
Hewlett Packard
Enterprise
enterprise.nxt
Hybrid IT
Operations
An expert guide
EDITED BY DAVID CHERNICOFF & RICHARD MCGILL MURPHY

TABLE OF CONTENTS: HYBRID IT OPERATIONS

Hybrid IT Operations

Introduction

By David Chernicoff and Richard McGill Murphy

Introduction

One significant advantage of implementing a hybrid IT solution is that integration with central IT brings the core strengths of legacy IT, including operations management, security, backup, and disaster recovery tools and processes.

In many ways, developing and deploying a hybrid IT environment is like assembling a jigsaw puzzle, often without a clear picture of what the final implementation will look like. The ease of deploying cloud components and implementing DevOps-driven solutions without going through central IT can result in services and solutions that are poorly aligned. While they may solve immediate problems, they are unlikely to scale well or fit together to drive future business growth.

Hybrid IT is emerging as the new normal for enterprise IT because it allows enterprises to choose the blend of apps, services, and platforms that work for their needs. These technologies are widely distributed, constantly changing, and delivered at massive scale. By extending the operational rigor of legacy IT from the data center to the hybrid cloud, businesses can build a strong backbone for the delivery of business services.

This report covers the basics of implementing and running an enterprise hybrid IT environment that extends the benefits of legacy IT to hybrid and cloud deployments, from delivering a secure environment, to budgetary considerations, to managing effectively across multiple platforms.

01

Keeping your hybrid IT safe and secure

By Paul Ferrill

- It's not traditional IT
- Evaluating your environment
- Lock it up
- Find the right tools

Securing your hybrid IT infrastructure

Is your organization launching a hybrid IT project? If so, you might want to take a good long look at the policies and tools you're currently using, as they might be insecure.

Key pieces worth reviewing include identity and how you both assign and monitor privileges; boundaries between systems under your control and those outside your purview; and policies and procedures written with basic assumptions that might not be valid.

Security risks rank as one of the highest concerns for IT administrators when implementing a hybrid IT environment. While old security tools and techniques don't necessarily need to be replaced, they may need augmentation through the use of additional tools to address the new landscape. Identity remains a key component for controlling access to corporate resources, but it must be expanded across all platforms.

> The key is security in depth: You put a lock on the door (firewall), but you also put an alarm on that same door (IPS) and security cameras throughout the building (IDS), monitored by human beings.

Regular backups are one of the best things any organization can do to protect itself from attacks. Backups should be tested on a regular basis to ensure both integrity and validity of the data. Many ransomware attacks have been thwarted with a recent backup of import databases. For mission-critical business data, these backups, or snapshots, should be accomplished at a frequency dictated by the volatility of the data.

Traditional security IT

The world of security is a constantly changing landscape in which an agile approach is required to stay ahead of the threats. According to the chief information security officer at a large healthcare organization who spoke on condition of anonymity, "While we still rely on things like firewalls and VPNs, we do use some of the newer intrusion detection system [IDS] and intrusion protection system [IPS] products for monitoring and alerting."

The key is security in depth: You put a lock on the door (firewall), but you also put an alarm on that same door (IPS) and security cameras throughout the building (IDS), monitored by human beings. The analogy between physical and information security is a bit of a stretch but does bring out the idea of using both tools and people for the best coverage. Automated tools can take you a long way, but having a human in the loop still adds a capability that computer systems can't quite match.

Microsoft's Active Directory (AD) is the authoritative identity source for most, if not all, IT organizations of any size. The challenge is how to leverage that resource in a way that makes sense and addresses any potential issues from both a security and performance perspective. Many tools exist to help with the monitoring and managing of your AD infrastructure.

Both Microsoft and VMware have recently incorporated encryption for their virtual machine disks and the transport used to move or migrate VMs from one host to another. The encrypted disks require a key stored in a repository. The key is provided only when the system connects and authenticates through the use of a certificate. This prevents any attacker with physical access from simply copying a virtual disk.

Assessing the landscape

Looking to shore up your security posture? First take stock of the current condition of the components, starting with administrator privileges and the policies and procedures surrounding the granting and monitoring of admin rights. While this might require new tools, it also implies a need to understand the moving parts and limitations of each. If your company integrates with third-party providers, you must identify any potential risks that require mitigation.

"A number of identity issues must be considered in the hybrid IT scenario," says Edward Haletky, a principal analyst at TVP Strategy covering cloud, security, and DevOps, among other topics. "First, as an IT organization you still need control of the process, and that becomes the problem. Getting authenticated is not the real issue, but who has access. Role-based access control is limited in many cloud services, with some offering only admin or not—with no in between."

Conducting a comprehensive security audit should be one of the first tasks on your list. This includes evaluating firewall policies and any existing applications that require a firewall rule. You should examine all users and existing privileges to ensure no administrative rights were given to someone who didn't need them. Local administrator rights on servers need to be reviewed to determine if any users or service accounts should be removed.

Ask any IT administrator about the effects of so-called shadow IT on their security and you're bound to get a few horror stories. In today's environment, this extends to unauthorized usage of cloud apps as well. A comprehensive monitoring or logging tool can help identify shadow IT situations and bring them under corporate control. While many of them might be harmless, they can introduce vulnerabilities if left unchecked.

Battening down the hatches

Using existing tools and capabilities smartly is still a good approach. "Our primary method for allowing external vendors access to our network centers around known IP addresses and SAML," says the healthcare CISO. Security Assertion Markup Language, SAML 2.0, is supported by all the big cloud providers, including Amazon Web Services, Google's cloud services, and Microsoft's Azure.

