CREATING TOMORROW'S DATA CENTERS

MI

MODERN INFRASTRUCTURE

The Big Data Debate

Dreaming of Converged Endpoint Management

Eunice: The Road to Oz

JULY 2013

The Trouble with Private Cloud

Why is adoption so slow?

TechTarget

Two Steps Forward, One Step Back

Overall, these are pretty good days for IT professionals. There are lots of new technologies coming down the pike that are bringing innovation and efficiency, and stable-to-rising budgets to take advantage of them. And when things are going well, it's easier to take minor setbacks in stride: "Two steps forward, one step back," as the saying goes. Still, rarely does a new technology make its way into the organization without causing problems for IT.

Take mobile devices. While no one with an iPhone in his or her pocket disputes its value, the proliferation of these devices is no picnic for IT managers, writes Jim Furbush in "Beyond Endpoint Management." Endpoint management tools used to focus on a single device—the company-owned PC—to handle tasks such as patch management, asset reporting, OS updates and application distribution. Now, endpoint management tools are being called upon to wrangle a variety of end-user owned laptops, tablets and smartphones, with more devices on the way. The discipline now comprises security, application delivery, identity management and profile management—a hugely expanded list of responsibilities.

Virtualization has also put IT organizations ahead by two steps, arguably even more. But there, too, IT is a victim of its own success. Virtualization is pervasive in enterprise IT shops, so the next logical thing to do is to build a private cloud on top of it. But as I found in my article "The Problem with Private Cloud," that's easier said than done. While a private cloud looks good on paper, the tools to build one on top of existing virtual infrastructure can be complex and expensive, and they may not deliver the same level of value—consolidation, agility—that virtualization first brought to the table.

Some organizations may choose to skip the private cloud exercise altogether and go directly to the public cloud. In particular, VMware's new vCloud Hybrid Service looks like it might be a good way for VMware shops to jump on the cloud bandwagon and start reaping the benefits of shared, pay-by-the-drink infrastructure (see "Wooing the VMware Faithful with vCloud Hybrid Service"). It certainly seems like an easier pill for IT pros to swallow than Amazon Web Services. But if IT does adopt VMware's new public cloud, dollars to doughnuts they'll uncover more than a few unpleasantries they hadn't counted on.

All this is par for the course, and as long as "Two steps forward, one step back" doesn't devolve into "One step forward, two steps back," we are still ahead of the game. ■

CURRENTS 7/13

All the news you can use about modern infrastructure

HOME

ONE ON ONE

SDN in Its Infancy

Software-defined networking has been touted as a technology that can finally solve many of the issues that networks have traditionally posed for virtualized environments. SDN can help move data center environments closer to the cloud by bringing the network to parity with other aspects of virtualized infrastructure. SDN is still in its infancy, however.

Bob Muglia, executive vice president of software and solutions at Juniper Networks, talked with *Modern Infrastructure* about the emergence of SDN, what it means for enterprise data centers, and why IT managers should care about this still-evolving technology.

Prior to joining Juniper, Muglia spent more than two decades at Microsoft. Most recently he was the president of the Server and Tools division and worked on some of the company's key virtualization- and cloud-based technologies. As Muglia noted, "When we were building System Center and Hyper-V, and when we were building Azure, . . . the lack of dynamic capabilities in the network was a significant bottleneck."

What are the potential benefits of SDN for an enterprise-class data center?

We have watched companies move from an environment with physical infrastructure to one where they have a virtualized infrastructure. They are now looking for more of the dynamic attributes they can derive from a cloud-based infrastructure.

The difference [with a cloud-based infrastructure] is the need for the network, the whole system, to respond. In a virtual infrastructure, when a business unit wants to put up an application, they send out an email to the IT department, who says, "Great. We'll get this up within the virtual system in a couple of days."

In a cloud-based infrastructure, they go to a portal, they talk about the resources required and 10 minutes later the system has spun up all those resources that are

available for use.

And in that world, the underlying elements of the infrastructure—the compute, the storage, the network—all have to dynamically respond to the changes required. And what we have seen is, the compute infrastructure first, then storage infrastructure, have been quite responsive to the needs of the business, but the network has not. So that is the business problem that SDN is focused on solving: to make the network dynamically respond and to meet the needs of the business.

Why has traditional, physical networking been a bottleneck?

There are a lot of elements in the network, a lot of devices—way more than there are storage devices, for example. Basically each device is configured independently and the configuration data is stored and mastered within that device. So there really is no centralized point of control, no standardized way to automate that process. That has been the fundamental issue.

Are there benefits in terms of cost and performance?

Yes, the biggest benefit of cost will be the operational cost that is reduced and the people cost associated with managing the system. If you look at any IT shop, the dominant cost is the people side. [So] you will see a reduction of cost because there is a lot of manual complexity associated with running the network.

Are there factors that might hinder acceptance of SDN, such as in the area of standards?


BOB MUGLIA

The biggest issue that will hold back adoption is the state of maturity within the industry around this. To say it is immature would be an understatement. This is the early, early, early days of it. Less than 1% of all data center systems are SDN-enabled.

Over time, as the industry goes through its typical maturation process, standards will emerge. [And] while OpenFlow is an interesting standard, there are many others that need to emerge as well.

One is XMPP. We think there are standards existing today that can be used effectively, standards for the data path in particular, such as MPLS over GRE or VXLAN. Those standards are usable as-is.

At this point, there isn't any one protocol where we sit back and say, "This must be the answer." The issue isn't so much the protocols as the architecture of the solution. The more macro point is that we do need to see broad-based adoption of some standards.

For those getting ready to implement an SDN initiative, what sorts of things should they think about?

1. Centralize management. This is a very achievable thing to do in a shorter period of time and generates

significant benefit in operational cost savings, to move away from a distributed, typical approach of managing each individual network element as an independent unit.

2. Evaluate network services. Today these services are chained together in a physical way; I have an application delivery controller box, I have a static firewall box—different boxes that are connected. [Shops can look at] which of those services could be virtualized, pull them out and begin working with virtualized services.

3. Centralize controller architectures. This is where many of the benefits of SDN are fully realized. But it is still fairly early for this.

4. Optimize the usage of network and security hardware to deliver high performance.

Are there certain types of shops where SDN is not the answer?

Unless you are contemplating installing a private cloud, SDN doesn't make sense. There is some minimum threshold of servers that are needed to justify a private cloud. Truthfully, no one knows what that is today. I can easily tell you if you have 500 or 1,000 servers, I think it makes sense. Fifty servers colocated in a data center, well, you know, there is a good chance it could benefit you. If you have 50 servers located across 10 branch offices nationwide, then [that's] probably not enough.

—LAUREN HORWITZ AND ED SCANNELL

Do you run a private cloud?

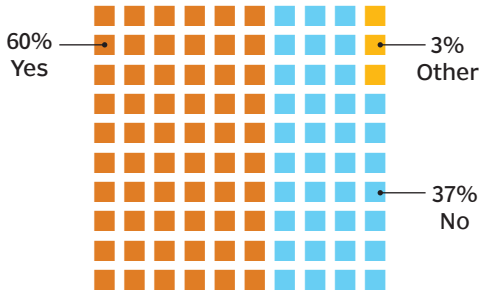

Will Desktop as a Service drive more virtual desktop adoption?

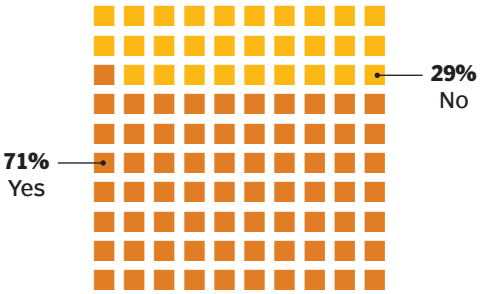

FIRST LOOK

Apple iOS 7

Apple's iOS 7 has a whole new design, with clean, colorful icons that seem to float above the background. But as an IT professional, you probably don't care. You want to know if there are any features that pose security concerns and if there are new management capabilities. The answer to both those questions is a resounding yes.

