the very infrastructure designs they are responsible have not only technical considerations, but human and end user considerations, as well as legal, business process, and supply chain ones. Legal/regulatory compliance concerns are often more complex than technical ones, and the upcoming decade will place IT teams squarely in the middle of any number of legal battles. If you have any doubt about this, ask anyone in the IT department at Equifax about the complexities of legal considerations in the modern data center.

In addition, business processes have to align with technology, and there has always been something of a chicken-and-egg problem there. IT staff have always been required to know the newest technologies and how to build, buy, install and manage them, but the constant treadmill of simply knowing what exists is already beyond what any one individual can possibly keep track of, let alone knowing how to implement any of it.

Cloud computing has shown the world that consuming IT infrastructure can be simple. Hitachi Vantara and VMware are demonstrating that installing, configuring, and managing IT infrastructure can also be simple. Complexity is a fact of life in IT, but as the complexity of providing IT increases in areas not directly related to the technology itself, the complexity of implementing that technology should decrease to match.

Multiple Moving Parts From Multiple Vendors on Multiple Infrastructures = Problems!

Anytime you have multiple vendors in the mix, things can go wrong in bizarre but ultimately very human ways. One of the most famous examples of this within the IT community is that of circular vendor finger pointing.

Every IT practitioner who has been around for more than a year will have a story of trying to find support for an application or device that doesn't work and getting the runaround in the process. Applications run on an operating system, which may run on a hypervisor, which runs on a server, which connects to a network, and which may also have external storage. So on and so forth.

The application itself may be the problem, as may any of the bits upon which it relies. When we start talking about cloud applications, there are rather a lot more bits to consider, and hybrid applications make it more complex still. The large, distributed multicloud nightmares discussed in earlier chapters of this book could involve components from hundreds of vendors, and there's no way of knowing which components cloud providers are using or whether they even matter.



Navigating the politics of vendors is, in many ways, its own job. Does it need to be your job?

Complicating matters further is that vendors update asynchronously of one another, which can cause interesting problems that have a tendency to cascade. API drift—in which an application or application component is written for one version of an API but ends up executing against a different version of that API—is increasing a very real problem.

And then there's that fuzzy, grey, and multi-decade area where some applications are legacy, and others are cloud-native, and trying to find the money to modernize applications that "are working fine" is a fight in itself, which is to say nothing of the challenge of IT supply chain management.

Navigating the politics of vendors is, in many ways, its own job. Does it need to be your job?

The Value of Vendors

A good vendor is one that views their customers' relationships as partnerships. It is a relationship in which the vendor is not simply attempting to sell a product or service; rather, the vendor strives to understand the challenges of their customers and to come up with real-world solutions.

The value of a vendor is found in that vendor's focus: are vendors trying to push a way of thinking upon you, complete with SKUs and billable hours? Or are they seeking to find out what you do, why you do it, and want what you want to do in the future so that they can build it for you?



With more than 25 years of operational experience,

Hitachi Vantara provides technology solutions, skilled manpower, experience, and necessary support framework to act as one-stop shop for customers' private, public, and hybrid cloud needs.

Today, the hybrid cloud seems an easier goal than ever before, but this is in many ways an illusion. Hybrid cloud computing—to say nothing of multicloud computing—is terribly complex, and requires an experienced partner to help you navigate the complex planning, implementation, governance, and post-deployment phase.

Intelligent Operations: Hitachi and VMware

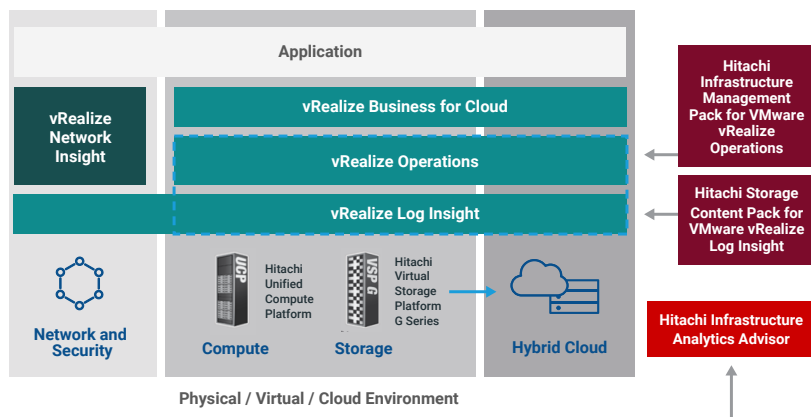


Figure 4-2: The combination of Hitachi and VMware vRealize offers powerful advantages for IT.