Microsoft has invested a great deal in its Azure platform and extending AD to support all of its cloud-based services. Azure AD currently handles upwards of 1.3 billion authentications per day. For existing Microsoft customers looking to expand their AD services to the cloud or other sites, Azure AD Connect is a good option. AD Federation Services provides the mechanism to configure a hybrid environment connecting an on-premises infrastructure with a number of complex scenarios.

For systems running the Windows Server operating system, it is possible to lock down many potential security risks using group policy objects (GPOs). Common threat scenarios here include privilege escalation, where an attacker attempts to gain access to sensitive information by gaining increased privileges with compromised credentials. GPOs allow an administrator to disable features like LDAP's simple bind, which permits unencrypted passwords.

> A comprehensive monitoring or logging tool can help identify shadow IT situations and bring them under corporate control. While many of them might be harmless, they can introduce vulnerabilities if left unchecked.

Microsoft's Advanced Threat Analytics product can detect plain-text passwords passed in an unsecure way using LDAP. This same product will scan for credential exposures through service accounts as well. "We have recently seen a rash of compromised databases based on NoSQL products like MongoDB," says Craig Young, principal security researcher at Tripwire, a provider of advanced security and compliance tools. "The issue centers around default installations, which are inherently insecure. These systems must be configured to explicitly deny remote access and change any default passwords." Young also recommends a tiered or layered security approach that includes both monitoring and active scanning.

Tools for the job

Gartner defines cloud access security brokers (CASBs) as "on-premises, or cloud-based security policy enforcement points, placed between cloud service consumers and cloud service providers to combine and interject enterprise security policies as the cloud-based resources are accessed." A good number of vendors exist in this space. Many deliver software as a service that mediates between corporate applications and cloud services.

Tripwire, for example, offers a security information and event management tool that establishes a baseline of normal activity and then monitors the system for large deviations from that norm. If the activity monitor detects abnormal behavior, it generates an alert to a designated recipient list. Other tools, such as universal threat monitoring appliances, can provide detailed inspection of network traffic to look for malicious behavior.

Encryption is built into many storage systems, as well as operating systems such as Windows Server. Microsoft's BitLocker technology encrypts at the drive level. This covers the issue of data at rest and protects against theft of a physical device. At the network level, most companies use some type of virtual private network (VPN) to protect traffic between partners and corporate assets such as laptops and the internal network. Microsoft has addressed the complexity of VPNs in later versions of its operating system with tools like DirectAccess and a feature called Work Folders. Both use certificate-based authentication to connect corporate assets back to internal resources.

Network and application monitoring tools also help with the monitoring and management of networked resources. Several companies sell both single and multifunction monitoring tools that provide reporting and agent-based monitoring to gain insight into local and remote resources. These tools have the ability to monitor both on-premises systems and cloud-based services.

Final thoughts

Hybrid IT security has many challenges, some of which require new policies and tools to keep the system safe. It also requires IT staff to learn new methods while maintaining legacy components, many of which will continue to operate as they always have. The key here is to know which pieces to keep and which ones to either remove or augment with something new. Finding vendors you can trust is key and must be a priority for any IT manager. ■

01 NOW TO NEXT

Hybrid IT and security: Lessons for leaders

Up-to-date security technologies that extend from the legacy environment are critical.

Using the public cloud means exposure to the most current threats. Security teams need to keep that firmly in mind and update and prepare accordingly.

Don’t be afraid to use security technologies built into your physical and virtual hosts.

02

Delivering services where and when you need them

By Andy DeBernardis

- Virtual machines can be a big help
- On-demand services change the way IT responds
- It's not just about speed
- Rapid provisioning and services are not a panacea

Workload placement

Should you rent, lease, or buy your next residence? To make a decision, you weigh a multitude of variables, talk to people you trust, and then make a decision that works for your personal situation.

Similar logic applies to workload placement decisions in a Hybrid IT environment. There are many deployment choices to consider and no one-size-fits-all approach, because different workloads have different characteristics and each business is unique. You need to weigh the key decision factors, compare different options, and then match the workload with the appropriate destination based on specific business objectives.

In a recent Frost & Sullivan survey, 51 percent of IT decision-makers cited "assessing the optimal deployment model for workloads" as a difficult challenge in implementing their Hybrid IT environment. It's complicated, and the landscape changes constantly. Not too long ago, traditional IT was the default platform. Here, the application, service, and workload were all nicely tied together. The terrain was familiar and mostly contained within your data center.

Workload placement patterns are shifting. According to a recent 451 Research study, 60 percent of workloads still reside in traditional or non-cloud infrastructure, compared with 40 percent for cloud workload deployments. The authors predict that "this mix will change dramatically in two years, essentially reversing the workload figures to 58 percent of workloads deployed in the cloud and 42 percent deployed on non-cloud infrastructure."

Keep in mind that not all clouds are public. Cloud is about delivering the agility and speed the business expects. If you adopt a Hybrid IT strategy, your environment will consist of different deployment environments, including traditional IT, on-prem private cloud, managed cloud, software as a service (SaaS), and public cloud. Each has pros and cons. It's a good idea to consider the full range of deployment options for each workload placement decision you make.

> According to a recent 451 Research study, 60 percent of workloads still reside in traditional or non-cloud infrastructure, compared with 40 percent for cloud workload deployments.

Choice, choices

Traditional IT consists of computing resources that reside within your data center and are configured to meet your particular needs.

Like traditional IT, private clouds are customized for each customer deployment and are part of your data center. Using automation and orchestration software integrated with hardware, developers and end users can access the services they need through a self-service catalog. Once a service is selected, resources are automatically and rapidly provisioned to meet workload and application development needs.

With public cloud, customers consume shared, scalable services that are delivered over the Internet.