Apple unveiled iOS 7 at June's Worldwide Developer Conference (WWDC), and details have been slow to trickle out. But CEO Tim Cook and company provided an overview. Here are three things you need to know:

1. AirDrop? More like ScareDrop!

The AirDrop feature lets people share files with others in proximity using peer-to-peer Wi-Fi. A user enables AirDrop when accessing a compatible file, and the names and photos of nearby friends will pop up. One click transfers that file to an iPhone or iPad. Apple continues to make sharing easier—it opened iCloud up to third-party apps in iOS 6—and that should continue to feed IT's concerns about data loss. It's unclear whether IT can disable AirDrop, but it's something to watch.

2. Automatic for the People

On Capitol Hill in May, Cook fielded a question from Sen. John McCain: Why do I always have to update my iPhone apps? Manual app updates have long annoyed iOS users, but their time is up. In iOS 7, apps will automatically update themselves in the background. It's a great convenience for users—McCain even tweeted his thanks to Cook after the WWDC keynote—but what if the new version of an app adds Dropbox integration or other features that could pose security and compliance problems? We don't yet know if IT can turn these off.

3. Apple Getting into MAM?

On the surface, the WWDC keynote didn't focus much on enterprise mobility management, but there were lots of clues if you paid attention. Apple showed a slide that mentioned App Store volume purchases, enterprise single sign-on and managed app configuration. An exec also said iOS 7 will include a per-app VPN feature. These are all components of mobile application management (MAM) technology, so you may want to hold off a MAM purchase until we know how much functionality will be built into the next operating system. **—COLIN STEELE**

"A few years ago, it was security through obscurity. People didn't know how to get to information and how to mismanage it. Hackers are getting smarter."

—EUGENE SHABLYGIN, CEO of WWPass, at the MIT Sloan CIO Symposium

"We don't give them what they ask for. We give them what they need."

—MICHAEL RELICH, CIO of Guess Inc., talking about IT responding to business users, at the MIT Sloan CIO Symposium

"We are trying to make a data center operating system that's agnostic to compute, network, storage and hypervisor."

—SARAN MANDAIR, senior director of platform engineering and operations for PayPal, on whether the company plans to rip and replace VMware's vSphere in the process of deploying OpenStack

"Native apps on mobile devices is the ultimate answer."

—SCOTT WRIGHT, technology architect, Marathon Oil, at Citrix Synergy

"Thin clients should be zero clients. ... Zero clients come from the fact that we as thin-client vendors weren't doing our jobs."

—IAN GEISER, Devon IT CTO, at Citrix Synergy

"There are forest people, trees people and then there are bark people."

—JOHN HALAMKA, CIO of Beth Israel Deaconess Medical Center, on those who work with IT health care standards, at the MIT Sloan CIO Symposium

NEWS IN REVIEW

Wooing the VMware Party Faithful with vCloud Hybrid Service

VMware's virtualization software is pervasive in enterprise IT shops, but, so far, the company hasn't been able to parlay that dominance into the cloud. Its newly launched vCloud Hybrid Service could change that, though—at least for existing customers.

After many years of declining to offer its own public cloud service, VMware's vCloud Hybrid Service became available to early-access customers in May and is slated for general availability in the third quarter of 2013.

In the past, VMware referred customers who wanted public cloud to vCloud Powered Services partners such as Bluelock and CSC using its vCloud Connector tool. That approach had its drawbacks, said Mitch Northcutt, executive vice president for strategy and services at AHEAD, an IT solutions provider in Chicago. With a vCloud Powered Services public cloud, "you have no inside-out federation, no seamless management," Northcutt said.

The VMware-branded Infrastructure as a Service (IaaS) offering, alternatively, lets VMware vSphere and vCloud Director shops extend their virtual environments into VMware's public cloud using existing VMware tools such as vCenter while offering automated replication, monitoring using vSphere vMotion, high availability and Distributed Resource Scheduler. Customers can pay for the service on a subscription or per-use basis.

The "vCloud Hybrid Service will be the path of least resistance," Northcutt said, pointing out that users won't face high transformation costs. Plus, he said, they'll be comfortable with the tool set and have trust in the VMware name.

Participants in the early-release program are excited about it.

"If VMware can get this right—which it looks as if they are—then organizations coming up to refreshing their infrastructure are going to find it difficult to resist," said Neil Smith, a virtualization architect with a U.K. hedge fund and a vCloud Hybrid Service beta user. "The large majority of [corporations] are moving from a Capex- to an Opex-type model, which fits the cloud service very nicely: one monthly bill and no elevated infrastructure maintenance costs after three years of running that hardware."

Among vCloud Hybrid Service's selling points are a "slick [user interface]" that is very easy to use, pay-by-the-hour billing and the ability to redeploy certain VMware licenses, Smith said. He envisions using vCloud Hybrid Service for developer workloads as well as for select internal applications that are evolving into Software as a Service-type services, when the service becomes available in Europe in the fourth quarter.

Meanwhile, the city of Melrose, Mass., is eyeing the

vCloud Hybrid Service as a disaster recovery (DR) target for its on-premises NetApp FlexPod environment, said Jorge Pazos, the city's CIO.

"As a municipality, I don't have a lot of luxury with geographic diversity," Pazos said. Currently, the city replicates select workloads to a DR site one mile away. "If we had a big event like a massive power outage, we could be in trouble," he said.

"IT'S A VERY, VERY FAINT SHADOW OF AWS."
—Carl Brooks, research analyst, 451 Research

The city also explored Amazon Web Services (AWS), Microsoft Azure, Rackspace and Bluelock for cloud-based DR, but vCloud Hybrid Service is more interoperable with the current environment, Pazos said. "With those, there's no smooth way to move workloads back and forth between them and on-premises."

Eventually, he hopes to see vCloud Hybrid Service integrated with VMware's Site Recovery Manager for fully automated business continuity. "That would be a dream come true."

The other advantage of vCloud Hybrid Service over other cloud providers is pricing predictability, Pazos said, thanks to the subscription-based pricing model. "With the other offerings, it's really hard to get a consistent sense of what your costs are going to be," he said. "It's easy to shoot yourself in the foot; if you leave a [virtual machine] running accidentally, your costs can shoot up."

VMware's Offering Isn't Amazon

But vCloud Hybrid Service should not be confused with the more established public clouds, said Carl Brooks, research analyst at 451 Research.

"It's a very, very faint shadow of AWS," Brooks said. "It's not an actual hyperscaled, globally available, elastic and apparently infinite platform." Instead, "VMware is renting space from a hosting provider and delivering vCloud," he said.

As such, vCloud Hybrid Service is probably not a fit for forward-looking developers who want to run complicated, scale-out Web applications or extensive high-performance computing workloads. But it may be exactly what mainstream IT users need.

"We do extensive surveys of enterprise IT, and the overwhelming subtext of those surveys is that [IT teams] are being harassed about cloud," Brooks said. "So if you've already got VMware, when people start hassling you, you can say, 'Fine, here's your server,' and it appears as a line item on your monthly bill. It lets them off the hook." **—ALEX BARRETT**

SUMMING IT UP

What are your plans for server virtualization over the next 12 months?