With more than 25 years of operational experience, Hitachi Vantara provides technology solutions, skilled manpower, experience, and necessary support framework to act as one-stop shop for customers' private, public, and hybrid cloud needs.

For customers who want an end-to-end solution but don't want to build it themselves, Hitachi Enterprise Cloud (HEC) with VMware vRealize offers customers the ability to implement a pre-engineered, pre-designed, and fully integrated solution with VMware. It's ready to go within 90 days and provides guaranteed, enterprise-class SLAs, a pre-designed service catalog for self-service, and consumption-based pricing to match actual costs with business usage. (**Figure 4-2**)

Hitachi, through both Hitachi Vantara and Hitachi Consulting, has extensive IT resources and a parent company (Hitachi, Ltd.) with vast real-world operation technology (OT) resources. Hitachi makes everything from bullet trains to nuclear reactors, and as a result, it has both a breadth and depth of real-world IT experience and OT resources to draw upon unmatched by any other provider.

The Value of Ecosystem Continuity

Not to put too fine a point on it, but life is a lot easier if you can use VMware everywhere. VMware is already the most popular enterprise virtualization provider, and it holds the pole position for a reason: VMware has the most comprehensive set of offerings around x86 virtualization, private cloud, and hybrid cloud of any of the enterprise virtualization providers available.

Part and parcel of this is that VMware has fully developed automation and orchestration capabilities and has been investing in these management applications for years. Today's VMware cloud offerings are significantly easier to use and more capable than those available even one year ago, and there is a great deal of continued investment slated for the foreseeable future.

With VMware making the software that stitches together both public and private clouds and Hitachi providing the benefit of real-world experience, the Hitachi Vantara-VMware partnership is difficult to beat.

VMware and Hitachi: Better Together

“Customers will be in a much-improved position to respond to the market opportunities with new capabilities, providing agility and simplicity. Newly gained flexibility in finding the right home for workloads, depending upon their profile, will improve operational efficiency, thus freeing up capital for innovative projects.”

Hitachi Vantara brings multicloud management as well as a strong portfolio of cloud migration technologies to the table. This allows customers to stop worrying about the technology itself and focus on the many other hard parts of making technology go.

More important are the intangibles of the Hitachi Vantara-VMware partnership. Hitachi Vantara has access to all the knowledge gained by the other arms of Hitachi, just as VMware brings its own vast

experience to the table. Both vendors have great experience in building and maintaining multi-vendor partnerships, which are important to the deployment of complex IT solutions involving multiple products from multiple vendors.



With VMware making the software that stitches together both public and private clouds and Hitachi providing the benefit of real-world experience, the Hitachi Vantara-VMware partnership is difficult to beat.

Hitachi Vantara and VMware test their solutions to ensure that they work together, and they make sure that third party products and services integrate properly. By operating at scale around the world, both vendors have experience in coping with crazy jurisdictional nightmares and bring the benefits of that knowledge to customers.

CHAPTER 5

Technology Considerations

What do you get when you combine the most reliable x86 virtualization available with technologies, technologists, and support teams that brought the world storage systems with a 100% availability guarantee?⁹ Legendary reliability. The Hitachi Vantara-VMware partnership is one of the few partnerships in tech that is truly greater than the sum of its parts.

Vendors as large as Hitachi or VMware have a diverse portfolio of products to offer. On their own, either vendor has a wealth of IT offerings that cover things cloudy, both public and private. Both have partnerships outside of their partnership with each other, and each competes with the other in addition to working together. Such is the nature of business, especially at enterprise scale.



Hitachi Vantara is a VMware TAP partner and a long-standing strategic vendor, one of the few VMware global OEMs. Hitachi Vantara supports VMware across a broad spectrum of products, and both organizations have created deep integrations across their portfolio.

The history of such a partnership matters. Hitachi Vantara is a VMware TAP partner and a long-standing strategic vendor, one of the few VMware global OEMs. Hitachi Vantara supports VMware across a broad

⁹ <https://www.hitachivantara.com/en-us/pdf/datasheet/ensure-data-availability-for-vsp-datasheet.pdf>

spectrum of products, and both organizations have created deep integrations across their portfolio.