Managed cloud services are a hosted model that provides the benefits of the cloud with access to services and tools to support the unique needs of each business. Unlike public cloud, managed cloud services are managed by a third-party resource.

SaaS is a model designed to support the needs of multiple customers. It includes shared, hosted software services like Salesforce and Office 365, delivered over the web and licensed by end users.

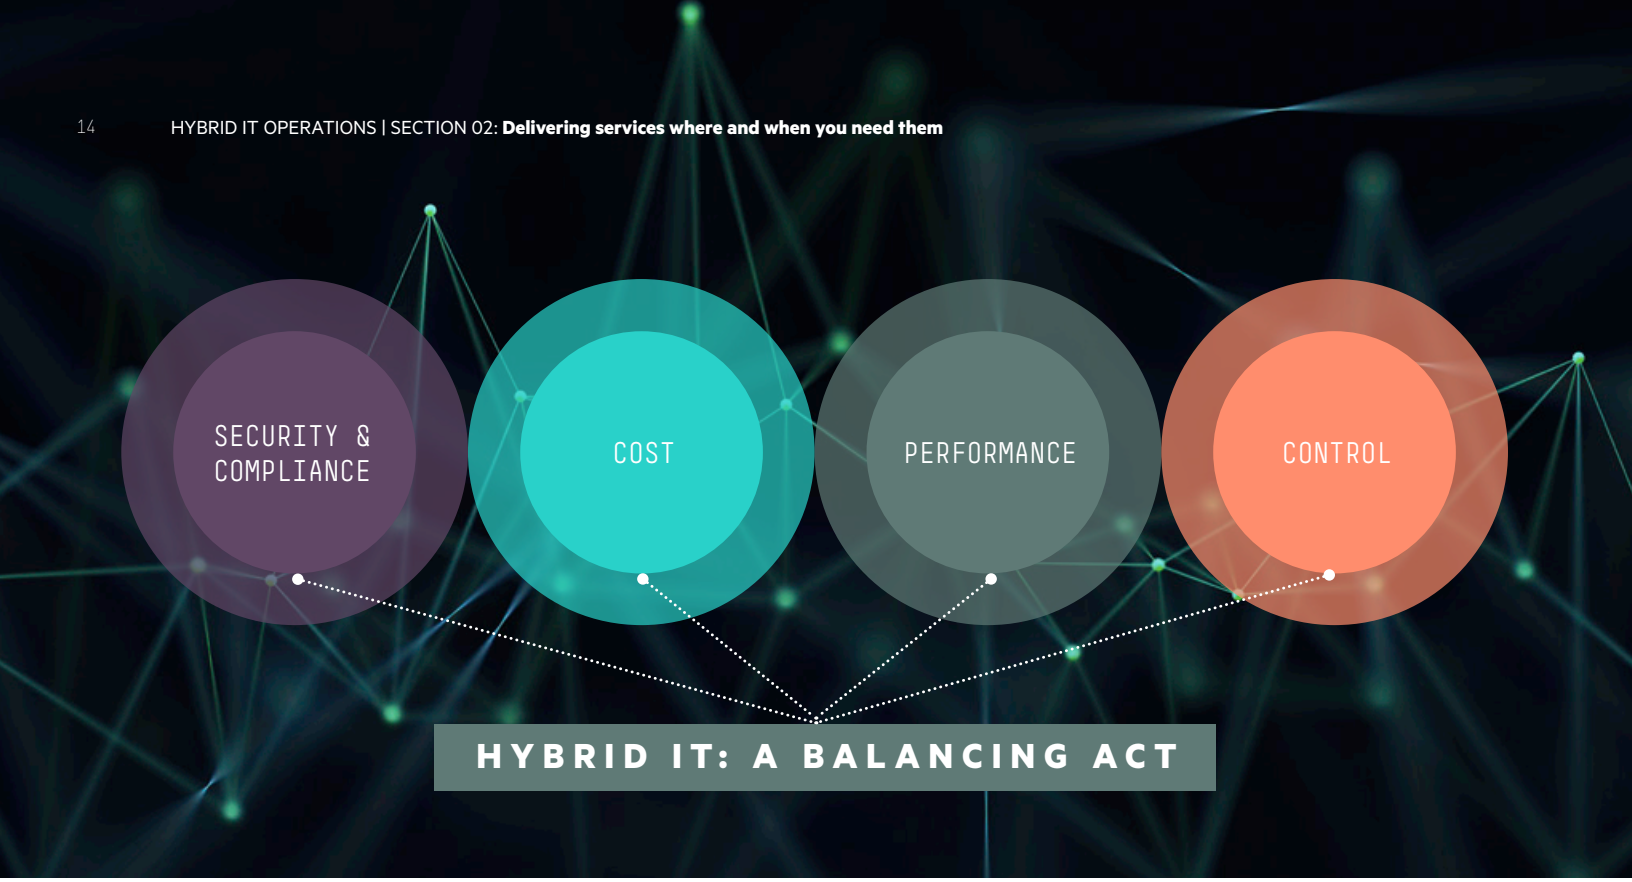

Success factors

Hybrid IT typically involves a mix of environments and many variables. You'll need to consider four main variables when making a workload placement decision.

1. **Hybrid IT security and compliance.** Each day brings new exposure and attack surfaces, increasingly sophisticated cyberattacks, and complex regulatory pressures. When considering off-prem environments, you need to know exactly what your vendor will do if there is a problem and understand up front what it would take to bring the workload back on-premises if needed.

   How much risk can you tolerate? What will it take to keep your enterprise secure and compliant? These are complex calculations that vary by geography and company. In the U.S., for example, you might determine that a certain workload is appropriate to run on a public cloud. In Europe, you might opt to run the same workload in an on-prem cloud environment, due to European data sovereignty rules and data sensitivity concerns in your business.