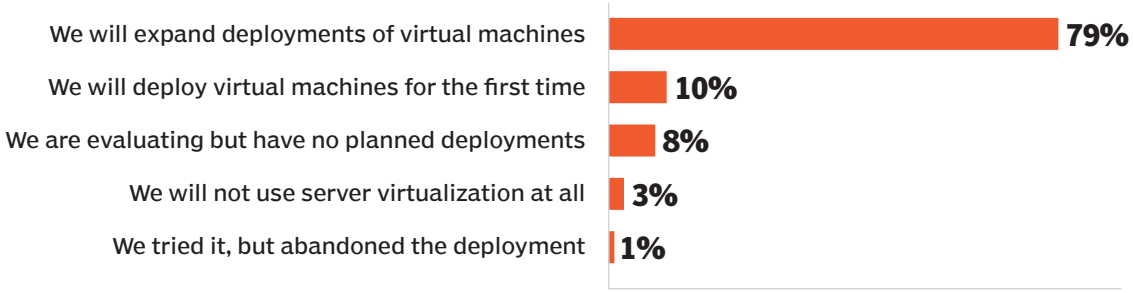

N=739; SOURCE: TECHTARGET'S DATA CENTER AND READERS' CHOICE 2013 SURVEY

86% We will assemble our own private cloud

14% We will purchase a dedicated system for private cloud

N=332; SOURCE: TECHTARGET'S DATA CENTER AND READERS' CHOICE 2013 SURVEY

What are the primary drivers behind your evaluation or adoption of integrated infrastructure technologies?

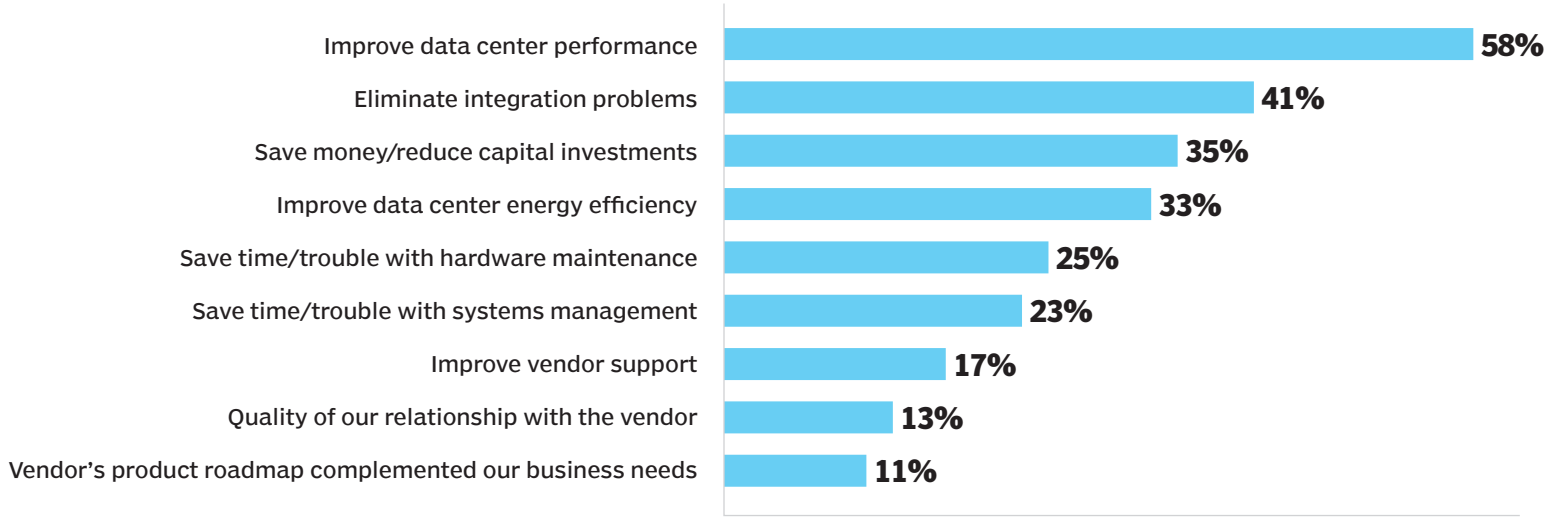

N=498; SOURCE: TECHTARGET'S DATA CENTER AND READERS' CHOICE 2013 SURVEY READERS COULD SELECT UP TO THREE RESPONSES.

SUMMING IT UP

What are your plans for cloud computing over the next 12 months?

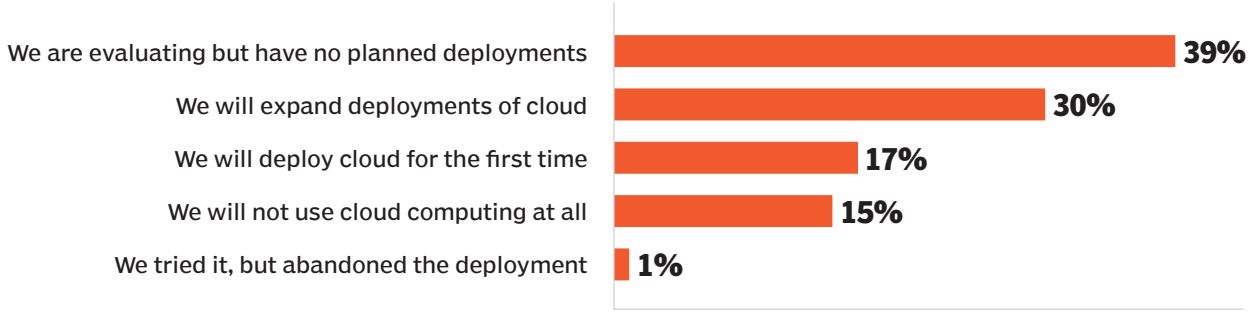

N=739; SOURCE: TECHTARGET'S DATA CENTER AND READERS' CHOICE 2013 SURVEY

How will your 2013 spending on integrated infrastructure products compare with 2012?

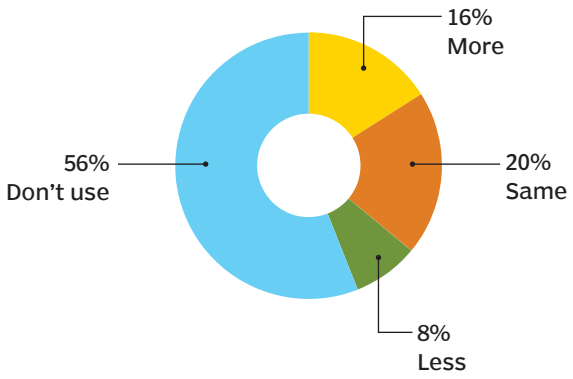

N=1,155; SOURCE: TECHTARGET'S DATA CENTER AND READERS' CHOICE 2013 SURVEY

What is your main reason for choosing your primary virtualization platform?

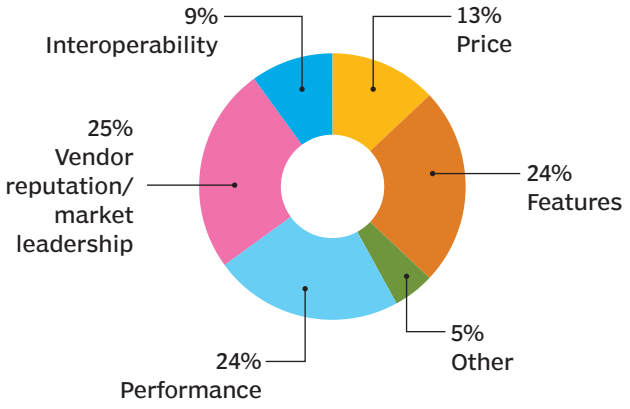

N=714; SOURCE: TECHTARGET'S DATA CENTER AND READERS' CHOICE 2013 SURVEY

UP FOR DEBATE

Big Data

Big data isn't always new data. Rather, it's data that has been accumulating, often gathering dust, over the course of many years. But new tools and processes are turning that data into live, actionable, and sometimes very revealing information.

But big data becoming Big Brother is a common concern. Many organizations already have the data they need—but they're afraid to use it. A group of MIT professors, all panelists at the MIT Sloan CIO Symposium in May, offered some enlightening—and chilling—potential uses for gathered and analyzed data. For researchers, big data tools offer access to a treasure trove of behavioral information, statistics and other hard numbers that are waiting to be discovered.

