



Hewlett Packard
Enterprise

enterprise.nxt

Foundations of **Hybrid IT**

EDITED BY DAVID CHERNICOFF & RICHARD MCGILL MURPHY

TABLE OF CONTENTS: FOUNDATIONS OF HYBRID IT



3

Executive summary


By David Chernicoff & Richard McGill Murphy



5

The future is hybrid: An interview with Craig Partridge, HPE Pointnext

By Pam Baker



12

Hybrid IT survival skills


By Steven J. Vaughan-Nichols



23

How to evolve your infrastructure to hybrid IT

By Ken Hess



35

How hybrid IT enables the software-defined data center

By Alyson Behr



45

Packaging your environment for hybrid IT

By Frank Ohlhorst



51

The edge vs. central IT

By Pedro Pereira

Executive summary

By David Chernicoff & Richard McGill Murphy

Hybrid IT provides the flexibility that enterprises need to navigate digital transformation. Hybrid infrastructure enables business growth, a strong ROI and a clear path to the future.

Executive summary

Today, most businesses use a mix of traditional IT and cloud services, combining elements from both ends of the spectrum to meet their business needs, often in an ad hoc fashion.

The strongest enterprises understand that they will need the right combination of next-generation apps, services and technologies that are widely distributed, continually changing and delivered at massive scale. In this future, IT will deliver services when and where they are required, using the appropriate technology for each workload.

In this report, we describe the foundational technologies that make hybrid IT an effective business differentiator.

While certain basic IT functions are fairly common, the most successful business advantages derived from IT are accomplished by those who can deliver the unique solution that best fits the business need and maintains a competitive business advantage. Hybrid IT allows a business to mix and match legacy IT, cutting-edge hyperconverged and composable infrastructures, and the latest in software-defined platforms, as needed, to deliver the desired result.

Across all of this, IT will be tasked to integrate, secure and govern these environments in order to optimize how services are created, delivered and consumed. Traditional IT, managed services, and private and public cloud must combine to power apps and services that meet changing business demands. These services will be deliverable from core IT all the way to the network edge.

Key business goals for IT:

- Increase customer engagement, share of wallet, satisfaction and loyalty
- Create new areas of profitable growth and differentiation
- Reduce risk and lower operational costs by accelerating and driving efficiency

01

The future is hybrid

An interview with Craig Partridge, HPE Pointnext

By Pam Baker

-
- Why hybrid IT matters
 - How IoT and edge computing are affected
 - About the cloud
 - How much time do we have?

Why hybrid IT matters

Hybrid IT isn't just a buzzword or a passing trend. Neither is it a steppingstone to the transition to some future form of IT.

"It's not that traditional or transformational IT becomes irrelevant," says Craig Partridge, director, Data Center Platform Consulting, [HPE Pointnext](#). "It's more of a symbiotic relationship between idea incubation and idea implementation."

Here's more strategic thinking from Partridge on why the hybrid IT model is so crucial and how to get the most from it. The term *hybrid IT* refers to an operating model aimed at harnessing the power of the corporate digital experience. It differs significantly from the traditional IT model of ruthlessly cutting costs in the search for ever more efficiencies. For example, in the traditional IT model, using cloud services is primarily a cost-saving decision, but with the hybrid IT model, using cloud services is more about building agility, flexibility and innovation at top speed.

"What's different now is data," says Partridge. "Hybrid IT is about exploration, about finding your place in the market, about finding a new, more competitive business model. It's about disruption."

But, he says, ultimately it's about more than even those worthy and crucial goals.

"In the end, once you've hit gold, you must use that information, those insights, to transform the organization," Partridge explains. "That requires the expertise in your run organization to handle risk, compliance, governance, security, reliability, efficiencies and other things necessary to the optimum state of business."



We sat down with Partridge to discover more about why the hybrid IT model is so crucial and how to get the most from it. Here are the insights that came from that discussion.

Q: Why does moving to the new hybrid IT model right now matter if we're just going to return to some semblance of traditional IT in the end?

Partridge: The world has already changed and continues to change rapidly. Physical and digital boundaries are blurring more every day. While traditional IT is well scoped, it's not set up to explore this undefined but exciting, new reality.

Traditional IT is geared toward delivering projects over many months or years. It's not set up to deliver new projects very fast and in quick succession. A new model for IT was needed that would enable smaller steps or sprints, rather than marathons, in order to explore and exploit this new corporate digital experience.

Hybrid IT is about exploration, about finding your place in the market, about finding a new, more competitive business model. It's about disruption.

Q: At first blush, hybrid IT could appear to be more about managing big data than a makeover of IT. For example, maybe it incorporates more high-performance computing (HPC). Explain the relationship of big data and HPC with hybrid IT.

Partridge: HPC is the engine for understanding data. HPC and big data are at the heart, but just understanding data isn't enough. New models have to be designed and delivered so that the information and insights can be exploited and monetized.

Materializing what you learned from data requires a new IT operating model. That new model is hybrid IT, and it is uniquely designed to find ideas in the data and then move those ideas into a business reality.

Q: So, hybrid IT completes the big data equation? It completes the process of moving from actionable insights to completing the action?

Partridge: It's about more than using data we have now and making decisions based on that. It's also about strategically reaching out to harvest or create more data in order to explore and exploit possibilities even further.

Q: Then hybrid IT also incorporates the Internet of Things (IoT) and analytics on the edge?

Partridge: The hybrid IT operating model definitely incorporates data from all types of sensors, the Internet of Things and analytics at the edge. But data can be analyzed at the edge for fast outputs and brought back in for more in-depth analysis.


The key here is the agility, flexibility and speed built into the hybrid IT model. It allows you to prototype, test, explore and experiment fast. Traditional IT is typically more of a guardian than an instigator.

Q: Does that mean that efficiency is no longer an IT priority?

Partridge: As everyone in IT knows, change is the No. 1 cause for downtime. Until everything became digital and data driven, IT was managed to ensure no failure. Traditional IT was focused on delivering more for less. That is, being ruthlessly efficient and reliable, always cutting costs, building business cases and ensuring ROI.

But a creative tension results when you're trying to stick to that model and still operate and innovate at the speeds it takes to be competitive today. Priorities in IT must be realigned to fit the demands of business today. That means speed, agility and flexibility become the top priorities.

Resolving creative tensions between old and new priorities can't be fully answered by technology. People within the business must make some very important, even bold, business decisions. If they don't, creative tension becomes a barrier instead of the path to opportunities.



The world has already changed and continues to change rapidly.

Q: Which leads us to the question of who is the run organization in this new operating model?

Partridge: Yes, that's been one of the hottest questions today. Is your run organization — the people keeping everything running — your internal IT, or is it Microsoft, Amazon and other cloud service providers?

That question is important because you need both a run organization and an organization focused on data and the innovation that it brings. That's why many organizations now have a CDO — chief data officer — to keep the digital vision attuned to revenue opportunities and materializing the means to capture those. The CDO creates a new style of IT that is separate from the run organization.

It used to be that the CIO or CTO also served as the CDO, but that's becoming less common.

Q: Isn't moving the run organization outside, placing efficiencies on cloud service vendors, for example, a smart play?

Partridge: It can be. But what happens when data becomes a critical asset? When data is driving your revenue, caution and more control are prudent. This is another reason the hybrid IT operating model is so critical. It gives you more and faster control on how, when and where you use data and materialize the results. As the appetite for risk changes, securing data-driven revenue becomes a rising priority.

Right now, organizations don't know what they don't know, what their R&D dollars will turn up. But there is a data imperative, and so separating the IT exploratory arm from the run organization is usually a good idea. There are many ways to do that internally and externally.

Right now,
organizations
don't know
what they don't
know, what their
R&D dollars
will turn up.



Q: How long does it take to switch to the new hybrid IT operating model?

Partridge: The answer to that depends on a number of factors. Large organizations may take longer simply because their size tends to slow the rate of change. However, that's not always the case. Companies of any size, even the largest of organizations, that have the will and ambition at the board level often make a fast transition to the new model. HPE is even an example of that.

Having a clear vision — we're talking about a CDO here with a clear vision on the new style of IT he or she wishes to create — helps speed the transformation, too. You need a clear goal and a specific strategy from the start.

And a good deal depends on how disruptive that vision is. For example, in car manufacturing, the connected car system is highly disruptive and calls for an accelerated transformation to support that disruption. Industries in such highly disruptive scenarios tend to have companies making rapid progress toward a new IT operating model.

Hybrid IT is the foundation for the future

As organizations begin to delve deeper into data and what Partridge calls “the digital experience,” the demands on tools and IT operating models will tax their limits. It is already obvious that the traditional IT model is insufficient on its own, but as Partridge points out, it still plays a vital role now and will for the foreseeable future. However, in a world that is now and forever data driven, IT must adapt to an exploratory mode that can more efficiently find and exploit new opportunities. IT must find ways to both be and combat the disruptor. A hybrid IT operating model is a strong foundational step in the right direction. ■

Companies that have the will and ambition at the board level often make a fast transition to the new hybrid IT model.

01

**NOW
TO
NEXT**

Lessons for leaders : Making the most of hybrid IT

Look at hybrid IT as a new operating model.