2. **Total cost of ownership.** Comparing TCO across platforms can feel like comparing apples and puppies. It's fairly easy to calculate TCO for a group of dedicated applications across a given set of servers, switches, and storage. Even then, it can be hard to compute actual cost per application. Cost for public cloud is often calculated on a price per minute for each virtual machine (VM)/minute, but those prices can be volatile. With managed cloud, the price is calculated on a VM/hour basis, with a set price for each contract. And there are many cost factors to consider with private cloud. They include hardware, software, data center costs, business continuity needs, and personnel costs, among others.

Don't assume that public cloud is always the cheapest option. Let's say you plan to build an enterprise-grade private cloud or invest in software-defined infrastructure to speed application delivery. When you go to present your investment business case, be prepared to demonstrate that the TCO is either less than, equal to, or greater than the public cloud. And don't focus on cost alone. Figure out the price for each option, and be ready to differentiate them in business value terms. Take the time to define an apples-to-apples comparison that you can use to inform your decision.

3. **Application performance.** According to Frost & Sullivan, "60 percent of IT decision-makers cite concern about application performance as a reason not to move a workload to public cloud." Public cloud is accessed over the Internet, and the infrastructure is shared. Virtual machines compete for the same resources, introducing latency potential. But who determines resource prioritization in a shared environment? Consider whether your workload requires low-latency performance. There are workloads, such as high-frequency trading in the financial services industry, where application performance needs to support high-speed transactions and delays are unacceptable. Determine what you consider an acceptable delay and whether your workload can tolerate inconsistent latency. Depending on your application performance needs, a private cloud may deliver the performance you ultimately require.

4. **Control is a factor that cuts across almost all of the other variables.** The right level of control varies with company culture and policies. IT seeks to stay relevant by providing business users what they need, when they need it, at the right economics. You need to provision the right resources quickly to match the needs of each workload. You need to manage the workload from end to end, and you may need to manage it centrally. According to Frost & Sullivan, 80 percent of IT decision-makers cite "loss of control over application" as a reason to avoid deploying a workload in the public cloud.

Which of your workloads distinguishes you from the competition? For these workloads, how much control do you or your line-of-business counterparts feel you need? Would you be comfortable running them with a third party? Just because you can put a workload in the public cloud doesn't mean you will or should. Different workloads have different needs. 451 Research shows both current and projected deployment trends broken down by workload type, including email and collaborative, web and media, data and analytics, application development, business applications, and shared IT workloads. I encourage you to take a look at this paper if you're interested in the specific trends occurring by workload type or want to learn more about how to make an informed decision.

Also, data center technology has evolved rapidly in recent years. There are ways to get the speed and agility of public cloud for a variety of workloads using traditional IT. Software-defined options such as composable and hyperconverged infrastructure with cloud management software can rapidly deliver traditional and cloud-native applications. You should also consider on-demand services that combine the agility and economics of public cloud with the security and performance of on-prem IT. ■

02 NOW TO NEXT

Where to place your workloads: Lessons for leaders

Just because you can place your workload on the cloud doesn't mean you should.

Carefully consider regulatory issues that can impact your workload placement.

Comparing TCO between the various solutions can be very difficult. Make sure your comparison points are accurate.

03

Cross-platform management and monitoring

By Amy Newman

- It's all about the process
- Many challenges, multiple solutions
- Finding the right answers
- Evaluating cost versus results

How holistic management benefits your hybrid IT environment

The IT landscape has changed dramatically in the past decade. Today's modern enterprises are acutely feeling the transformation, as businesses gravitate toward software as a service (SaaS) instead of packaged software, and cloud computing continues to hit numerous metrics that reflect its growing maturity.

Cloud by the numbers

Cloud use, both public and private, is pervasive. A survey conducted by CompTIA found that 90 percent of respondents have migrated at least some of their infrastructure to the cloud, and 71 percent are using the cloud in full production or for noncritical uses. The remaining respondents are still experimenting with cloud or have transformed their IT organization with cloud.

Far from having peaked, the cloud market continues to exhibit growth. Data from the Synergy Research Group reveals the cloud services and infrastructure market grew 25 percent on an annualized basis in late 2016, reaching $148 billion. Infrastructure as a service (IaaS) and platform as a service (PaaS) had the highest growth rate at 53 percent, followed by hosted private cloud infrastructure services at 35 percent and enterprise SaaS at 34 percent.

The Synergy survey notes that 2016 represented a sea change in terms of spending: Cloud services spending overtook that of cloud infrastructure hardware and software. In aggregate, the cloud services markets are now growing three times more quickly than cloud infrastructure hardware and software.

75%

of IoT adopters will turn to outside firms for help in strategy, planning, development, implementation, and/or management by 2018.

Source: IDC.

However, despite this growth and an increase in enterprises taking a cloud-first philosophy, few organizations will rely on public cloud 100 percent. With the exception of startups, organizations rarely undertake greenfield IT initiatives. Long-standing investments in legacy infrastructure and storage keep some applications and corresponding data housed in the data center, and regulatory and other compliance requirements restrict some applications and data to remaining on-premises.

Some organizations will address this with private cloud, and indeed, the Synergy study notes that despite the fact that public cloud spending is growing much more rapidly, private cloud spend accounted for more than half of total cloud spend.

CompTIA reported a similar preference for private cloud: Nearly half of survey respondents said they have invested in only private cloud. However, hybrid cloud was also popular and deployed by more than one-quarter of respondents. In addition, nearly three-quarters of respondents have SaaS applications, and IaaS and PaaS are also popular choices.