For instance, at the moment, "new treatments for advance cancers are based on the intuition of doctors," said Dimitris Bertsimas, professor of operations research at MIT. "Drug combinations are used over and over again. But can there be new combinations?" That information was easy to find. An existing database of treatments and drugs dates back nearly 40 years.

The Operations Research Center at MIT, which Bertsimas co-directs, created a database with natural language processes, pulling information from academic oncology papers. The outcome was a graph charting survival against toxicity for treatments that had been used through the years. That allows doctors to personalize treatment, Bertsimas said, eliminating guesswork in favor of hard numbers.

Masses of mobile phone data already exist, said panelist Alex "Sandy" Pentland, director of MIT's Human Dynamics Lab and Media Lab Entrepreneurship Program. "Companies gather it but are scared to use it," he said. "It's politically controversial. But we have to start to see that cell phone and public data is for the public good."

"CELL PHONE AND PUBLIC DATA IS FOR THE PUBLIC GOOD."
—Sandy Pentland, MIT professor

European telecom company Orange started a data commons, which put company data into the public domain to encourage business use. In one case, commute times were reduced by suggesting rearranged bus routes based on mobile phone location data, Pentland said.

Other cell phone data projects helped reduce infectious diseases and create a real-time census map of the Ivory Coast—a country where citizens hadn't been counted in years. That data also let researchers see ethnic boundaries, which had long been unclear.

The mobile data gathered isn't necessarily personal

smartphone data, but rather data on when phone owners moved from one cell tower to another. "There are statistical relationships between the pattern of calls and the pattern of mobility," Pentland said. "It's the heart of what that data is about."

A Fine Line

Who can argue with curing cancer and encouraging world peace? But analyzing big data can quickly turn into snooping, researchers said. Predicting crimes before they happen is one example. Police could deploy extra officers based on common crime location information that's been cross-referenced against people of interest sending text messages to one another, Pentland said. The texts themselves wouldn't be revealed—simply the fact that the messages are being sent at all.

Ensuring privacy is possible when individual information is stripped from its associated data, said Andrew Lo, professor of finance at MIT. "To prevent big data from becoming Big Brother," he said, "we have to protect the privacy of individuals, and we have the tools to do that." Cryptographic methods could allow researchers to get aggregate statistics from encrypted personal information stored on a central server while still ensuring privacy, he said.

The sky's the limit for actionable intelligence based on big data in fields such as health care, consumer behavior and organizational management, according to the MIT researchers.

Andrew Lo and Sandy Pentland at the MIT Sloan CIO Symposium.

"You can track patterns of conversations [among employees] with name badges and automate an office layout and productivity, automate interactions to produce greater output," Pentland said.

Till then, it's cultural, not computational, concerns that reign. "I'm very optimistic that a lot of the things around privacy and data ownership will be taken care of," Pentland said. He envisioned an opt-in framework, where subjects have the ability to audit which information they share.

"The real challenge here is making the data available," Pentland said. "People have to be willing to share their data and make it available."—**CHRISTINE CIGNOLI**

IN THE MIX
BOB PLANKERS

Where AWS Beats On-Premises

HOME

I TEND TO be pro-private data center. Some have compared my views to those of the Luddites, the 19th-century textile artisans who protested the mechanization of the textile industry by destroying the machinery itself. I don't advocate destruction, but I do believe there is tremendous value in on-premises data centers in the face of issues in the public cloud—issues like control, transparency and even affordable secure connectivity.

Ironically, one area where the public cloud is rapidly emerging as the winner is in secure computing, specifically in environments that must comply with regulations such as PCI guidelines, or even worse: comply with the U.S. Federal Risk Authorization Management Program, or FedRAMP. These are areas that are troublesome for on-premises data centers, especially if most of the workloads are not secure. The secured environments require enormous amounts of duplicate infrastructure (including separate data centers), physical security and infrastructure controls that most organizations are completely unprepared for.

Companies and universities that do certain types of research funded through federal grants may not have a choice but to comply, as the Federal Information Security Management Act (FISMA) mandates compliance with particular standards. As a result, many organizations see the economics of continuing their research becoming very unfavorable. Many have stopped their work altogether, and their researchers have gone elsewhere, including to foreign countries where the controls are looser.

Security at Scale

Amazon Web Services (AWS) recently announced that they have been deemed compliant with FedRAMP guidelines, for FISMA "low" and "moderate" levels, corresponding to the same levels in the United States' National Institute of Standards and Technology (NIST) SP800-53 guidelines, the federally mandated rule book for implementing such guidelines. These rule books are enormous, though, and they are often just guidelines, which complicates matters. In security, the idea

of "compensating controls" means that it's OK to avoid a mandated type of security control, as long as there are other methods in place to achieve the underlying goal.

AWS IS A DE FACTO STANDARD IN CLOUD COMPUTING.

This makes certification of environments difficult, and makes the process very subjective. Being able to automatically pass a large part of the subjective certification process through the use of commodity services saves substantial time and money.

The Amazon announcement is huge because of the scale of AWS, too. The ability to buy FISMA-compliant infrastructure as a service that scales and is interoperable with all sorts of management tools is a giant step forward. Previously there were only two certified cloud providers: CGI Federal and Autonomic Resources LLC. Eighty more providers have applied for certification, but they face a serious competitive challenge because AWS is a de facto standard in cloud computing. But the big impact of this move can be found in AWS' frequently asked questions: "Will compliance with FedRAMP increase AWS service costs?" Their answer is "No, there are no additional costs." And with that answer, on-premises, secure private and hybrid clouds have died, to be replaced on their next refresh cycle with the public cloud. ■

BOB PLANKERS *is a virtualization and cloud architect at a major Midwestern university. He is also the author of* The Lone Sysadmin *blog.*

END-USER ADVOCATE
BRIAN MADDEN

The Year of VDI is Finally Here (Really!)

HOME

How many times have we heard that this is the year of VDI? Probably every year since 2006. Until recently, I had yet to see evidence that virtual desktop infrastructure would ever be ready for prime time. (I even co-wrote a book called *The VDI Delusion*, in which we essentially spent 240 pages explaining why VDI is not that great.)

But that has changed this year with the advent of two new technologies.

The first involves storage. I always argued that VDI had to be about persistent 1:1 disk images, but that doing so with traditional server storage was prohibitively expensive. This is why most VDI deployments have focused on nonpersistent "shared" images, and why VDI adoption has remained marginal. Now, vendors such as Atlantis Computing, GreenBytes, Tegile Systems (and probably 20 others) offer block-level, single-instance primary storage. This means that you can now build a VDI environment for hundreds of users—with each user having a unique disk image—for a price that just last year would have bought you only a shared-image system.

The second major breakthrough of the past year has been in the area of graphics performance. Multiple vendors now have plug-in cards for the remote VDI host servers that can offload the processing, encoding and compression of the remote protocol display streams. Adding one of these cards to a server means that it's possible to dedicate all that processing power to the graphics experience without diminishing the user experience by taxing existing CPUs. These plug-in cards also allow for GPU virtualization, which means applications such as 3-D CAD, Photoshop and video editing are now possible via VDI.

VDI: Coming to a Desktop Near You

These improvements in storage and graphics capabilities mean that virtual desktop infrastructure is now applicable in far more situations than it was before. This sentiment is corroborated by the number of large—more than 10,000-seat—deployments we've seen kick off over the past year. (In fact we've titled the second edition of our

VDI Delusion book *The New VDI Reality* to better reflect our current sentiment around VDI.)

To be clear, I'm not suggesting that the future of the desktop is VDI, or that all the world's Windows desktops should move to VDI. Over the past few years, we might have seen VDI makes sense for only 5% to 10% of the enterprise desktops out there. Now that VDI can support 1:1 persistent disk images and graphically intense applications, VDI becomes a valid option for 40% or 50% of enterprise desktops. Yay for 2013. ■

BRIAN MADDEN *is an opinionated, supertechnical, fiercely independent desktop virtualization and consumerization expert. Write to him at* bmadden@techtarget.com

THE PROBLEM WITH PRIVATE CLOUD

What will it take for private cloud adoption to skyrocket?