Understand how hybrid IT can help you take better advantage of the data your business collects.

Leaders need to provide the vision and guidance to make a successful transition from traditional IT models.

02

Hybrid IT survival skills

By Steven J. Vaughan-Nichols

-
- How to deploy a cloud
 - Infrastructure-as-a-service
 - Platform-as-a-service
 - Benefits and concerns

Everything you need to know about clouds and hybrid IT

Let's go over the basics again. [The National Institute of Standards and Technology](#) spelled out [cloud computing](#) for us years ago: "Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."

[Clouds must have five essential characteristics](#): On-demand self-service, broad network access, resource pooling, rapid elasticity or expansion, and measured service. On these fundamentals, cloud designers have created a whole new IT approach.

There are three ways to deploy a cloud: public, private and hybrid.

Public clouds, such as [Amazon Web Services \(AWS\)](#), [Google Compute Platform](#) and [Microsoft Azure](#), are controlled by their owners, from whom you rent computing services. In contrast, a private cloud runs on your own servers using cloud software such as [Nextcloud](#), [OpenStack](#) or [VMware's vSphere](#).



As you might guess, a hybrid cloud bridges the gap between public and private using your own special mix of public and private cloud services.

However, as clouds have continued to develop, the spaces between public and private models have shrunk. Case in point: You can use Azure as a public cloud or deploy it in-house as a private cloud and a hybrid cloud. For example, the [HPE | Microsoft Azure Stack solution](#) can connect the Azure public cloud with Azure running on top of your existing servers.

From a business perspective, the closer you are to using public clouds, the more your IT spend shifts to operating expenses (OpEx) versus capital expenses (CapEx). As the top brass sees it, public cloud-based technology promises to replace high CapEx with lower OpEx.

How much of a savings is this? Mark Pietrasanta, chief technology officer at [Aquilent](#), a government IT solutions provider, says that compared with typical data centers, customers can realize “[dramatically reduced costs](#) — often 75 to 90 percent savings.”

Even if you don’t realize that kind of savings, your CFO gets more predictable IT operating expenses. That’s always a good thing.




“Compared with typical data centers, [hybrid IT] customers can realize dramatically reduced costs — often 75 to 90 percent savings.”


Mark Pietrasanta, chief technology officer at Aquilent



■ Hybrid IT data center
■ Typical data center

BUILDING BLOCKS OF CLOUD SERVICE

-  SaaS
-  PaaS
-  IaaS



No matter which cloud model you use, there are three main ways to consume cloud resources: Infrastructure-as-a-service (IaaS), Platform-as-a-service (PaaS) and Software-as-a-service (SaaS).

You've heard of many other kinds of cloud services. For example, there's data as a service, test environment as a service, desktop as a service, and API as a service. Those are all corner cases. For practical purposes, IaaS, PaaS, and SaaS are the only ones you need to worry about.

IaaS

Infrastructure-as-a-service (IaaS) is the building block for all the other cloud services. With IaaS, users are given access to physical or, more commonly, virtual servers. These provide file storage (think Dropbox) and other fundamental computer services, such as firewalls, load balancers, virtual LANS, and domain name servers.

Name a public cloud service and chances are it's IaaS. Besides the aforementioned public clouds, [Rackspace](#), [IBM Cloud](#), [Oracle Cloud](#) and [VMware vCloud Air](#) are also significant players.

The primary advantage of IaaS is scalability. With an IaaS hybrid cloud, you can quickly ramp up additional resources. For example, if your storage needs balloon from the terabytes of a private IaaS

to a public cloud's petabytes, you can master the expansion simply by paying for more room using a hybrid model.

An IaaS offers more than just storage. It can also provide shared storage space for your workgroups.

As Microsoft has pointed out, an IaaS is more than just disks in the cloud: “Applications that require complete control (e.g., disaster recovery services that need to clone the drive by capturing disk I/O at the driver level, software licensing service based on the virtual MAC address) can be built on IaaS. The developers and IT professionals have access to the complete app platform stack, user mode subsystems, and kernel-level control so that the VM can be customized to the needs of the business domains they serve.”

If you use any kind of cloud, underneath it all you'll find an IaaS. It may be invisible to your users, but it's there.


PaaS

Take an IaaS and add a software development stack to it. Now you have a Platform-as-a-service (PaaS). Examples include AWS Elastic Beanstalk, Red Hat OpenShift and HPE Helion Stackato.

With a PaaS, you can take your existing in-house applications and migrate them to a cloud. This makes your programs more elastic. Let's say your user base goes from hundreds to thousands. With a PaaS, your resources can expand to meet client needs without overprovisioning.


A PaaS solution can often require a lot of work from your developer team. For example, if your applications rely on local file systems, expect locally stored data to be persistent between restarts, or rely on a media access control (MAC) address for licensing. All will require serious rework before they can be cloud-deployed.

You should also consider your existing software stack when choosing a PaaS cloud provider. For example, if you use JBoss for your corporate Java EE server, your cloud must support Red Hat OpenShift.



Enterprise adoption of IaaS as the primary environment for workloads will jump from 10% in 2015 to 51% in 2018.

Source: McKinsey, September 2016.



PaaS will be the fastest-growing IT service line with a CAGR of 37% between 2014 and 2020.

Source: Ovum, November 2016.

Additionally, rewriting system diagnostics, accommodating resource governance in a multi-tenant setting and implementing software metering can be troublesome. Finally, once you've moved an application to a PaaS, porting it back to its older model is not easy. In short, look before you leap when moving your program to a PaaS.

On the other hand, after your team has moved the application to the cloud, maintaining it should be easier. Cloud architectures, with their reliance on Representational State Transfer (REST) web services, are relatively simple to maintain.

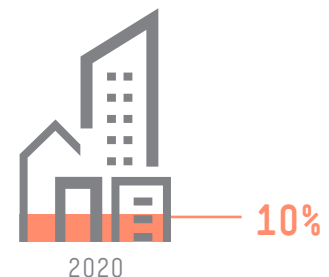

SaaS

Once you've used a PaaS to build pre-packaged applications, you have Software-as-a-Service (SaaS). SaaS started as a variation of client-server computing and then took a left turn into application service providers in the late 1990s and early 2000s.

SaaS became mainstream because of two factors. First, the cloud architecture's elasticity made it possible to deliver applications without manual intervention. And second, the web browser became the universal user interface. So we now have [Google Docs](#), [Office 365](#), [Salesforce CRM](#) and a horde of other applications.

Whether you're using someone else's application or running your own, there are numerous advantages to using SaaS. These include:

- **Fast deployments:** SaaS software solutions can be implemented in a matter of weeks versus months. Adrian McDonald, [EMC's EMEA president](#), claims that, thanks to the cloud, the average time for new application deployments will be reduced by more than 20 percent. Once you've finalized a SaaS program design, you can easily deploy it. How? By using the cloud to simplify the complexities of ordering, configuring, staging and scheduling to a SaaS-based web interface.
- **Universal access:** Thanks to its reliance on the web, if you or your customers have access to the internet, they can work with your programs.



LARGE ORGANIZATIONS USING SAAS PLATFORMS

By 2018, 4% — and by 2020, 10% — of large organizations will run all their CRM applications on SaaS platforms.

Source: Gartner, October 2016.

Cloud pros and cons

Besides the financial benefits of a cloud model, there are three other major reasons to use Everything-as-a-service (XaaS) with your existing servers and a public cloud in a hybrid model:

1. Faster time to market: Instead of hours in manual setup time, it takes less than a minute to set up a new server automatically. In addition, most of the design has already been validated because it's based on reference architectures. All you need to do is customize it to your unique needs.

2. Improved ongoing support: For most as-a-service models, ongoing support and maintenance services such as service desk, tiered support, network monitoring and third-party vendor management can be handled by the cloud provider. In addition, the ability to leverage resources at cloud scale lets you support-as-a-service solutions more cheaply than those of self-managed, one-off programs.

3. Resilience: With a hybrid model, your IT infrastructure is stored not only on your site but also on a public cloud. The meteorite that strikes your data center is unlikely to also hit your cloud provider's data center. This way, even if a disaster strikes your business, you still have access to your data and programs. That means you can always be sure you'll be able to get your company back up and running.



That's the good news. Now, here are some things that can go wrong with a cloud-based IT infrastructure.

1. Security concerns. The more you turn your resources over to the public cloud, the less control you have over them. You are, after all, leasing servers instead of owning them. This can be a real worry. For example, the [Cloud Security Alliance](#) lists no fewer than a dozen serious security concerns in its [Treacherous 12](#) list. Most of these are far more dangerous to public cloud consumers.

That doesn't mean clouds can't be secured. The CIA knows a thing or two about security, and the agency [entrusts its secrets to a private AWS cloud](#). According to Jill Tummler Singer, former deputy CIO at the CIA: “By [keeping the cloud inside your firewalls](#), you can focus your strongest intrusion detection and prevention sensors on your perimeter, thus gaining significant advantage over the most common attack vector: the internet.”

This is another reason why the hybrid model can work for you. By keeping your company's crown jewels inside a private cloud and leaving day-to-day materials on a public cloud, you can balance security and costs saved.