Addressing the challenges of hybrid IT

This leaves a growing number of enterprises with a hybrid mix of on-prem hardware, applications, cloud-based solutions, and data located both in the data center and cloud. Such environments present a unique set of challenges. Chief among them is managing these disparate components as one cohesive infrastructure.

Managing resources across those clouds and services should be a seamless exercise. To the end user, there should be no difference in accessing on-prem or cloud-based resources. In reality, a hybrid infrastructure introduces a host of management and monitoring complexities and, with that, potential security issues. Compatibility, integration, and control also have the potential to introduce risk and impact the user experience.

Some enterprises opt to cobble together a mixture of siloed tools and build their own management solution. However, this is a Band-Aid approach that will only get increasingly complicated and more difficult to maintain with time.

Others turn to data center infrastructure management tools (DCIM). However, while some DCIM tools extend some degree of visibility to the cloud, they are primarily designed for data center monitoring, and often come up short when SaaS and cloud environments are added to the mix. These tools are available as software solutions from vendors such as Nlyte and hardware vendors

such as Schneider Electric, which offers StruxureWare management software. These solutions are tailored to manage data center facilities and can be optimized for specific vertical market users, such as healthcare providers.

Fortunately, tools are emerging to help organizations more easily integrate and manage the entire IT environment from a single solution. Such solutions provide a holistic view, often through a single pain of glass that offers visibility into activity across your entire IT infrastructure. These solutions range from Hewlett Packard Enterprise's Helion, which looks to manage "Any Cloud, Any Environment, Any Technology, for Any Application," to more specialized products like Red Hat CloudForms, which offers support for multiple virtualization platforms, and cloud providers or the Dell EMC suite of cloud management solutions.

Deploying the right tool provides insights that enable your organization to improve automation, orchestration, and security, leading to improved performance. But knowing what tool is optimal for your organization requires a deep understanding of your infrastructure.

Thus, before considering monitoring and management solutions, take some time to consider the role of these capabilities in your cloud environment and the impact changes will have on any intricacies and dependencies.