When it comes to private cloud adoption, you could be forgiven for thinking it's been somewhat of a disappointment.

Pundits have long insisted that widespread private cloud buildouts are a foregone conclusion, but reality begs to differ. While the technology to build a private cloud has been available for several years, uptake has been slow.

At first glance, private cloud adoption figures seem relatively robust. Forrester Research, for example, said that 31% of its enterprise customers claim to have a private cloud in place, and another 17% plan to build one over the next 12 months.

But upon closer examination, only 13% of those organizations that report having a private cloud have a "true" private cloud, said Lauren Nelson, an analyst with the firm.

"Most have adopted a software solution that improves their management capabilities," Nelson said. More often than not, those so-called private clouds don't include some key characteristics of a cloud—for example, multi-tenancy, end-user self-service or metered usage.

HOME

Part of private cloud's problem could be IT's loosey-goosey interpretation about what it is—and therefore what it brings to the table. Experts define a private cloud as dedicated resources running behind the firewall that are organized into an Infrastructure as a Service (IaaS) cloud computing platform. The National Institute of Standards and Technology, in turn, says that an IaaS cloud must include five essential characteristics to be a true cloud: on-demand self-service, broad network access, resource pooling, rapid elasticity and measured service.

IT's working definition for private cloud is quite different, said Aneel Lakhani, Gartner Inc.'s research director for virtualization and cloud infrastructure.

"What people call private cloud is a vast array of things," Lakhani said. "I've seen private cloud refer to everything from, 'We have a data center' to 'We've built something that looks a lot like Amazon [Web Services].'"

Tough Acts to Follow

That's in contrast to the wildfire that was virtualization over the past decade.

"People got into virtual infrastructure very quickly and enthusiastically—not because the vendors told them so," said Jonathan Eunice, principal IT adviser at Illuminata Inc., in Nashua, N.H. Those adopting virtualization were consolidating servers with abandon and implementing and systematizing shared services such as backup—even tackling disaster recovery, a luxury that used to be reserved for only the most mission-critical applications of the most deep-pocketed organizations.

Compared with the tremendous cost and agility benefits that virtualization brought to even the smallest of IT shops, private cloud's benefits are more subtle. "Academically, I get that there's a difference. But in my mind, there's no necessary difference between a virtual infrastructure and a private cloud," Eunice said. "Most of the benefits of building a private cloud have already been achieved before you get there."

"WHAT PEOPLE CALL PRIVATE CLOUD IS A VAST ARRAY OF THINGS."
—Aneel Lakhani, research director, Gartner Inc.

Meanwhile, demand for public cloud is surging, as evidenced by industry leader Amazon Web Services (AWS), whose 2012 revenue was estimated at $2 billion; that figure is projected to double this year. That growth isn't limited to AWS. The top 10 cloud providers grew by 37% in 2012, according to a recent report by financial services firm Robert W. Baird and Co. Inc.

At the same time, early adopters say that their private cloud efforts have been worth the time and money. And IT executives just laugh when public cloud vendors like AWS CTO Werner Vogels deride private clouds as "false

clouds" that are designed to get enterprises to buy more hardware.

"I also don't ask my barber if I need a haircut, because I know the answer," said Scott Blanchette, CIO at Vanguard Health Systems in Nashville, Tenn., during a panel at the MIT Sloan CIO Symposium in May.

While it may not have set the world on fire, "private cloud is here to stay," Blanchette said.

The Case for Private Cloud

IT shops that have implemented private cloud software wouldn't want it any other way.

"Business owners have gone from going to the data center to see the flashing lights to consuming Infrastructure as a Service," said Neil Smith, virtualization architect at a U.K.-based hedge fund that was an early user of VMware vCloud Automation Center (VCAC), derived from VMware's 2012 acquisition of DynamicOps.

With its dominance in the virtualization market, VMware was an early leader in private clouds. In 2009, the company began down the private cloud path with Project Redwood, which in 2010 was formally introduced as vCloud Director, and is now a centerpiece of vCloud Suite, which has replaced VMware Lab Manager.

In particular, VCAC automates and speeds up the process of provisioning individual servers, Smith said, and chargeback and showback capabilities have demonstrated to business users that IT is not free. "It's influenced them to clean up after themselves," he said.

However, Smith said that his organization declined to implement other elements of VMware's vCloud suite, such as vCloud Director, because "we didn't feel that it brought us much." Likewise, the firm is looking beyond VMware for advanced application provisioning services, eyeing software from ElasticBox that can span in-house and public clouds. VCAC claims application provisioning capabilities, but in Smith's experience, "it's still quite a manual intervention there."

But many organizations cite high costs for commercial private cloud solutions such as vCloud as a significant barrier to adoption, Forrester's Nelson said. Prices for a full-fledged private cloud offering from VMware or one of the "Big Four" enterprise management vendors can easily cost $500,000 or more in licenses and professional services, she said, making it cost-prohibitive for small or medium-sized organizations.

The Open Source Option

Not surprisingly, a number of open source private cloud technologies have emerged as alternatives. They feature varying degrees of maturity and adoption. Among these open source platforms are offerings from Eucalyptus, Citrix CloudStack, OpenNebula and OpenStack.

For some users of these platforms, commercial private cloud software was never even in the picture.

When it came time to build a private cloud, marketing software startup HubSpot in Cambridge, Mass., disregarded VMware because of its historically high costs, but also because of its legacy as a virtualization vendor, said Jim O'Neill, HubSpot's CTO.

"A lot of people associate private cloud with VMware, but there's a really big difference between virtualization and cloud," he said. Cloud isn't so much about being able to aggregate and manage virtualization hosts, "it's about having a programmable infrastructure," he said.

"We see the cloud as being about the advent of an API [application programming interface] through which you can control capacity and provisioning of resources, such that you can automatically move workloads and script changes," O'Neill said.

HubSpot began building an OpenStack-based private cloud in 2011 to complement workloads running on Amazon. Of the 2,000 or so instances it runs steady state, about 90% still run on AWS, with the remainder on the internal private cloud. HubSpot hopes to increase the number of workloads running internally to about 30% of the total over the coming year, focusing on data-intensive apps and leaving emerging and high-growth applications on AWS.

"The goal is to have the right cloud for the right workload," O'Neill said.

But O'Neill also acknowledged that OpenStack isn't exactly bulletproof. "OpenStack [Nova] compute and [Swift] storage services are solid," he said, "but networking is evolving, like with all clouds." He is still looking forward to more from OpenStack Quantum, a software-defined networking project that will allow users to script virtual networks via API. Furthermore, implementation isn't always easy. "We're still in the first inning; some assembly may be required," O'Neill said.

The Middle Path

Between the high-cost commercial private cloud packages, and low-cost do-it-yourself open source platforms, a sea of smaller players is vying for the attention of would-be private cloud implementers, all promising to facilitate the implementation of a private cloud more easily or for less money.

"We feel that private cloud is very popular, but a lot of SMBs don't have the skill sets to stand up a private cloud," said Paul Speciale, CMO at cloud provider Appcara, which focuses on making it easy to layer application services onto a private cloud. "Even for a [specialist], it takes a full day of work, and they know what they're doing."

As such, there's a veritable cottage industry of companies offering some level of open source private cloud bundles, said Gartner's Lakhani. At last count, Gartner tracks more than 80 companies that make some form of private cloud software.

For instance, there's Appcore, which packages CloudStack private clouds on top of in-house hardware; Piston Cloud Computing for OpenStack shops; and CloudWeavers, a U.K. company whose software on a pluggable USB stick auto-discovers and configures an OpenNebula private cloud.

Some private cloud startups are taking their cues from as a Service and managed service providers when approaching customers.