“By keeping the cloud inside your firewalls, you can focus your strongest intrusion detection and prevention sensors on your perimeter, thus gaining significant advantage over the most common attack vector: the internet.”

Jill Tummler Singer, CIO at the U.S. Central Intelligence Agency

2. Compliance. Security is hard. Compliance with security regulations can be harder. Ask healthcare providers about toeing the line for the Health Insurance Portability and Accountability Act (HIPAA). Ask financial service firms about jumping through Gramm-Leach-Bliley hoops. Ask accountants about their Sarbanes-Oxley nightmares. Oh, and don't forget about financial companies that must comply microsecond by microsecond with Payment Card Industry Data Security Standard (PCI-DSS) regulations. It's not easy.

Making matters even more complex, public cloud providers aren't governed by the same rules that apply to your vertical company. According to Jim Whalen, senior analyst at [The Taneja Group](#): "HIPAA requires that organizations providing business services to regulated customers must qualify as a 'business associate.' Cloud providers aren't subject to the same level of compliance and reporting that apply to healthcare businesses." In short, the buck stops with you, not your cloud provider.

There are public clouds that are authorized to handle authorized data types. For example, [Cerner](#) can handle electronic health records. But your run-of-the-mill public cloud can't be trusted with this data. You'll need to keep it within the confines of your private cloud.

3. Performance. A cloud-based application is only as good as your users' internet connectivity. According to the latest [Akamai Internet](#) report, 80% of U.S. internet users now have an average connection speed faster than 4 Mbps, but fewer than half (46%) have 10 Mbps. If your application is graphics heavy or requires video, you're going to have trouble keeping your customers happy.


This balancing act of cost, flexibility, control, risk and security is pushing many organizations to adopt hybrid IT strategies.

This isn't just the hybrid cloud. True, a properly implemented hybrid cloud enables you to benefit from both private and public cloud models. For example, you can use your private cloud with its faster internal network and increased privacy for mission-critical, private data. Concurrently, you can use the public cloud network for more innocuous tasks such as sharing resources, backing up private data and hosting low-security data.

But you don't need a cloud for everything. Lydia Leong, distinguished analyst at [Gartner](#), says, "[Not everything can or should be cloud](#). Customers have plenty of applications that they probably will never move to the cloud — or at least will not move to the cloud in any kind of reasonable time frame."

That's where hybrid IT comes in. As HPE puts it: "This is about having the right infrastructure optimized for each of your applications, old or new. But it isn't just in your data center. It isn't just in the cloud. Your infrastructure has to be everywhere, at the right cost, at the right performance, with the right management, at the right scale. And it all has to work together seamlessly. A hybrid infrastructure is one that seamlessly combines public cloud, private cloud and traditional IT."

Average connection speed for
U.S. internet users



Your infrastructure has to be everywhere, at the right cost, at the right performance, with the right management, at the right scale. And it all has to work together seamlessly.

Leong adds, “Lots of people are still running mainframe systems, for instance, that are unlikely to end up going to the cloud, at least until the applications themselves are replaced. For those systems where there is no pressing need to move them onto a cloud infrastructure, it can be better to avoid the expense of doing so until necessary.”

The two models can, of course, work together. While your application may still live in a mainframe, there’s no reason it can’t be presented to users on the web via SaaS. As ever, interoperability and integration will be an important part of central IT’s job.

Managing the proper balance of cloud and traditional IT isn’t easy. But with an organized hybrid IT plan, you can move your company into the future using both 21st century cloud and your existing technology with its 20th century roots. ■

02

**NOW
TO
NEXT****Lessons for leaders: Hybrid IT survival skills**

Understanding what the different clouds and cloud services can bring to the table means first realizing your own infrastructure and business needs.

Clouds are not a universal solution; know the potential problems as well as advantages.

Proper planning is mandatory for the successful implementation of cloud and hybrid IT solutions.

03

How to evolve your infrastructure to hybrid IT

By Ken Hess

-
- Things must change
 - Backups and disaster recovery
 - Tools and security
 - Convergence
 - Why hybrid IT?

Evolution in action

Data centers have always been cold, noisy, poorly lit and filled with rows of racked server systems and storage arrays. While the technology inside those racks has evolved over the past 50 years, the purpose has never changed: to provide a single, climate-controlled, secure space to house computing power for customer access. That's a data center's life in a nutshell.

Data centers have also always been available as multi-tenant properties that provide space, computing power and environmental protection to a range of users and customers. Traditionally, companies either built and maintained their own data centers or rented space from a colocation data center provider, where they placed their own servers, storage and network equipment. They paid for network bandwidth, a portion of the power and cooling, and maintenance. Some companies also maintained their own systems by acquiring physical data center access.

So if nothing has really changed in the past 50 years, why all the buzz about hyperconvergence and the evolution toward hybrid IT? The answer is simple: Servers no longer serve single applications to customers. Since the introduction of virtualization in the 1990s, companies have been converting



physical systems to virtual ones by the thousands. Virtualization's popularity has grown to the point where owning the underlying hardware might not make sense for every workload.


Companies have also begun leasing virtual private servers (VPSs) that perform and behave the same as their familiar physical counterparts. Cloud computing and as-a-service offerings have further distilled computing down to workloads and applications. In many cases, companies no longer need to deploy servers, provision storage, connect network resources and provide environmental protection. Instead, they focus on deploying applications, creating portals and building customer interfaces.

The legacy data center model made up of disparate hardware components is evolving into a software-defined, converged, hyperconverged and virtualized version of its former self. The data center's functions haven't changed, but its components continue to evolve.

Increasingly, chief technology officers realize that fully localized or on-prem infrastructures make their organizations less agile in an ever-changing marketplace. Yet they face a range of issues, from existing data center investment to regulations that force them to keep direct control of data. This is where hybrid IT becomes a clear path to the future. Moving to a hybrid cloud infrastructure is a stepwise process. It generally follows an evolutionary path that begins with moving less-critical services and infrastructure to the cloud.

Public clouds, such as Amazon Web Services (AWS), Google Compute Platform and Microsoft Azure, are controlled by their owners, from whom you rent computing services. In contrast, a private cloud runs on your own servers using cloud software such as Nextcloud, OpenStack or VMware's vSphere.

As you might guess, a hybrid cloud bridges the gap between public and private using your own special mix of public and private cloud services.



In three years, rack-level hyperconverged and hyperscale bundles will account for 30% of server/storage/network deployments, driving changes in power and cooling design.

Source: IDC, November 2016.

Backups and disaster recovery

The first step toward hybrid IT often addresses backups and disaster recovery. For many enterprises, this step is an easy transition that probably won't impact regular business operations should something go wrong.

Moving to a hybrid IT environment used to mean that you had a cloud-based disaster recovery (DR) site, existing "just in case" your primary site experienced a complete outage. For backups and disaster recovery, the cloud has a distinct advantage: mission-critical data and files are stored off-site.

Cloud providers use multiple geographically remote locations to store your data for redundancy and safety, while also providing high-bandwidth access to that data.

Cloud solutions are cost-effective. You're not paying for leased systems, rack space, environmental security, power, cooling or physical security on an individual system basis; therefore, providers can charge much less for services on a per-use basis. You can also save money by eliminating traditional tape, disk and transportation costs associated with managed backups.

The cloud is both vertically and horizontally scalable. Since cloud providers have multiple data centers and thousands of systems, you can easily and almost infinitely increase your capacity without disruption. Cloud scalability means that you have the option to scale up or to scale out.


The most significant issue with using the cloud for disaster recovery is initially synchronizing your production data with your DR data. Synchronization requires a significant amount of bandwidth for initial seeding, which is why most businesses perform a complete backup on-site and then transport that backup to the cloud provider for restore. Incremental synchronization and backup require less bandwidth, but demands vary greatly between businesses and among business cycles.

Moving to a hybrid IT environment used to mean that you had a cloud-based disaster recovery (DR) site, existing "just in case" your primary site experienced a complete outage. For backups and DR, the cloud has a distinct advantage: mission-critical data and files are stored off-site.

The stepwise path to hybrid IT

The next step in a hybrid IT transition typically involves development and test environments (DevTest). Companies typically select these environments because of their less-than-critical nature. The business doesn't necessarily depend on 100% uptime of development and test systems. The cloud is the perfect match for DevTest not only because of the non-critical nature of those systems, but also because administrators and developers can create and destroy systems as needed.

Using cloud-based systems, developers, testers and administrators can use copies of production data without risk to production or to customer access. Cloud providers also enable administrators to manage their systems centrally using a web browser and without the need for virtual private network (VPN) access or a special out-of-band network. Having this capability in-house — to easily spin up and spin down test systems at will for DevOps personnel — might also decrease payouts to third-party services such as Amazon Web Services (AWS) and Microsoft's Azure platform.



Tools to ease the transition

Many hardware and software vendors offer management tools and singular interfaces to make the transition to a hybrid IT environment easier. This means companies can transition to a hybrid IT business model without changing their management tools or their infrastructure choices. Administrators can learn a single set of tools, which apply to on-premises and to cloud resources alike. Vendors also offer consulting services to help make the transition to hybrid IT easier.