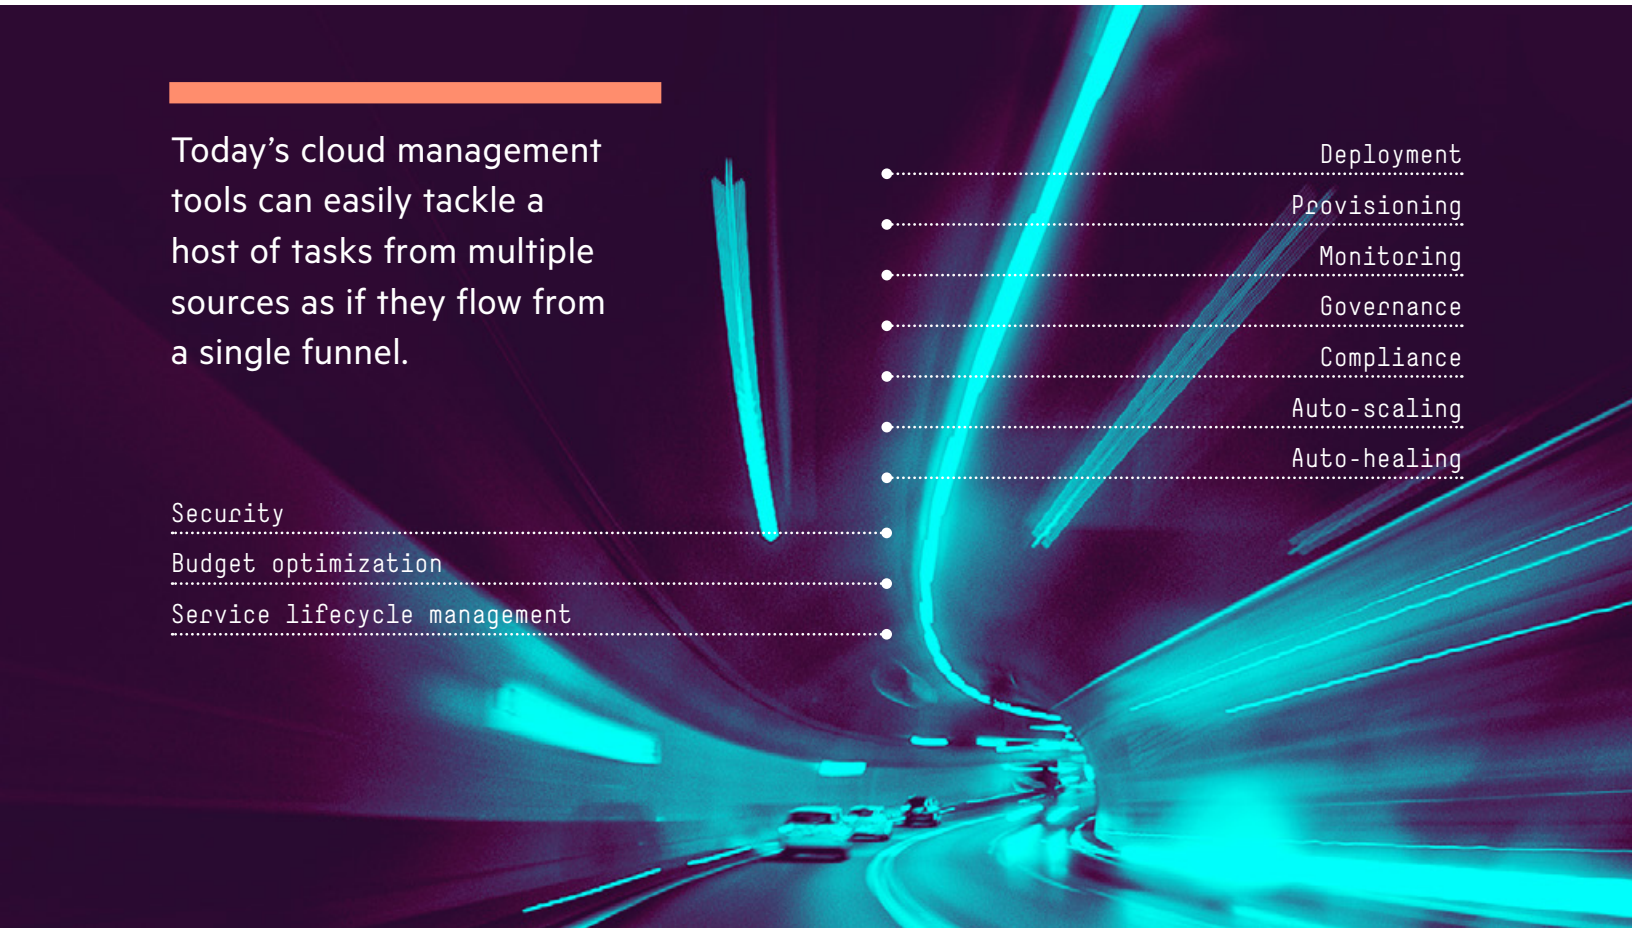

What to look for in a solution

Cloud management and monitoring tools have grown in sophistication alongside the cloud. Today's management tools can easily tackle a host of tasks, including deployment and provisioning of resources, monitoring, governance, compliance, auto-scaling, and auto-healing. Many include security and budget optimization features, as well as service lifecycle management.

Being able to perform these tasks from a unified tool with a single console goes a long way toward reducing complexity and simplifying the environment, making it easy to streamline actions like cloud bursting or disaster recovery.

> The ability to manage applications and services from a single console reduces complexity, streamlining operations management across the entire infrastructure.

Monitoring this complex infrastructure from a single console allows it to be managed as a single entity. A cohesive view results in a more holistic perspective and reduces the time that must be allocated to administration and troubleshooting tasks. This ultimately leads to improvements in integration, performance, usability, security, and cost.

Integration: The beauty of a single solution is that all data and metrics are derived as though they flow from a single funnel. For this to work, a great deal of integration must go on under the hood. Data is gathered from multiple sources that follow different standards and protocols. The data must be integrated from different vendors and locations so it is accessible as though from a single source. This will facilitate a set of consistent metrics across all platforms and provide a clearer picture for setting automation and governance policies.

Performance impact: Improved performance is a key selling point for a unified monitoring and management tool. After the tool has pulled and integrated data, it should be able to analyze log data, correlate events, and alert staff about incidents that need attention. The tool should offer a variety of methods to warn of impending problems as well as issue alerts when problems occur.

Policy-driven automation tools should also be part of its functionality, with governance policies that drive this automation carefully crafted and customized. Consider devising a variety of monitoring profiles, each with its own set of priorities and privileges to differentiate access to data and applications. Automation driven by such policies reduces the chance of errors and frees up staff to focus on other activities.

Usability: The ability to manage applications and services from a single console reduces complexity, streamlining operations management across the entire infrastructure. Automating IT tasks and processes further enhances usability by reducing errors, downtime, and risk of noncompliance. This ease of use enables IT to focus on the entire infrastructure and make changes proactively to prevent errors or downtime.

Security: Security is a top concern for enterprises regardless of where data and applications reside. While some organizations keep sensitive information on-prem or in a private cloud for compliance reasons or to mitigate cloud security concerns, this alone is not enough, and it does not eliminate the need for good security practices. Adequate security measures must be implemented to ensure data is secure regardless of location. With a unified management tool, consistent and location-agnostic access control policies based on data and application sensitivity can be established.

In addition, management tools that provide automation and ensure adherence to governance policies go a long way toward ensuring security by reducing potential for human error and

APPLICATIONS IN THE CLOUD

Percentage saying the following types of IT systems currently reside inside the cloud.