Gravitas of New York, N.Y., targets hedge funds with its Gravitas Private Cloud running on a partitioned VCE

Vblock. Appko delivers a private cloud on top of private converged infrastructure through its self-service portal called MyCloudPod that is delivered as Software as a Service (SaaS), over the Internet. Metacloud in Pasadena, Calif., uses the SaaS metaphor but builds an in-house OpenStack private cloud that it monitors and manages remotely.

It's the Use Case, Stupid

Simplicity notwithstanding, an equally important matter for organizations is to consider their goals in building a private cloud.

The answer to that question varies a bit, said Jay Litkey, CEO at Embotics, a virtualization and cloud-management software provider.

"What an organization wants from a private cloud depends on what business they are in; it's not a one size fits all," he said.

Generally speaking, Litkey said, private cloud deployments tend to follow three broad use cases: improved service provisioning through automation; self-service for test and development environments; and help in moving to a shared-services model, especially within the education and government markets.

"We used to see a lot of distributed IT," Litkey said. "Now, IT wants to be the central service provider, and bring all its tenants into a multi-tenant private cloud."

Alternately, developers working on a new application type will help bring a private cloud online, said Carlo Daffara, an engineer with CloudWeavers.

"Most organizations do not want to build a private cloud per se; they want to build a platform for a specific application," he said. In particular, Daffara said that many shops build a private cloud when they are looking to launch a major new application development project based on, say, NoSQL or scalable Web services.

In the future, that's probably what will push private cloud from the shadows and into the mainstream, said Illuminata's Eunice: the increasing popularity of scale-out, grid-style applications such as big data analytics.

"Not all apps are created equal, and a lot of the apps that run so well on service provider clouds and grids are not necessarily the same [enterprise] apps that run on traditional scale-up infrastructure" (i.e., your average relational database or ERP system), Eunice said.

That new class of apps—video capture, big data analytics, select HPC workloads—"a lot of them used to be 'don't care' applications," he said. "Now, they're running Facebook or Google or Bank of America's customer service portal, and a good portion of them run very well on a cloudy infrastructure." ■

ALEX BARRETT *is editor in chief of* Modern Infrastructure.
Write to her at abarrett@techtarget.com.

BEYOND ENDPOINT MANAGEMENT

Endpoints are a quickly moving target, and the technology market is trying to improve its aim.

Back when the office was contained within four walls and employees had only one computing device that never left the building, endpoint management was fairly simple. Endpoint devices—basically users' PCs—needed to comply with specific criteria before users were granted access to network resources.

But there's no longer just one endpoint per employee. What used to just be a single corporate-provided Windows PC to access data and applications has ballooned into a complicated mess of ownership scenarios and different operating systems. Now, users might have any combination of the following to do their jobs: a Windows PC provided by the organization, a personal Apple computer at home and a smartphone or tablet.

The evolution of endpoint management, then, is really the story of the evolution of the employee workspace

HOME

(desktop, applications and data) and all the management and security tools that have been developed to support and control it.

For most IT pros, untangling those methods to securely deliver data to employees' endpoints can be daunting because it has evolved so much over the past decade—and rightfully so, said John Little, chief technology officer at Venture Technologies in Ridgeland, Miss.

"When everything was a corporate asset, you could control the information by controlling the device," Little said. "Information was stored on the corporate servers, and it wasn't going anywhere. [Only] once that [intellectual property] started to leave the perimeter because of new devices and cloud apps did we realize we really needed to control the leakage of that information."

The transformation of the employee workspace has increased at warp speed with the adoption of cloud, social and mobile technologies. These tectonic shifts have forced IT to examine how best to secure data and deliver services beyond the firewall.

Endpoint security and management systems can be purchased as software or as dedicated appliances. Such systems can enable IT to discover, manage and control computing devices that request access to the enterprise network. Endpoints that do not comply with corporate policy can be controlled, to varying degrees.

These management systems, such as IBM's Tivoli Endpoint Manager, Novell's ZENWorks Endpoint Security Management and LANDesk Management Suite, can come with a variety of features including automated patch management, asset reporting, application distribution and OS migration tools. Beyond specific functions, though, these tools serve a greater purpose.

IT's objective has always been controlling the security of its data and delivering services to employees as efficiently as possible, said Kent Christensen, director of data center services at Datalink Corp. in Eden Prairie, Minn.

"It's really that simple and difficult," he said.

Endpoint Management Is a Tangled Web

Unfortunately, there are so many pieces to this complex puzzle that no approach works like a Swiss Army knife today. It also doesn't help that easy management of user workspaces is a moving target. New products, features, processes and philosophical approaches are being introduced all the time.

How an organization decides to manage and secure endpoints depends entirely upon myriad unique factors, said Jason Blackett, product manager of endpoint solutions at Novell Inc.

For starters, desktop virtualization has allowed many organizations to abstract the desktop layer off the endpoint device for central management in a data center.

Sunrise Health Region, a health care organization based in the Saskatchewan province of Canada, conducted a 3,000-seat virtual desktop rollout in 2012 to improve clinician application login times and provide a centrally managed, single desktop golden image.

"We are introducing a new way for secure printing,

single sign-on for applications and smart cards for quick and easy logging on," said Sheranga Jayasinghe, Sunrise Health Region's director of IT, who oversees the delivery of services for 3,500 employees and approximately 120,000 patients.

ENDPOINT MANAGEMENT HAS BEGUN TO CONVERGE WITH IDENTITY MANAGEMENT, SO SOME IT SHOPS ARE MANAGING DATA, IDENTITY AND ACCESS TO SERVICES.

The organization wanted to move away from buying new PCs, which it hasn't had to do in six months, in favor of thin-client terminals. Sunrise Health Region also wanted to allow remote or mobile employees access to those same corporate resources on a tablet or smartphone.

Instead of just focusing on devices or even applications, Sunrise Health used a combination of Citrix's XenDesktop, AppSense's user profile management and Unidesk for application layering for its desktop management. It also used identity-based capabilities that were flexible enough for employees to access from anywhere and anytime.

Improvements in infrastructure reduced application system loads, cutting workstation login times from hours per day to as little as five to 10 minutes, Jayasinghe said. And because access was tied to each user's identity through a smart card reader, only authorized people could log on, protecting patient privacy and further ensuring the security of medical records.

Virtual desktop infrastructure, however, can be expensive and time-consuming for most organizations to establish, which is among the reasons adoption rates continue to hover around 10% to 15%, according to expert estimates.

Sunrise Health's efforts will cost $5 million over a five-year roadmap, but most of the return on investment is not quantifiable in hard dollars.

"Right now, we are fighting fires and being reactive instead of proactive," said Jayasinghe. "We anticipate that will be different once the rollout is complete."

A Future Without Endpoints

Endpoint management has begun to converge with identity management, so some IT shops are managing data, identity and access to services rather than a physical device.

That shift makes the actual management of an endpoint device—tablet, laptop, desktop, thin client, smartphone, whatever future device—irrelevant in the grand scheme of things, said Benjamin Robbins, principal at Palador Inc., a mobile consultancy in Seattle.

For example, an employee who travels frequently is rarely, if ever, accessing corporate resources behind

the company firewall. That worker logs in from remote locations on unsecure wireless networks provided by coffee shops, hotels and airports from various mobile computing devices. Yet, even under those circumstances, road warriors still expect access to corporate resources because organizations expect them to be productive, said Krish Kupathil, CEO at AgreeYa Mobility, a mobile software vendor in Mountain View, Calif.

If organizations are concerned only with managing the physical device, they won't be managing the stuff that enables secure mobile productivity, he said.

With mobile productivity increasing and potential data security threats at every turn, organizations can no longer focus on managing just physical devices, which was easy to do through BlackBerry Enterprise Server when BlackBerry smartphones were the only mobile game in town, Robbins said.

Now, it's better to approach management and security using several layers like application management, identity access controls and some light device management, he added.

Unlike traditional Windows PCs, mobile devices aren't capable of the same extensive management and security features yet, Little said. "They are definitely getting better but not quite up to the same level," he added.