Hewlett Packard Enterprise (HPE) OneView, for example, makes your transition easier by allowing administrators to manage physical, virtual and cloud environments in the same management console application. Using OneView management plugins, administrators can also monitor compute, storage

and network components using centralized dashboards. Administrators gain the flexibility to provision, manage and control virtual resources, such as expanding storage volumes to add more capacity to a computing environment.

Dell's OpenManage Essentials (OME) Version 2.0 is the latest incarnation of a console solution that monitors Dell and third-party hardware. It also provides full lifecycle server management. OME is a free download from Dell.com that easily installs on a standard Microsoft Windows server, but licensing is fee-based. Using OME, administrators can deploy operating systems to bare-metal servers, establish baseline configurations, verify and report on configuration compliance, track system health and manage configuration drift.

And when you're ready to move further into hybrid cloud infrastructure, HPE's Hyper Converged Systems, for example, deliver preloaded Microsoft Cloud System Standard software that includes Azure-connected services such as Azure Backup, Site Recovery, Operation Management Suite and Windows Azure Pack. Or you can opt for a VMware vSphere starter cloud solution with HPE's Helion CloudSystem offering.

HPE's Simplivity provides the industry's most complete hyperconverged platform. It includes the OmniCube converged storage solution and the OmniStack Data Virtualization Platform, which is now in its third generation.

Dell's flagship hyperconverged offering is the Dell EMC VxRail appliance, which provides configuration flexibility, seamless integration into VMware environments, streamlined deployment and lifecycle management. Dell supports a range of appliances.

The Nutanix Enterprise Cloud Platform (ECP) allows businesses to start small and scale performance and capacity as needed. The Nutanix storage management solution, Prism, helps you limit or eliminate storage overprovisioning; Nutanix ECP integrates with existing Hyper-V and VMware environments.

Many hardware and software vendors offer management tools and singular interfaces to make the transition to a hybrid IT environment easier.

Security

There's a lot of concern surrounding cloud security, but the truth is that businesses should handle cloud security like all other security — using least privilege, strong passwords and two-factor authentication. The security issues that have plagued cloud solutions have less to do with the provider's security and more to do with the customer's assumptions that all security responsibility falls upon the provider, which is not the case. Cloud customers must ensure that virtual private servers (VPS), all customer-managed infrastructure systems and all customer-developed applications adhere to strict security guidelines. These systems generally fall outside of the provider's realm of security responsibility.

Check your provider's service level agreement (SLA) and services agreement to find out the specific boundaries of its responsibilities and yours. For example, VPS patching, software installation and application maintenance typically are your responsibility, while your provider guarantees 100% uptime for those systems and services. It cannot guarantee security for applications that you develop or for the information those applications access.

Converged, hyperconverged and fabric-based infrastructure

Converged infrastructure consolidates the four essential and formerly disparate components of data center computing functions into a single chassis or housing: compute, networking, storage and virtualization. The emergence of converged infrastructure has been driven by the desire to minimize complexity. Formerly, provisioning network, storage and compute were chores that could take weeks for a single application. Converged infrastructure allows administrators to deploy applications and workloads far more rapidly. Depending on a company's governance procedures, the entire process can take as little as a few days. Other significant benefits of converged infrastructure are centralized management of the entire infrastructure and the fact that the components all work together as a single unit.

Hyperconverged infrastructure has further evolved from the requirements to better leverage administrative personnel and to add better support ancillary services such as backup, snapshotting, data deduplication and WAN optimization. Hyperconverged architectures rely heavily on a software-defined approach to converged computing, which means software and hardware are tightly integrated and cannot be separated from one another. Look at storage, for example. In standard infrastructure architecture, storage and compute are usually separate hardware components connected via cabling.

Four data center computing functions:

1. Compute
2. Networking
3. Storage
4. Virtualization

Hyperconvergence places storage in direct contact with the compute layer via a communications bus. This directly attached storage configuration makes it easier to create a single storage pool that every server in the hyperconverged appliance can access.

Dell EMC claims that its VxRail appliance line-up is the “only fully integrated, preconfigured, and pretested VMware hyperconverged infrastructure appliance family on the market.” Since Dell EMC is the majority shareholder in VMware, it makes sense that its tight integration with VMware’s Virtual SAN and vSphere software makes a good case for its all-in-one hyperconverged solution. The Dell EMC appliance features a single point of contact for hardware and software, plus its call-home and proactive two-way remote connection for remote monitoring, diagnosis and repair.

Another approach is the HPE Hyper Converged 380, an example of HPE’s computer architectural hyperconvergence. An IT generalist can manage the infrastructure from a mobile device app, deploy virtual machines in minutes, and add storage and compute power to workloads in minutes. The HPE Hyper Converged 380 is an all-in-one compute, software-defined storage and intelligent virtualization solution built on the world’s best-selling server, the HPE ProLiant DL380. Administrators manage, monitor and control this system and its resources through HPE OneView. OneView’s automation engine uses a template-based approach to infrastructure deployment that greatly speeds up the configuration of server, fabric and storage.

Though the terminology used by different companies may differ, the underlying goals remain the same. Gartner and others use the terms “fabric-based computing” and “fabric-based infrastructure” to describe converged infrastructure and hyperconverged infrastructure, respectively.

Hyperconvergence places storage in direct contact with the compute layer.

Why hybrid IT?

The big question is, “Why would my company want to move to a hybrid infrastructure?” After all, if you’ve done business the same way for the past 10, 20, 30 or more years, you might ask why you should change now. Chances are very good that you have changed the way you do business over the life of your business. You have to change with the times.

On-prem infrastructure is expensive, requires a lot of maintenance and continually requires upgrades. A hybrid environment can cut many of these standard costs to zero. For example, when using a cloud service, you’re not responsible for maintaining the underlying hardware or refreshing the hardware every few years.

The promise of lower cost doesn’t stop with maintenance; it extends to personnel as well. Having fewer tasks require fewer people to perform them. Also, a consolidated set of tools for management is a big savings. Before 2015, consolidated tools were practically nonexistent, and administrators had to learn one or more sets of new tools to manage and to deal with cloud services. Having to invest a great deal of time and money into learning new tools made the cloud look like a very poor investment. That has changed. Vendors now offer single tool sets with which to manage on-prem and cloud infrastructure and services.

INNOVATION



Through 2019, **every \$1** that enterprises invest in innovation will require **an additional \$7** in modernization.

Source: Gartner, November 2016.

MODERNIZATION



Rapid service provisioning is a hybrid IT attribute that benefits IT administrators as well as the business. Marketing teams don't have to plan for the IT requirements of campaigns months in advance. Businesses can respond quicker to market changes and tap new revenue streams as they emerge, rather than months after they've peaked.

Once internal governance requirements have been met, IT administrators can provision an entire compute infrastructure within minutes. Once the campaign or other reason for the spin-up has subsided, IT staff can remove the additional services and discontinue any added charges. That scenario is impossible with traditional infrastructure because of purchase lead time, provisioning, governance and deployment time requirements.

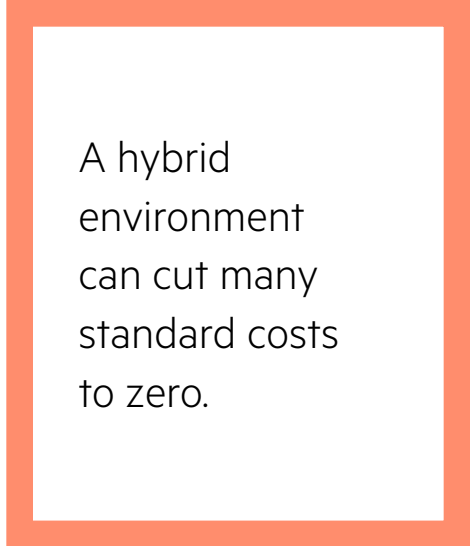
However, simplicity of deployment doesn't imply lax governance rules. In fact, the opposite might be the case, particularly when dealing with certain types of regulatory compliance issues.

Additionally, once you've made the investment in new hardware to handle the demands of a marketing campaign or a new line of business, that new software licensing and additional resources are yours to keep whether they're used or not.

If your IT staff runs low on on-prem capacity, they can augment by adding cloud resources. When you're out of floor space for new racks, that's a problem, but one that you can remedy. But if you're out of electrical capacity, that's a problem that you can't necessarily resolve if your data center has no more capacity to use. A hybrid cloud solves both problems by allowing you to spin up when you need capacity and spin down when you don't. And possibly no floor space is required.

A hybrid IT environment is also regulatory compliant. Workloads that must stay on-site can run on your internal cloud while you conserve internal capacity by running other workloads in the public cloud. Hybrid IT solutions leverage open-source technology to increase flexibility and avoid vendor lock-in, which is good news for businesses that want to remain somewhat vendor-agnostic.

Caveat: Although many industry pundits laud vendor agnosticism, you need to select contracted services carefully and ensure that their personnel understand and adhere to regulatory compliance requirements. A vendor-agnostic solution doesn't guarantee compliance with regulatory requirements such as HIPAA and PCI. On the whole, mainstream vendors adhere to strict security and regulatory compliance requirements both for hardware and software.