Source: Harvard Business Review Analytic Services Survey, February 2017

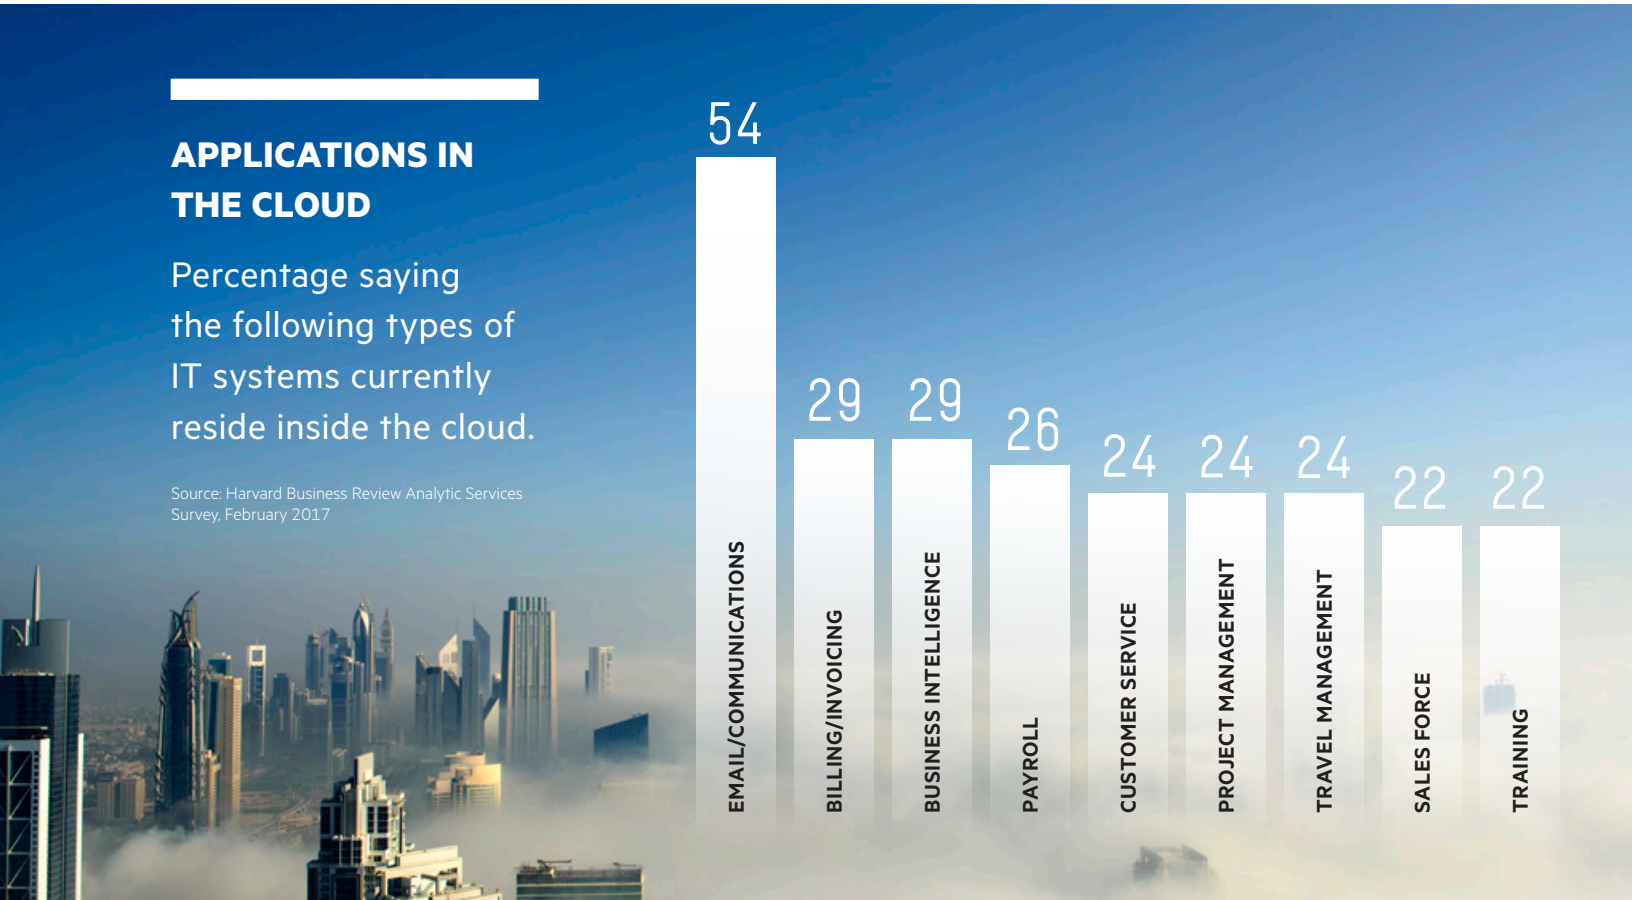

preventing unsanctioned actions. Further, patching and updating application can also be handled via automation, which ensures the most up-to-date versions of software are running and reduces potential breaches.

Cost: Finally, consider costs. The lure of lower costs (whether it actually pans out or not) drives many companies to the public cloud or helps persuade them to deploy a SaaS-based solution. But a monitoring and management tool, especially one comprehensive enough to oversee an entire infrastructure, adds costs to the mix. While this cost is sure to be higher than that of a single solution, it will likely be far less than the aggregate cost of multiple point solutions jerry-rigged together. In addition, the benefits afforded by a comprehensive solution have the potential to reduce risk and operational costs in terms of time and resources saved through automation, governance, and more.

A hybrid infrastructure enables enterprises to build an IT ecosystem consisting of best-of-breed components.

Consider also that a single-solution product may offer the ability to integrate with large-scale operations management platforms from vendors such as BMC, HPE, and ServiceNow.

A hybrid infrastructure enables enterprises to build an IT ecosystem consisting of best-of-breed components. While the benefits are many, the result is an increasingly complex hodgepodge of solutions that are best managed from the single console of a unified monitoring and management tool. For many enterprises, finding the right tool kit to manage this mixed environment is fast becoming a key differentiator for success. ■

03 NOW TO NEXT

Hybrid IT management: Lessons for leaders

Integrated management tools will allow for a more optimized hybrid infrastructure.

Building management in from the start simplifies the growth of your enterprise environment.

Quality tools will address concerns ranging from traditional systems management issues to governance and compliance issues.

04

Evaluating your infrastructure: Finding the core apps and services

By Scott Koegler

- How to distinguish core vs. critical apps
- Where do the applications reside?
- Looking at the integration of your services
- Finding the right balance

Defining your core application and services mix

Enterprises today are dealing with change on multiple fronts. Increasingly, their employees are technologically savvy, geographically distributed, and highly mobile. Central IT manages everything from legacy software running on aging computing platforms to the latest cloud-based apps. Business units sometimes implement new apps without direction from IT. And communications have become more complex than what existing services were designed for, including voice, video, text, and document storage.

Many established companies face younger competitors that are unburdened by legacy systems, meaning they need a clear plan for transitioning to a more agile infrastructure and application portfolio. This outline defines the most critical steps needed to identify your core applications and the mix of services needed to support both apps and enterprise operations.