As a result, organizations will probably need different tools to manage and secure the mobile stack for the time being. That might not be the case for long, however. The line between mobile endpoint management and traditional PC endpoint management is beginning to blur even for vendors. Companies that started out on the mobile side have expanded into PCs and vice versa.

"Because of all these new technologies, we think IT will be managing all of it from an identity and context point of view and layering in security on top of that," said Ben Goodman, VMware's end-user computing evangelist. It's the main reason the company, best known for its virtualization products, has pulled together physical, virtual and mobile management products into the just-released Horizon Suite.

Typically, companies take the tools and techniques from the previous era and apply them to the new—"even the first cars looked a lot like a horse and buggy," Robbins said. VMware, however, is using what it has learned in mobile management for other products, he said.

"Some of things we're learning from iOS app management, we are bringing them back into ThinApp," Goodman said. "App virtualization on Windows has until now been about the ease of install or compatibility or patching. It hasn't been a management or control technology, but we think it could be." ■

JAMES FURBUSH *is a news reporter covering consumerization. Email him at jfurbush@techtarget.com, or follow @JamesFurbush on Twitter.*

DATA CENTER ARCHITECTURE GETS A SHAKEUP

Changes are coming in data center network and server architecture.

Remember when data centers had separate racks, staffs and management tools for servers, storage, routers and other networking infrastructure? Those days seem like a fond memory now. That setup worked well when applications were relatively separate or they made use of local server resources such as RAM and disk and had few reasons to connect to the Internet.

But, with the intersection of intensive virtualization and the growth of Internet applications that touch on Web, database and other cloud-based services, things have gotten more complex. A wider range of computing services is available to satisfy these various applications, including more compute-intensive applications that involve big data and higher-volume and higher-density

HOME

server configurations. So, today there are many more options for data center server and network architecture—which trickles down to affect overall data center design.

There is no longer a single right way to build out your data center. With so many choices available, data centers have more flexibility. There are three primary schemes for how to deploy your servers, which also reflect three different network architectures.

Data Center Architecture Options

■ **The traditional data center.** These data centers employ purpose-built name-brand servers from tier-one vendors with storage area networks (SANs) running Windows or Linux. Today, the vast majority of data centers feature this as the gear, with roots that trace back to the early days of the client/server era. The key advantage of such an architecture is that it's well known and IT staffers have extensive experience with it. In terms of networking, traditional servers are connected via commodity network switches running 10 Gigabit Ethernet or InfiniBand. Servers and switches use proprietary management software but can be easily upgraded or swapped out for other vendors' gear. Each equipment rack has its own network switch, and these switches are connected to a backbone core switch.

HP's ProLiant DL2000 Multi Node Server

■ **Hyperscale servers.** The hyperscale server label refers to new kinds of servers that are customized for particular data center needs—even the racks are wider than the traditional 19-inch mounts that were the standard. This accommodates more components across the motherboards of these PCs. They are also assembled from common components that can be easily swapped out when failures occur. Think of having a single server with a dozen hard drives on a motherboard with multiple power suppliers for redundant operations.

This is the architecture of choice in 100% cloud-based businesses such as Facebook, Amazon and Google, and also for building a new breed of supercomputers. Typically, these servers run some form of Linux and are now sold by both traditional server vendors such as Hewlett-Packard, with its ProLiant DL2000 Multi Node Server, along with components available from several suppliers.

Hyperscale server networks can use some traditional top-of-rack network switches, but

Facebook's networking standards are part of the Open Compute Project and call for New Photonic Connectors and embedded optical modules. Both Facebook and Google are building new data centers in Iowa using similar designs. The advantage of these new server and network designs is that you can reduce power losses with direct current (DC)-powered drives and save time on troubleshooting failed components because every server is uniform. You can also scale up capacity in smaller increments.

■ **Virtual or converged clusters.** These systems use proprietary server hardware that isn't interchangeable and only upgradable in large increments of storage or compute power. The servers have very high internal memory usage of hundreds of gigabytes of RAM. Cisco's Unified Computing System, Dell's Active Infrastructure and IBM's PureSystems are examples of converged infrastructures.

Virtual or converged clusters use integrated network switches, computing and storage blade servers and specially made chassis to assemble all components. The model typically runs over multiple 10 GbE connections, and the advantage is that you can eliminate many steps to provision and bring online new capacity and quickly connect to your infrastructure. The disadvantage is mostly cost.

In the future, the chances are that your data center will combine two or three of these approaches. "There are brand-new workloads that we have never seen before," said Andrew Butler, an analyst at Gartner Inc. "Amazon and Google's requirements aren't like banks or traditional IT customers. They want to run these workloads with large in-memory databases or over large-scale clusters, and these don't fit the traditional data center architectures."

This new breed of cloud apps is even more challenging, because it changes the focus from raw computing power to a more efficient use of electric power. "Power is everything these days," said the CTO of a major cloud services hosting provider. "We design our data centers for the lowest [power usage effectiveness, or PUE] ratings possible." Hyperscale servers use entirely new power distribution methods for the maximum power efficiency possible. For example, hyperscale servers use 480 volt (V) DC power supplies and 48 V DC battery backups, along with direct evaporative cooling, meaning that no chillers or compressors are needed for cooling these data centers. Organizations that have deployed these kinds of hyperscale architectures include Rackspace and several Wall Street firms, relying on 1U to 3U dual-socketed AMD-based servers with up to 96 GB of RAM and direct-attached hard drives.

Part of the challenge of the hyperscale approach is real estate. "If you can build facilities where real estate costs are low, and if you are running monolithic business applications or cloud services, these low-density data centers make sense," said Jack Pouchet, the VP of business development for Emerson Network Power. "However, when it comes to HPC [high-performance computing],

big data, Hadoop and other massive number-crunching analytics, we likely will see the continued push toward higher-density architectures."

Keeping Tradition Alive

For smaller data centers or more modest needs, traditional architecture still makes the most sense. It is familiar and the least costly to purchase. "But you really need to understand your applications' consumption model and requirements," said Ken Owens, cloud CTO for Savvis, a large international cloud services provider that operates multiple data centers. "If you are in a more traditional environment and have typical business applications that require limited network services, you probably want to stick with traditional Intel services, maybe adding a blade server to simplify your provisioning."

It is easy to find the traditionalists. Take EasyStreet, a Portland, Ore.-based data center provider. It uses traditional server architecture completely, buying standard Dell and HP servers to host mostly small local businesses. "We don't have any blades now and don't anticipate buying any in the future either," said Rich Bader, EasyStreet's CEO. "Even our customers that are running 50 or 100 virtual machines still don't really warrant using blades. Blades carry a cost premium that doesn't make sense for us, and the economics just don't work."

Tony Stirk is the president of Iron Horse, which operates another data center for smaller customers in Springfield, Va. He has stuck with traditional architecture as well. "Blades have a large number of proprietary

What about the network?

WHILE THE Open Compute Project has lots of information on server and "open rack" designs for an all-DC power, white-box collection of computers, it is noticeably silent on how these computers are going to be connected together.

In a blog post in May, the organization stated that "we're still connecting [these new computer designs] to the outside world using black-box switches that haven't been designed for deployment at scale and don't allow consumers to modify or replace the software that runs on them."

To that end, the Open Compute Project is "developing a specification and a reference box for an open, OS-agnostic top-of-rack switch. Najam Ahmad, who runs the network engineering team at Facebook, has volunteered to lead the project." A wide variety of organizations will join in, including Big Switch Networks, Broadcom, Cumulus Networks, Intel and VMware. That is quite a lineup.

But it is also quite a challenge. The new network switches will have to run on DC power, like the rest of the gear sitting below them. They will need to support software-defined networking, which is itself an evolving collection of standards. And it has to fit a wide variety of workloads and use cases too, all while being vendor-neutral.

parts which limit competition and make your business more dependent on a specific manufacturer," Stirk said, "You have limited configuration, networking and storage choices."