A hybrid environment can cut many standard costs to zero.

Conclusion

Hybrid IT offers resiliency, reliability, agility and frugality. It offers control while minimizing the hassles of maintenance and refresh. It also provides your IT staff a familiar environment to work in, by sharing the same tools as your on-prem, hyperconverged infrastructure. You can enhance your business continuity plans via cloud-based backup, disaster recovery, snapshots and remote copy integration. Hybrid IT also increases business agility by allowing you to expand and contract infrastructure as needed.

But a hybrid IT or hybrid cloud infrastructure isn't just cheaper, more agile or more resilient than an all on-prem infrastructure. It's also easier to manage, which takes the stress out of business continuity by providing a management interface that extends and transcends workload location. The hybrid IT model enables your IT team to become an internal service provider that ensures optimal performance, domain security and regulatory compliance. ■



03

**NOW
TO
NEXT****Lessons for leaders: Evolution in action**

The path to hybrid IT is dictated by business needs.

Hybrid IT deployment requires buy-in from both business and IT.

Finding the right balance between CapEx and OpEx can define your hybrid IT strategy.

How hybrid IT enables the software-defined data center

By Alyson Behr

- Understanding the virtual world
- Containerization
- Software-defined everything
- Orchestration and automation
- Understanding the path
- Take your time and define your goals

Hybrid IT in the software-defined data center

“When you come to a fork in the road, take it,” said Yogi Berra. So it is with the path to a virtualized hybrid IT environment.

Not so long ago, static architecture ruled using client-server computing. Along came mobile devices and their content, server virtualization, big data and the cloud — all of which changed enterprise data center traffic patterns. The evolution of today’s data center networking demands requires massive compute power to handle not only the increased classic server to client north-south traffic, but also the new east-west, machine-to-machine communication that occurs within the data center and public and private clouds.

Enterprises are embracing new software-defined concepts and complementary technologies that have been developed to take full advantage of the promise a virtualized software-defined data center (SDDC) holds. It’s a complex path to achieving a fully virtualized hybrid IT environment, but worth taking given the performance, agility and cost benefits. Here are high-level descriptions of the major concepts and technologies you’ll need to understand as you move forward.

Computer virtualization

Also known as software-defined compute, computer virtualization is the fundamental technology SDDC is based on. It requires x86 servers and has become the industry-standard technology that most data centers use. It’s important because in the old world, traditionally deployed servers used only around 10 to 15% of their capacity. Virtualization decouples, or separates, memory and CPU from physical hardware, allowing those unused resources to be used wherever they’re required. Decoupling creates an individual container, or a virtual machine (VM), for applications and its operating systems that can live and run at the same time on an x86 physical server. The result is that most of that server’s capacity is put to use. Computer virtualization increases your server efficiency, delivering higher performance and higher availability, all at a lower cost.

By 2019,

30%

of the global storage array capacity installed in enterprise data centers will be deployed with SDS or hyperconverged integrated system architectures based on x86 hardware systems.

Source: Gartner, April 2016.

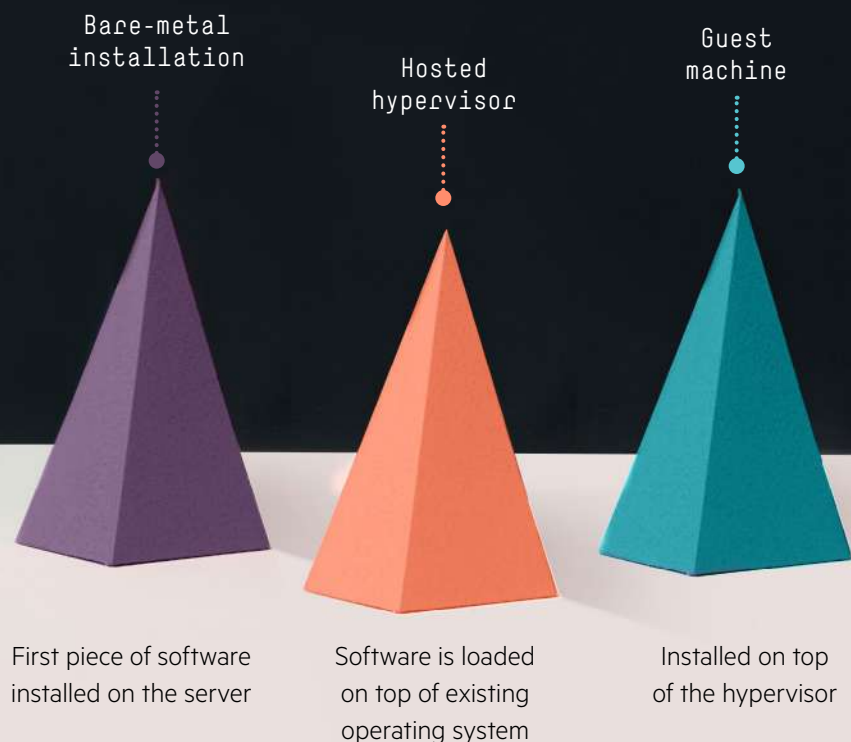
Hypervisors

A fundamental piece of computer virtualization technology is the hypervisor. This piece of software lets physical devices share their resources among VMs running as guests on that physical hardware, or host machine. According to Wikipedia, “A hypervisor or virtual machine monitor (VMM) is a piece of computer software, firmware or hardware that creates and runs VMs. A computer on which a hypervisor runs one or more VMs is called a host machine, and each VM is called a guest machine.”

There are a few different types of hypervisor. One type that is often used on production systems is implemented as a bare-metal installation. It’s the first piece of software to be installed on the server as the operating system, and it also becomes the hypervisor. It communicates directly with the underlying physical server hardware whose resources are subsequently virtualized and delivered to the running VMs. Another type of hypervisor is a hosted hypervisor. In this scenario, the software is loaded on top of an already existing operating system. Resources need to take an extra hop to pass through the VM — and even then, latency is minimal.

A third option is called a guest machine, also known as a virtual machine. In this case, it’s the workload that’s installed on top of the hypervisor. It may be a virtual appliance, an operating system or a virtualization-ready workload. It behaves unilaterally, as if it is its own system with dedicated resources. The virtualization technology lets multiple VMs run on top of that physical host, while resources are shared with other VMs.

TYPES OF HYPERVISOR



Containers

Containers and VMs are often thought of as being the same thing. In fact, they are similar but have crucial pluses and minuses. In an ITworld article, [Steven J. Vaughan-Nichols says](#), “The whole point of a container is to run a single application. The more functionality you stick into a container, the more likely it is you should’ve been using a virtual machine in the first place.” He goes on to explain that “VMs take up a lot of system resources. Each VM runs not just a full copy of an operating system, but a virtual copy of all the hardware that the operating system needs to run. This quickly adds up to a lot of RAM and CPU cycles. [In contrast, all that a container requires](#) is enough of an operating system, supporting programs and libraries, and system resources to run a specific program.”

The bottom line? “Generally speaking,” he says, “you want to use containers to run a single application and VMs to run multiple applications.”

“VMs take up a lot of system resources. Each VM runs not just a full copy of an operating system, but a virtual copy of all the hardware that the operating system needs to run.”

Steven J. Vaughan-Nichols

Software-defined storage (SDS)

SDS is storage deployed as software. Your applications and underlying storage services share hardware resources. SDS is the second step, behind software-defined compute, to achieving an SDDC infrastructure. There are a couple of ways to leverage SDS: You can use either a cost-optimizing approach with virtual storage appliances, which are purely software colocated with your applications, or employ a service-level optimization approach that uses a dedicated, virtualized multi-tenant system that works well for large enterprise traffic requirements. SDS is highly scalable and housed on industry-standard servers, and allows you to do away with dedicated arrays.

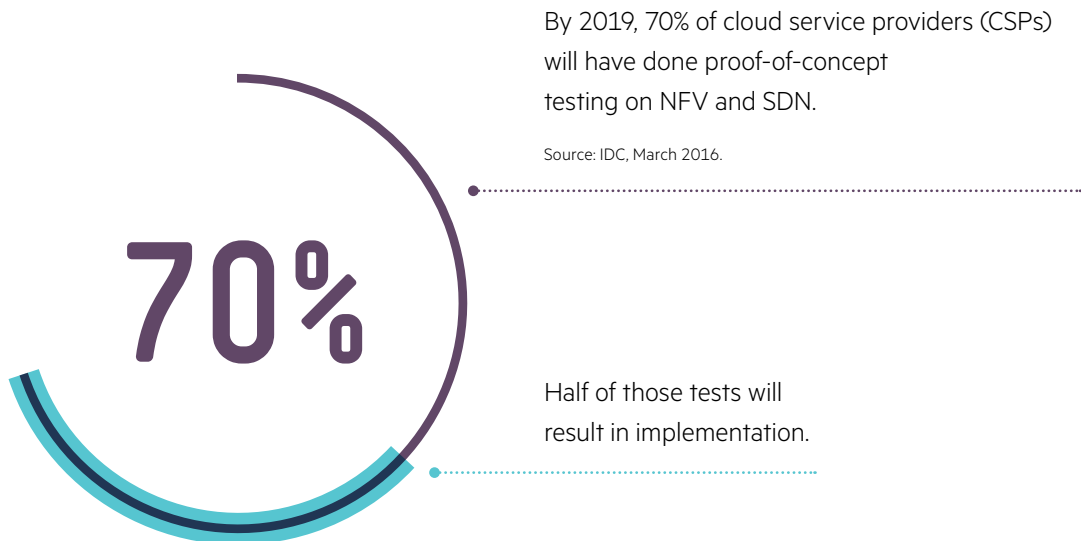
Software-defined networking (SDN)

Also a must-have piece of the SDDC puzzle, SDN addresses the needs of modern compute environments. SDN lets network administrators manage network services by abstracting lower-level functionality. It accomplishes this by decoupling the control plane, which is the system that decides

where traffic is sent, from the data plane, which is the underlying system that forwards traffic to the selected destination. The network control becomes directly programmable, and the underlying infrastructure can be abstracted for applications and services. Other benefits include increased network agility and the ability to centrally manage it. If it is implemented through open standards, it becomes a vendor-neutral, standards-based proposition. SD-WAN applies this technology to wide area networks.