Time and effort have allowed deep expertise to develop in this partnership—expertise that’s not only technical, but cultural and interpersonal as well. In addition to the breadth of VMware adapters, APIs, and plug-ins, both organizations have a history of working together on projects both big and small. But what are the technologies they’re bringing to the table?

Raw Speed

Everyone wants to show off their flagship toys, and Hitachi Vantara believes it can bring the best value and highest performing private cloud solutions to the table. VMware provides the ability to manage these systems across multiple clouds, ensuring that your workloads have the performance they need, both on-premises and off.

Achieving performance worthy of a little braggadocio for any vendor requires the use of NVMe SSDs, and Hitachi Vantara can bring all-NVMe¹⁰ products to the table when you need them. Not to be outdone, VMware offers its own unique capabilities regarding Intel Optane¹¹ to really kick it up a notch, ensuring that, if you’ve got the need, it’ll bring the speed.

VMware Cloud Foundation and vRealize

The core of anything multicloud involving VMware is VMware Cloud Foundation and VMware vRealize. VMware has been working slowly and steadily on cloud management software for over a decade. The result is noteworthy.

¹⁰ <https://www.hitachivantara.com/en-us/pdf/solution-profile/transform-it-with-all-nvme-hyperconvergence-solution-profile.pdf>

¹¹ <https://community.hitachivantara.com/people/DISINGH/blog/2018/09/26/hitachi-vmware-vs-an-intel-optane-nvme-turbocharged-hyperconverged-infrastructure>

Infrastructure Automation for VMware

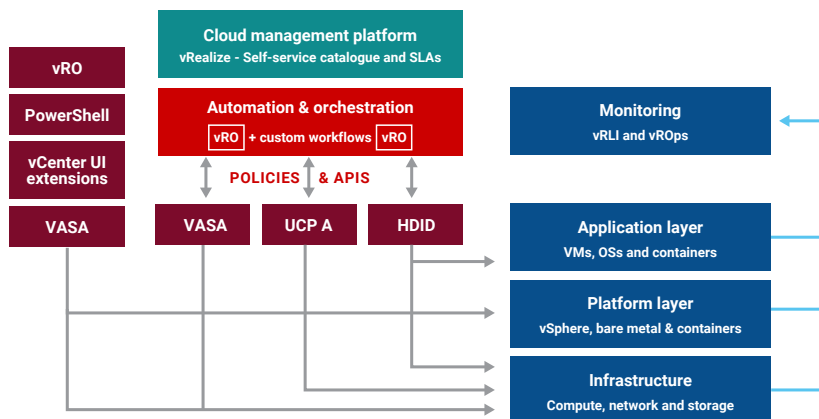


Figure 5-1: If you can't sufficiently automate your infrastructure, you can't modernize your operations. VMware and Hitachi specialize in helping companies do just that.

VMware has built the most approachable, easy to use, and simple to manage multicloud management product stack available today. VMware ESXi is the most widely used hypervisor in the enterprise for a reason: beyond the feature set of the hypervisor, the management capabilities are incomparable.

VMware has the most intuitive graphical user interface (GUI) of any multicloud cloud management solution available. PowerCLI offers extremely versatile and powerful command line administrative capabilities, whereas VMware's feature-rich and well-documented APIs have made its products easy to integrate with, for vendors and customers alike (**Figure 5-1**).

These are not mere marketing tick-boxes. A cursory trawl of both the trade press and community blogs will show not only enthusiastic support for VMware cloud foundation, PowerCLI, and the VMware APIs, but innumerable examples of the virtualization community contributing code, automation runbooks, and more.

Organizations and individuals do not use VMware Cloud Foundation because they have no other choice. They use it because it has proven, time and again, to give them back hours every day.

Hitachi Enterprise Cloud

Hitachi Enterprise Cloud (HEC) is Hitachi Vantara's cloud services brand. Hitachi Vantara offers on-premises/private cloud capabilities via Hitachi Unified Cloud Platform (UCP) products, as well as public cloud services via their REAN Cloud solutions.¹²



Hitachi Enterprise Cloud offers numerous “as-a-service” offerings. These include IaaS, disaster recovery as a service (DRaaS), backup as a service (BaaS), and object storage as a service (OaaS).