Starting New Traditions

But as your data center grows, traditional architectures may increase operating costs, either for power efficiency or for supporting higher computing densities. Take Contegix, a St. Louis-based data center provider. "We are deploying converged infrastructures for both our private and public clouds," said Matt Porter, Contegix CEO. "Both of these scenarios are aligned to the slicing of large pools of computing resources for small, specific needs." He plans to continue using both traditional and converged infrastructures as appropriate for customers.

This increase in scale is one reason that software-defined networks are becoming more essential, because existing virtual local area networks, or LANs, can't grow quickly. The same could be said about SANs, too.

For IT professionals making the transition from traditional to other architectures, it could be just a matter of timing which they choose. Newer architectures are still considered as replacements for traditional ones. Gartner's Butler agrees. "We are still at a relatively early stage for multinode [a.k.a. hyperscale] servers. Their growth has been meteoric and they are already outselling blades in terms of installed numbers, but not in terms of actual dollars yet. There is a huge appetite for these kinds of servers, and they are being included in many mainstream data centers."

And with its emphasis on the cloud and with the consumerization of IT, Savvis has changed the gear it is buying. "We have had to redesign our infrastructure to be configured on demand and be more flexible," Owens said. "Customers no longer know what they need or how they want the infrastructure to support their apps, so at Savvis we are mixing blades and SANs with multinode servers to better meet our customers' needs and provide more intelligence at the app level to make our data centers more flexible." That said, Owens predicts that multinode racks could account for 80% of Savvis' floor space in the near future, which is a big change from roughly that percentage of traditional and virtualized servers today.

This mixture of approaches is probably what the future portends, said Philip Molter, the CTO of Data Foundry, a major data center service provider in Austin, Texas. "If Google or Facebook had to answer to an angry customer every time an instance went down because of a 0.05% increase in failure rate, they might choose a different deployment strategy." ■

DAVID STROM *is an expert on network and Internet technologies and has written and spoken on topics such as VoiP, network management, wireless and Web services for more than 25 years. He has had several editorial management positions for both print and online properties in the enthusiast, gaming, IT, network, channel and electronics industries.*

ARE WE THERE YET?
JONATHAN EUNICE

The Road to Oz

HOME

Remember when Dorothy and friends finally—finally!—made it to the Emerald City? It was opulent, but the Great and Powerful Oz wasn't quite as expected.

In today's IT, The Almighty Cloud is our Oz. We venture toward it, full of hope that it will magically solve all known IT problems. It won't, of course. But also remember: However untrue Oz's hype, the wizard ultimately helped Dorothy and friends. They got where they needed to go, if not how they initially expected. Cloud computing is the same. It's transformative, but first we must pull back the curtains.

There Is No Cloud

The cloud is not a real thing. It's an abstraction for flexible, network-savvy, on-demand services. There are entirely real and successful implementations, but in essence it's a style of operation. And unlike the Internet, there is no single unified cloud. Not every endpoint can meaningfully communicate with every other endpoint.

There are *many* cloud service providers—Amazon Web Services, Cloud Foundry, Engine Yard, Force.com, GoGrid, Google, Heroku, Hewlett-Packard, IBM, Joyent, Microsoft, Rackspace, Savvis, and Verizon Terremark, to name just a few infrastructure and platform options. They offer varying application programming interfaces and abilities. Some specifications, such as OpenStack and vCloud, aim for standardization and interoperability, but it's very early days.

THERE IS NO SINGLE UNIFIED CLOUD. NOT EVERY ENDPOINT CAN MEANINGFULLY COMMUNICATE WITH EVERY OTHER ENDPOINT.

Meanwhile, public clouds hosted by service providers, available for a fee to all comers, are sometimes portrayed as the one true way. Enterprise-hosted implementations (private clouds) are thus false. Cloudistas spurn virtual

infrastructure even more.

The "true cloud" mythos is great attack marketing for the service providers, but for customers, outcomes outweigh labels. External services have advantages, including minimal up-front expenses, great "trialability," economies of scale and the benefit of letting someone else worry about this stuff. Some uses, such as content distribution and off-site backup, make perfect external services.

External services have drawbacks, however. They include the requirement to always send data over the network, the higher latency of faraway resources, a variable pricing structure, and genuine security, auditability and regulatory compliance concerns.

As a developer, I love public clouds. But those in regulated industries, with latency- and bandwidth-sensitive apps, where even temporary service unavailability is unacceptable, or who are already comfortable managing Capex-Opex tradeoffs, see things differently. For them, cloud needs to be private, or a public-private hybrid.

Cloud Elasticity Is Overstated

Software as a Service and Platform as a Service offerings are genuinely elastic. Buy however much you want. Infrastructure as a Service (IaaS), though, *isn't* very elastic—even for services such as Elastic Computer Cloud (EC2) where "elastic" is part of the name.

You can buy as many EC2 servers as you like, fire them up, and you're off! You can later shut them down, lowering or eliminating ongoing charges. That's all great. But EC2 is a utility service, with individual servers sold only in fixed configurations. Once chosen, resources never flex over an instance's lifetime.

VIRTUAL INFRASTRUCTURE IS THEREFORE MORE ELASTIC THAN CORRESPONDING CLOUD SERVICES, IN IMPORTANT AND PRACTICAL WAYS.

I've made the mistake of describing EC2 as elastic to clients. It's then very hard to explain why the instances are so inflexible, and that they'll need load balancers, scale-out clustering and the like to support a truly elastic service. Clients eventually get it, but they're disappointed.

It's even worse when clients already have virtual infrastructure. VMware, Hyper-V, XenServer, and IBM PowerVM all allow precise CPU, memory and I/O configurations. Those resources can be easily added or removed—sometimes even while applications are running. Virtualization platforms share resources across instances and can even live-migrate workloads. EC2 and most Infrastructure as a Service competitors do none of those things. Virtual infrastructure is therefore *more elastic* than corresponding cloud services, in important and practical ways.

Over the Rainbow

To be clear: Like so many others, I love cloud computing, with good reason. The option to externally host infrastructure, services and apps is great for developers, many operators and the industry as a whole. At a higher level, cloud describes everyone's aspiration for more a flexible, efficient, dynamic IT, and it usefully benchmarks all deployments.

WHEN YOU'RE TRAVELING TO ELASTIC CITY, TAKE IN THE WONDROUS SIGHTS—BUT PAY GREAT ATTENTION TO THAT MAN BEHIND THE CURTAIN.

But all that "Cloud solves everything" rhetoric helps no one. So when you're traveling to Elastic City, take in the wondrous sights—but pay great attention to that man behind the curtain! ■

JONATHAN EUNICE *is a principal IT adviser at the analyst firm Illuminata Inc.*

TechTarget

Modern Infrastructure is a SearchDataCenter.com publication.

Margie Semilof, *Editorial Director*

Alex Barrett, *Editor in Chief*

Lauren Horwitz, *Executive Editor*

Christine Cignoli, *Senior Site Editor*

Phil Sweeney, *Managing Editor*

Eugene Demaitre, *Associate Managing Editor*

Laura Aberle, *Associate Features Editor*

Linda Koury, *Director of Online Design*

Rebecca Kitchens, *Publisher*, rkitchens@techtarget.com

TechTarget, 275 Grove Street, Newton, MA 02466
www.techtarget.com

© 2013 TechTarget Inc. No part of this publication may be transmitted or reproduced in any form or by any means without written permission from the publisher. TechTarget reprints are available through The YGS Group.

About TechTarget: TechTarget publishes media for information technology professionals. More than 100 focused websites enable quick access to a deep store of news, advice and analysis about the technologies, products and processes crucial to your job. Our live and virtual events give you direct access to independent expert commentary and advice. At IT Knowledge Exchange, our social community, you can get advice and share solutions with peers and experts.

COVER PHOTOGRAPH: RIHARDZZ/ISTOCKPHOTO