Network functions virtualization (NFV)

NFV uses virtualization technology to decouple network functions and create individual virtualized network functions (VNFs) that IT can connect as desired. Think of individual appliances like firewalls, load balancers or WAN accelerators being virtualized, and thus no longer needing the unwieldy hardware. A VNF consists of VMs, or containers, running different software and processes on top of standard infrastructure, rather than having hardware appliances for each network function. NFV is similar to SDN but different. It is not dependent on SDN, so you could potentially employ a VNF on existing networks without SDN. That said, its performance and cost benefits in an SDDC are clear, and several vendors are developing NFV/SDN platforms.




The road to Rome

As the saying goes, “All roads lead to Rome.” That couldn’t be more true when describing how to get your enterprise to a harmonious place in a virtualized hybrid IT environment. Margaret Dawson, senior director of global product marketing at Red Hat, says, “Hybrid IT is not uniform. There’s so much you can do to put you on that path and help you better integrate this hybrid and very heterogeneous environment. It becomes complex, not only hybrid in terms of structure — meaning physical infrastructure, virtual, private cloud, public cloud — but within those footprints of a hybrid IT environment, there are very different technologies that may or may not work together.”

Dawson continues, “There are steps you can take to modernize by literally looking at each layer of the stack. Take storage. Can you move to more of a distributed software-defined storage environment instead of a traditional storage model? Look at things that can be used on your traditional infrastructure and also work better in a virtualized infrastructure as you begin to move toward the cloud.”

Atchison Frazer, a seasoned chief marketing officer and strategic marketing executive for several well-regarded IT infrastructure management and security companies, describes the issues and current landscapes: “What I’m seeing is an overlay approach where you take the concept of containerization and overlay it with all the other legacy components, including the applications. In a sense, you’re creating a containerized fabric so that all the security and network policies, and even dynamic changes, are done through this layer. Otherwise, you get into the whole cross-silo equation that’s way too complex. Right now, enterprises are running 85 to 90 percent legacy apps.”

There is no single path to a fully integrated hybrid IT environment, but there are steps you can take to inform a good decision. Determining which path to take depends on understanding your specific network and the line-of-business applications you are running on it.



Use a map

Out of the gate, the first task IT has in planning its transition is to create a network map of all infrastructure hardware, each asset's resource expenditure, and the apps running on it. Pay special attention to business-critical financial and enterprise resource planning (ERP) apps. Take into consideration all the mobile devices, and whether they belong to the company or an employee, and catalog all of those apps too.

This puts you in a position to be able to remediate factors like bad architecture, cloud sprawl and third-party applications that were deployed in the same virtualization layer that are degrading strategic business applications you want to move to the cloud. It will also make it clear if you've over-provisioned your hardware, which can help you avoid performance degradation of business-critical apps or having apps you're not even aware of running in the data center.

"Know which apps, or VM workloads, you're going to migrate, and [make sure] they're healthy," says Frazer. "Get rid of end-of-life or legacy apps not being used."

Beware cloud sprawl

Many IT organizations are not aware of the extent to which other business units have embraced the ease and simplicity of infrastructure-as-a -service (IaaS). IT initially thought IaaS was an easy way to let small workgroups have more compute power when they needed it. But that thinking has come around to bite them because these instances have exploded, draining budgets and compute resources.

"Know which apps, or VM workloads, you're going to migrate, and make sure they're healthy. Get rid of end-of-life or legacy apps not being used."

Atchison Frazer, CMO for IT infrastructure management and security companies

What goes first?

So which applications do you migrate over first? There are two schools of thought. Either migrate the workloads where you'll get the biggest bang for your buck first, or experiment with non-strategic apps until you and your business partners' comfort level is high enough to move the more business-critical apps.

You're not going to migrate wholesale. There are plenty of businesses that have applications still running on a mainframe. In this case, there may be apps you've invested so much in that you decide to build a new website front end and keep them on the mainframe. In other cases, it may make more sense to move those apps to more modern infrastructure that includes containerization and virtualization. "It's a workload-by-workload decision-making process focusing on what's best for the business," says Dawson. "It's not just a technology decision about what's new and sexy. It's about which workloads will have the biggest impact on the business."



Be choosy

Armed with your detailed map of hardware, software and technology requirements, you can begin to take a hard look at vendor solutions. Examine your hardware from both a software and operating system perspective. Choose any new physical infrastructure that will allow you to have that consistent foundation across all of those footprints.

"The more you can achieve unification, the easier it will be to scale and get the cost benefit, and also have common security and compliance policies," Dawson points out. "You will have fewer interoperability issues, and your transition will be more seamless."

When a vendor makes your short list, ask them to perform a proof-of-concept exercise prior to you buying in. You may find that no single vendor solves all your issues, and you'll need to look at open source options.

The open road

A main challenge to enterprises moving into an SDN, SDDC environment is interoperability between legacy equipment that is not at end of life and new infrastructure. Frazer notes, “The issue in software-defined has always been, ‘Whose standard?’ There are several companies marketing SDN controllers, each with its own standard. Now there’s an open-source movement. Using its methodology, you avoid lots of issues, including vendor lock-in.”

Dawson agrees: “Having an open API infrastructure is critical. Getting away from proprietary SDKs makes it a more open environment where the pieces can integrate more easily.”

The [OpenDaylight](#) community — a group of developers, service providers and end users — is collectively committed to providing programmable, interoperable networks and tools based on OpenDaylight (ODL), a widely implemented open-source SDN controller. The ODL platform adheres to open standards and uses open APIs. ODL users can choose the features, applications, protocols and plug-ins they need from several vendors’ product lines and then provide connectivity for service providers and customers. You can find out more about the ODL platform [here](#). It’s worth researching.

The end of the line?

The SDDC is not simply a stop on the way to a fully cloud-deployed infrastructure. We’ll be in this hybrid world with the physical, virtual, private and public cloud footprints for a long time. It’s not an either/or situation. Rather, SDDC is absolutely on the path to cloud, and cloud is a key part of the future of the data center. ■

“Having an open API infrastructure is critical. Getting away from proprietary SDKs makes it a more open environment where the pieces can integrate more easily.”

Margaret Dawson, senior director of global product marketing at Red Hat

04

**NOW
TO
NEXT****Lessons for leaders: How hybrid IT enables the software-defined data center**

Implementing an SDDC requires understanding not only the services provided by your existing data center but also your future plans.

Committing to a software-defined infrastructure requires buy-in both from IT and business units, and should be handled with a firm plan laid out.

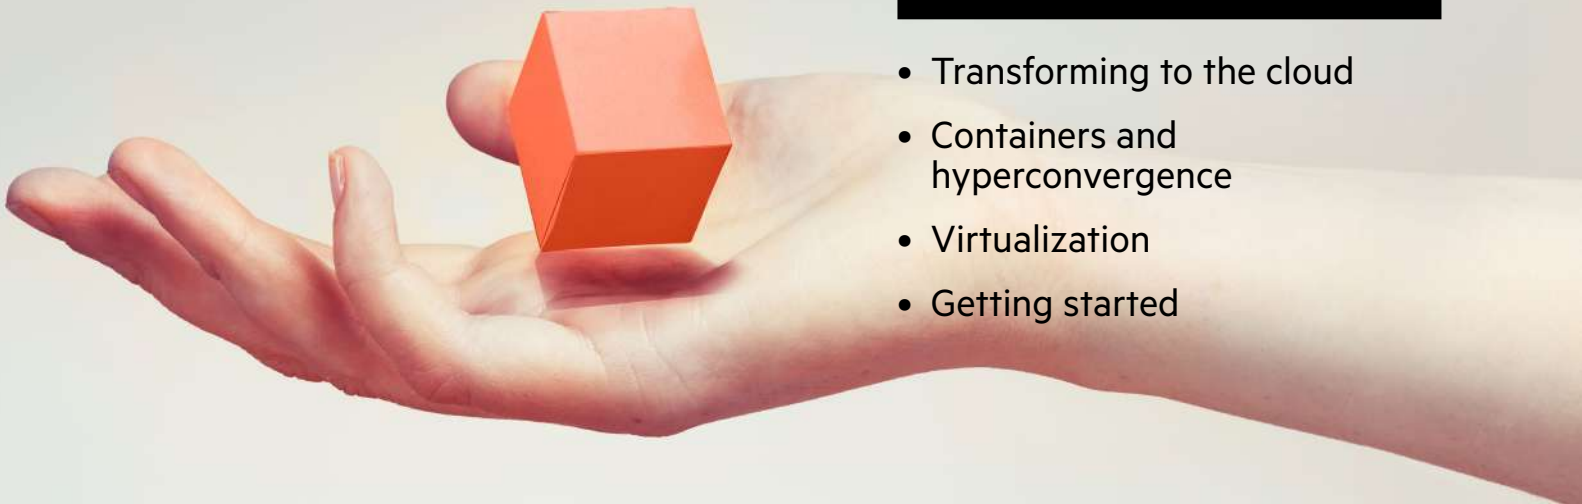
Hybrid IT lets business maximize the benefits of the SDDC.