Hitachi's UCP offerings range from the hyperconverged Hitachi Unified Platform HC (UCP HC¹³) to the rackscale Unified Compute Platform RS (UCP RS¹⁴), all the way to the converged infrastructure Unified Compute Platform CI¹⁵. Whether you want a private cloud that is only a few servers wide or you want to measure your cloud footprint in acres, Hitachi Vantara has a UCP offering that is battle tested to meet your needs.

In addition to VMware Cloud Foundation and vRealize capabilities, Hitachi Vantara offers additional management capabilities that further increase ease of use. More importantly, Hitachi Vantara—formerly

¹² <https://www.reancloud.com/>

¹³ <https://www.hitachivantara.com/en-us/products/converged-systems/unified-compute-platform-hc-series.html>

¹⁴ <https://www.hitachivantara.com/en-us/products/converged-systems/unified-compute-platform-rs-series.html>

¹⁵ <https://www.hitachivantara.com/en-us/products/converged-systems/unified-compute-platform-ci-series.html>

Hitachi Data Systems—bring award-winning global support to the table, helping your IT teams keep focused on the big picture, and not on keeping the lights on.

Hitachi Enterprise Cloud offers numerous “as-a-service” offerings. These include IaaS, disaster recovery as a service (DRaaS), backup as a service (BaaS), and object storage as a service (OaaS). Hitachi Enterprise Cloud also offers Hybrid Data Servicer (DP/Aspen) to those who need storage that transcends individual clouds.

For those seeking to get big with their data and perhaps crunch a few numbers, Hitachi Vantara offers 47Lining:¹⁶ big data managed services that go beyond simply spinning up Hadoop on someone else’s hardware.

Addressing Real-World Problems

Individual technologies don’t solve enterprise cloud problems. It is in their combination that solutions are found. Data locality, for example, is a problem frequently solved through the use of both Hitachi Enterprise Cloud hybrid cloud capabilities and VMware Cloud Foundation stretch cluster capability.

Together, these allow a single VMware cluster to span multiple data centers, both on-premises and off. Workloads can then be placed where they are needed and where jurisdictional issues allow the data they contain and/or process to exist.

Containers present another set of common problems for today’s enterprises. Containers are designed to house ephemeral workloads that are regularly created and destroyed. Despite the highly temporary nature of containerized workloads, they require persistent storage to be truly useful, in addition to management tools designed to handle their unique requirements.

¹⁶ <http://www.47lining.com/>

Here, VMware and Hitachi have an array of products to choose from. VMware offer their Photon Platform and VMware Pivotal Container Server (PKS) on VMware Cloud Foundation (VCF), whereas Hitachi offers Hitachi Enterprise Cloud Container Platform..¹⁷



Anyone can provide infrastructure, but today's enterprises want to go beyond just metal and blinking lights. They want a smart data center, and Hitachi Vantara delivers here as well.

Everyone needs reliable, high-performance storage,¹⁸ but if your data doesn't exist in at least two places, then it doesn't exist. Hitachi Vantara's solution for this is efficient data protection¹⁹ —a somewhat mundane but nonetheless vital offering that simply must exist, especially when talking to a hybrid cloud provider.

Anyone can provide infrastructure, but today's enterprises want to go beyond just metal and blinking lights. They want a **smart data center**,²⁰ and Hitachi Vantara delivers here as well. Investment in intelligence for IT operations is ongoing, with a fully modern AI Ops²¹ platform also being part of the portfolio.

¹⁷ <https://www.hitachivantara.com/en-us/pdf/solution-profile/enterprise-cloud-container-platform-solution-profile.pdf>

¹⁸ <https://www.hitachivantara.com/en-us/products/storage.html>

¹⁹ <https://www.hitachivantara.com/en-us/solutions/data-center-modernization/modern-data-protection.html>

²⁰ <https://www.hitachivantara.com/en-us/services/infrastructure-services/smart-data-center.html>

²¹ <https://www.hitachivantara.com/en-us/solutions/data-center-modernization/intelligent-operations.html>

The Big Takeaway

With the raft of technologies, branding, and buzzwords discussed, it's worth closing out this book with a simple, honest take on the Hitachi Vantara-VMware partnership. VMware has innumerable partners working to package and resell its cloud capabilities. Hitachi has multiple arms selling technology products and innumerable partners that it works with to do so. Over time, however, they've proven repeatedly to be better together.

Two partners, both with global reach, as well as decades of IT and OT experience await. Why not see what they can do for you?