05

Packaging your environment for hybrid IT

By Frank Ohlhorst



-
- Transforming to the cloud
 - Containers and hyperconvergence
 - Virtualization
 - Getting started

Using containerized IT to deliver a hybrid environment

In its 2015 FutureScape Report, IDC predicted that more than 80 percent of enterprise IT organizations would commit to hybrid cloud architectures by 2017, vastly changing enterprise IT operations. Whether or not that prediction holds true, enterprise IT organizations will still face the uncertainties of technological change predicated on cost, efficiency and, most important, digital transformation.

Transformation to cloud processes

The path to hybrid cloud solutions is full of technologies that claim to power digital transformation and hyperconvergence, while bringing economy and value to existing enterprise IT offerings. In many cases, technologies such as virtualization and software-defined infrastructure have become impediments to achieving forward motion on the path to digital transformation, due to their complexity along with deployment and management overhead.

Yet, all is not lost. Container technology is rapidly becoming the rising star in a field of solutions designed to drive digital transformation, hybrid environment adoption and hyperconvergence.

In many cases, technologies such as virtualization and software-defined infrastructure have become impediments to achieving forward motion on the path to digital transformation.

Containers: A brief history

The idea behind containerization arrived back in 1979 with Unix chroot, an operating system call that changed the root directory of a process and its subprocesses to a new location in the file system. Because this location is visible only to that process (and its subprocesses), it effectively isolates those processes.

From chroot, containerization evolved over the decades into modern containers, which offer the ability to decouple applications from an operating system. In other words, containers provide a lightweight, virtualized environment that isolates an application and its operating system dependencies into a portable, manageable runtime environment.

>50%

of new workloads
will be deployed into
containers in at least
one stage of the
application lifecycle
by 2018.

Source: Gartner, March 2016.

Containers and hyperconvergence: A marriage made in the cloud

Containers are now positioned to make a massive impact on enterprise IT. Container adoption is on an upswing and poised to challenge the stranglehold that virtualization solutions (virtual machines) currently have on the web application market. Containers and virtual machines (VMs) both abstract processes from hardware dependencies. The difference is that containers also abstract the process from its underlying operating system.

That means containers can be smaller and faster, and offer better isolation than a typical VM. Virtualization is designed to run workloads in operational environments that are isolated from the underlying hardware by an abstraction layer. That layer of abstraction allows physical servers to host multiple VMs, each of which can run an operating system, along with its associated applications.

By contrast, containers abstract the application from the operating system kernel. They isolate processes from the operating system by including only the operating system/application dependencies in the container, but still execute processes using the kernel. In other words, the kernel is still responsible for execution, but an abstraction layer isolates the process workloads from the kernel, protecting the kernel from poorly behaved code.

That's one of the most striking differences between a VM and a container. With a VM, a misbehaving application can corrupt the entire virtual machine, impacting any other processes that run on that same VM. Containerized applications impact only themselves when they misbehave, allowing the underlying kernel to continue running other containers.

What's more, launching a VM can be a resource-intensive process, one that introduces severe latency while the VM launches and loads its OS, applications and connectivity services.


Containers spin up much quicker than VMs, simply because all of the prerequisite elements, such as the physical hardware, OS, connectivity and so forth are already loaded and running.

Application isolation and quick spin-up make containers enormously valuable for enterprises that use cloud services, especially hybrid solutions that bridge cloud and on-premises solutions.

Containers offer the following benefits:

- **Small size.** A container includes only the application files and its dependencies, meaning it is many factors smaller than a VM.
- **Non-persistent.** Containers encapsulate everything needed to run the application. However, data and other elements can be stored elsewhere, such as in a database or some other form of persistent storage.
- **Portability.** Containers can be easily moved from one host to another just by shifting the container files.
- **Upgradability.** When an application needs to be updated, a newer version of the container can be quickly deployed.
- **Scale.** Containers can be duplicated and redeployed to address issues of scale.

These benefits make containers a good fit for hybrid cloud deployments, where services or applications may need to be deployed both on-prem and in the cloud, while offering the portability to freely move from one environment to another.



Getting started with containers

Containers don't come into being with a wish and a prayer. One of the first prerequisites is to deploy a cloud service that provides the necessary elements to launch a container orchestration system. Numerous hosts offer preconfigured cloud environments that can readily support containers.


Containers come in many flavors, with vendors such as Docker, Microsoft, Red Hat and many others offering their own versions of container deployment, orchestration and management. Choosing the right container environment starts with establishing what support your cloud service provider offers. If your provider doesn't support a particular flavor of container, you'll have a hard time moving the containers you create from an on-prem host to a cloud service.

You also need to understand what OS is required to run the containerized applications. Some containers can be hosted only on particular versions of Linux, such as CentOS or RHEL, meaning those operating systems must be supported by both your internal IT and external cloud hosts.

It's also important to ensure that your applications are compatible with the container environments. For example, Microsoft Windows applications will not run in containers that use the Linux kernel, and vice versa. Finally, you need to find out what technologies your IT department can support. Will deploying containers mean re-engineering existing solutions?

Much like any other technology, a successful container strategy requires appropriate planning, effective execution and proper management. That said, the benefits of containers cannot be underestimated. That's why containers are rapidly becoming the capstone of hybrid cloud implementations as well as the foundation of convergence. ■

SUCCESSFUL CONTAINER STRATEGY REQUIREMENTS



05

**NOW
TO
NEXT****Lessons for leaders: Packaging your environment for hybrid IT**

Enterprises need to distinguish between the appropriate use cases for VMs and containers.

Containers shouldn't be used in a vacuum; you need appropriate management tools in place.

The ideal container platform is OS-agnostic.

06

The edge vs. central IT

By Pedro Pereira

-
- Why the edge?
 - It's all about location
 - We want it now!
 - Balancing the public cloud vs. on-prem
 - Look to the future

Where do my apps and services belong: Edge vs. central IT

The Internet of Things (IoT) promises to connect everything that can be connected. But grasping its inevitability only gets you so far. Organizations must meticulously plan and execute their IoT implementations to produce the best outcomes possible.

You need a strategy to handle the inflow of information from sensors and monitors in the field, and to guide decisions about where to locate the compute power, storage and analytics engines essential to successful IoT implementations.

Why edge computing matters

So where do your apps and services belong? Some will inevitably reside in the cloud, but cloud infrastructures cannot efficiently handle the massive loads of data the IoT is expected to generate. Despite the cloud's scalability, cost-effectiveness and support for future architectures, latency issues can get in the way of the real-time processing necessary for IoT implementations.

That won't be happening at the core network, either. The cloud's *raison d'être*, after all, is to relieve central IT of ever-increasing demands for data processing, analysis and storage. Some other solution is needed between the core and cloud. And that's where edge computing comes in.

Edge computing allows you to place compute power closer to the action — the network edge. This is where many of the IoT's analytics and monitoring applications will reside to enable real-time decision-making. As IoT implementations get underway, a web of micro data centers will sprout at the edge. They will act as way stations between cloud servers, core IT and the vast networks

By 2017,

60%

of global manufacturers will use analytics to sense and analyze data from connected products and manufacturing to optimize increasingly complex portfolios of products.

Source: IDC, March 2016.

of sensors and monitors that capture and transmit data. Like a rail system with stops between major hubs, these micro data centers will ideally transform the IoT into a well-organized data delivery system.

Edge computing promises to play an essential role in the network of the future as it evolves to accommodate IoT needs. That network will be a hybrid combining cloud, edge and central IT components, with applications — or pieces of applications — residing in these distinct but integrated areas.

Location, location, location

As in real estate, edge computing comes down to location. The closer you place processing and data, the more agile your organization becomes. Now you don't have to wait for data to travel from the source for hundreds or thousands of miles to a cloud data center to be processed and redirected to a technician or analyst in front of a dashboard somewhere else.

Funneling data to the cloud potentially wastes precious seconds — or even minutes — that can make a real operational difference. A driverless car at an intersection can't wait several seconds for information from the cloud before it starts moving again. If the vehicle sits there too long waiting for data, it is bound to cause a traffic snag or even an accident.

As connected cars become more sophisticated, they will be able to communicate with each other about road and weather conditions. For instance, The DS Virgin Racing Team has teamed with Nokia Liquid Applications to use an LTE network to warn vehicles as they approach road hazards.

“Edge computing is used to extend existing cloud services into the highly distributed mobile base station environment so that road hazards can be recognized and warnings can be sent to nearby cars with extremely low latency,” according to a [Nokia blog](#). Google's Waze mobile navigation application performs similar services, albeit they require humans to inform the system about traffic slowdowns and potential road hazards.

“Streaming data is analyzed at the point of collection, providing real-time insight that allows [the DS Virgin Racing team] to make real-time adjustments to maximize the systems that control their car, and hopefully win the race.”

Kelly Pracht, senior manager of HPE Edgeline IoT Systems

Edge computing has a place not only on regular roadways but also on the racetrack, where cars running at 140 mph can transmit sensor data to the pit crew. This scenario is already a reality in Formula E, where the DS Virgin Racing team uses the compute power of a trackside data center provided by Hewlett Packard Enterprise to optimize car performance.

“Streaming data is analyzed at the point of collection, providing real-time insight that allows [the team] to make real-time adjustments to maximize the systems that control their car, and hopefully win the race,” says Kelly Pracht, senior manager of HPE Edgeline IoT Systems, in a [recent blog](#). “After the race, aggregate data is analyzed for deeper insights.”

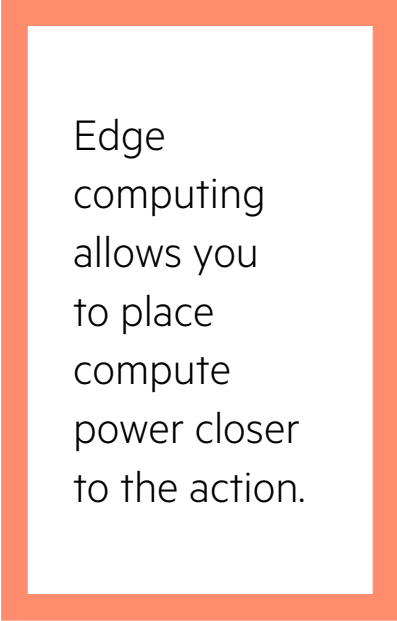
The power of immediacy

Away from roadways and racetracks, edge computing is starting to make a difference in other industries. For example, healthcare providers increasingly rely on connected devices that deliver vital information to applications monitored by medical personnel.

At-home monitoring devices track patients’ weight, blood pressure, heart rate, insulin levels and other metrics. The data travels to software monitoring systems that issue alerts to the smartphones, tablets and stationary monitors of nurses and doctors if intervention is needed. Any latency here is potentially a life-and-death situation. The same is true in [teleICU](#), which allows critical care medical personnel to connect remotely with intensive care unit patients through real-time audio, visual and data links.

Slow-loading screens or pixelated video images won’t cut it in these scenarios. However, not all edge computing instances come down to life and death. In retail environments, for example, the combination of Wi-Fi and smartphones can create internet-like shopping experiences.

A shopper who has previously registered for the store’s Wi-Fi connection will be recognized by the network as she walks in. Wi-Fi analytics software brings up relevant information such as previous purchases and time spent at the store. The system tracks the shopper through the store and sends promotional information to nearby digital displays or texts coupons to her smartphone. The goal is to get the customer to spend more money while feeling the retailer is attuned to her needs and wants.



Edge computing allows you to place compute power closer to the action.


Where the cloud excels

Edge data centers will be essential to IoT adoption in hybrid environments where real-time decisions are paramount. However, cloud infrastructures will still provide essential scalability, flexibility and storage for certain applications.

The cloud can handle massive volumes of data for which no immediate action is required. Analysts can mine that data later to identify patterns and trends that can be used for preventive maintenance and predictive purposes. For instance, cybersecurity solutions are being developed that identify the sources and methods of attacks to forecast future attacks, giving organizations a greater chance at preventing breaches.

Long term, large-scale data storage will remain an essential cloud function. So will web applications affected by seasonal fluctuations, such as retail websites that require extra capacity during the holidays or accountants who need to scale up during tax-filing season.

The cloud also makes sense for applications for which demand is hard to predict, along with testing environments and — increasingly — mobile app development and management. Cloud-based software development accelerates the development process and keeps down costs, helping organizations achieve the agility they need to compete in fast-paced markets.



BY 2019, AT LEAST

40%

of IoT-created data will be stored, processed, analyzed and acted upon close to or at the edge of the network.

Source: IDC, November 2016.

What to keep in-house


At least for the foreseeable future, certain applications will need to stay on-premises. There are compelling reasons for this. In some cases, it's more expensive to move applications to the cloud or replace them with cloud apps. Some executives still get nervous about giving up direct control of assets by moving them elsewhere. And there are lingering concerns about cloud security, privacy and regulatory compliance.

From a technical standpoint, a compelling case for keeping applications in-house can be made based on these criteria:

- Extensive redevelopment and integration needed for applications to run efficiently in a cloud environment
- Applications requiring extensive customization to meet corporate requirements
- Applications tightly linked to vast, complex databases
- Comparable cloud-based applications lack required functionality
- Mainframe applications that serve as hubs of data integration, such as enterprise service bus software, and can't be moved without also moving all dependent applications

Hybrid future

Hybrid environments combining edge, cloud and in-house assets will become as commonplace as client-server systems were not long ago. Years from now we won't even use the word "hybrid" to describe these environments. Instead we'll call them "the network."



06

**NOW
TO
NEXT****Lessons for leaders: The edge vs. central IT**

It is critical to understand how data is generated and consumed via your enterprise infrastructure.

The correct balance of edge and central IT will enhance IT agility while retaining security and reliability.

Proper implementation will reduce OpEx for services and applications by lowering networking costs and minimizing latency, and lead to more satisfied consumers.

CONTRIBUTORS



Pam Baker

A prolific and versatile writer, Pam Baker writes about technology, science, business and finance for leading print and online publications including InformationWeek, CIO, Institutional Investor, Fierce Markets Network, E-Commerce Times and many others. Her published credits include traditional books, e-books and several analytical studies on various technologies for research firms on two continents. Among other awards, Baker won international acclaim for her documentary on the paper-making industry, and is a member of the National Press Club and the Internet Press Guild (IPG).



Alyson Behr

Alyson Behr (@alysonbehr) is a content development expert, corporate communications consultant, as well as a seasoned technology editor, journalist and speaker. She is a regular contributor of technology features and product reviews to Computerworld, PC Magazine, and ArsTechnica. She is formerly editor in chief of Plane & Pilot Magazine and is a columnist for Contrails Magazine, the leading publication for executive jet owners.



David Chernicoff

David Chernicoff is a managing editor at Enterprise.nxt, an HPE site that covers the intersection of IT and business strategy. He brings close to 30 years of experience in IT to his writing and editing. After running testing labs for major magazines in the 1990s, he went off on his own, providing consulting services to business across the SMB market while writing books, magazine articles, and blogs on topics as diverse as desktop migration and data center energy efficiency. His experience ranges from database and software development to testing management, to being the CTO at a network management ISV.



Ken Hess

Kenneth "Ken" Hess is a freelance writer who is also a practicing technologist and has worked in enterprise data centers and with enterprise technology for more than 20 years.



Richard McGill Murphy

As the editor in chief of Hewlett Packard Enterprise, Richard McGill Murphy leads HPE publishing initiatives across all media platforms worldwide, working to articulate the company's point of view on digital transformation and the future of computing. Murphy started his career as a war correspondent in Afghanistan. He has covered technology, business, global affairs and popular culture for a wide range of media outlets, including Fortune, the New York Times Magazine, the New Republic and VH1.



Steven J. Vaughan-Nichols

Steven J. Vaughan-Nichols, aka sjvn, has been writing about technology and the business of technology since CP/M-80 was the cutting-edge PC operating system; 300bps was a fast internet connection; WordStar was the state-of-the-art word processor; and we liked it. His work has been published in everything from highly technical publications (IEEE Computer, ACM NetWorker, Byte) to business publications (eWEEK, InformationWeek, ZDNet) to popular technology (Computer Shopper, PC Magazine, PC World) to the mainstream press (Washington Post, San Francisco Chronicle, BusinessWeek).



Frank Ohlhorst

Frank Ohlhorst is an award-winning technology journalist and IT industry analyst, with extensive experience as a business consultant, editor, author and blogger. Ohlhorst works both with technology start-ups and established technology ventures, helping them to build channel programs, launch products, validate product quality, create marketing materials, and author case studies, e-books and white papers.



Pedro Pereira

Pedro Pereira is a New England-based writer who has covered the IT industry for two decades. He currently focuses on cloud, remote services, security and IoT.



**Read more about the tools you'll need to accelerate
your organization's transformation from now to next.**

Visit insights.hpe.com



**Hewlett Packard
Enterprise**

© Copyright 2017 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for HPE products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HPE shall not be liable for technical or editorial errors or omissions contained herein.

a00006640ENW, May 2017



SHARE THIS WITH YOUR NETWORK