The mix of applications and services in any organization grows quickly as different use cases are addressed with specific technological solutions. Over time, the technology infrastructure that supports those applications changes, adding complexity to the environment and complicating attempts to modernize older solutions and bring them up to current standards. You can't integrate

diverse applications in one swoop. You need to evaluate each app based on the value it generates and the level of difficulty involved in modernizing it. Once those evaluations have been made, you can start prioritizing core apps for modernization.

> Applications that support critical business operations vary from one business operation to another.

Identifying your core applications can be a complex task because applications that support critical business operations vary from one business operation to another. You can identify which applications are core and which are ancillary only when a full inventory of the applications has been completed. And conducting a full inventory is complicated by the fact that not every application in use is known to IT. Many business units have adopted software-as-a-service systems to meet immediate needs because they are offered free or at a low cost. And while it can be argued that these unauthorized systems are expendable, the departments that initiated their use may have grown to depend on them.

For all these reasons, you need a functional definition for "core application."

Core versus critical applications

While some applications may be critical to a particular function, job, or project, they are not necessarily core to the enterprise. Gartner defines enterprise applications as "designed to integrate computer systems that run all phases of an enterprise's operations to facilitate cooperation and coordination of work across the enterprise." While these enterprise applications can be considered critical to the company's operations, they may not be the only core applications in the portfolio.

At most established organizations, core applications reside on legacy hardware and have most likely been customized and expanded through the history of the company so that they are now part of a complex collection of multiple interconnected systems, the individual components of which have become nearly impossible to separate.

66%

of respondents are in a hybrid environment.

Source: Forrester – (Forrester Adoption Profile: Private Cloud in North America, March 2017).

Deloitte describes these combinations of core systems as "often a tangle of complexity and dependency that is daunting to try to comprehend, much less unwind." Deloitte reports that the cost of maintaining these core applications is the largest single line item in an IT budget. Migrating core apps to more efficient and updated infrastructure that is part of a hybrid environment can lower expenses and result in a more agile set of services that is more attuned to a modern workforce and the needs of a competitive organization.

So, while line-of-business applications are usually already clearly core applications, the ease with which non-central IT has been able to develop and deploy applications using cloud services without the knowledge of central IT means that it is still necessary to evaluate almost all existing applications and processes in order to determine where they belong in the long term.

The services mix

Enterprises that want to compete in disruptive environments must be able to respond quickly to competitive threats. Core applications that maintain the underpinnings of the business are typically baked into the fabric of operations and rely on legacy communication and data storage services that are not easily adapted to change.

The challenge for IT is to incorporate new services that can integrate data and processes that support current operations while at the same time delivering on the promise of agility, because the data in legacy systems is critical for use with new initiatives.

Technologies like cloud, analytics, and digital transformation efforts are generally seen as higher value with lower expense than are core systems. But the fact is that the data that runs the business resides in those core systems and the newer applications depend on access to and integration with the resources stored there. As businesses add advanced facilities, they are building a hybrid environment that spans legacy and advanced infrastructure, technologies, and applications. The core systems remain as foundations for newer initiatives.

The concept of flexible and scalable computing is a key reason to create a hybrid environment.

Efforts to enable a hybrid environment need to be undertaken with an eye toward the impact that data-driven technologies like customer support and targeted marketing can have on achieving company goals. Creating an effective combination of technologies to support every level of application requires the right skills and knowledge. Creating a private cloud environment can be the basis for establishing a set of standards to which future cloud services can be connected. And because the first step is private, it can be designed to closely integrate with existing legacy systems as well as existing IT management tools because IT has control of both environments.

The private cloud in combination with traditional IT can become the central hub connecting with public cloud services to make up the hybrid environment. The reverse can also be true, with apps that start in the public cloud moving back within the data center or to a public/private cloud mix.

The concept of flexible and scalable computing is a key reason to create a hybrid environment. It allows the enterprise to make incremental changes and experiment with new applications that can leverage existing core data in new ways without interfering with the operation of existing systems and at minimal expense.

As the proper balance of services and core applications takes shape, users will seamlessly access the resources of the enterprise through applications that make use of the resources of both existing and new applications. The hybrid infrastructure enables presentation of unified views of data that simplify and speed innovation because users are able to access the information they need without regard to what program or application created or maintains the data.

A well-structured hybrid environment also simplifies the allocation of security privileges through overarching access control, so users can reach the data they need but only as they are permitted. Overall, an increased utilization of resources adds value to the company by enabling it to leverage key assets to drive innovation and competition. ■

04 NOW TO NEXT

Defining core apps and services: Lessons for leaders

Start with an inventory of applications and services.

Core and *LOB* are not interchangeable terms.

Proper balance of public/private cloud and central IT are likely to be the most cost-effective solution.

05

Data security considerations in a hybrid IT world

By Esther Shein

- Information is critical
- Approach the problem from multiple angles
- Shadow IT remains a problem
- Finding all your data and maintaining control

Securing your data across multiple platforms

By now, moving at least some business processes to the cloud is not a question of if but when. So how do you keep your information safe while embracing all the benefits cloud computing offers?

Even if the enterprise is using private clouds and virtualization, your data may physically reside in infrastructure that is owned and operated by an external service provider.

When control is shifted to a third party that owns, operates, and manages infrastructure and computational resources, it is incumbent upon security professionals to put measures in place to maintain the safety of their data. It comes down to doing your research and due diligence, figuring out your threshold for risk, and not giving up all of the keys to the castle.

Ask questions, conduct audits

There is no single measure or technique that can keep a company's data secure, regardless of whether you use an on-premises data center or the cloud, notes Paul Hill, senior consultant at SystemExperts. "When using the cloud, an organization has to understand what responsibilities are outsourced to the cloud vendor and what will remain the responsibility of the organization," he says.

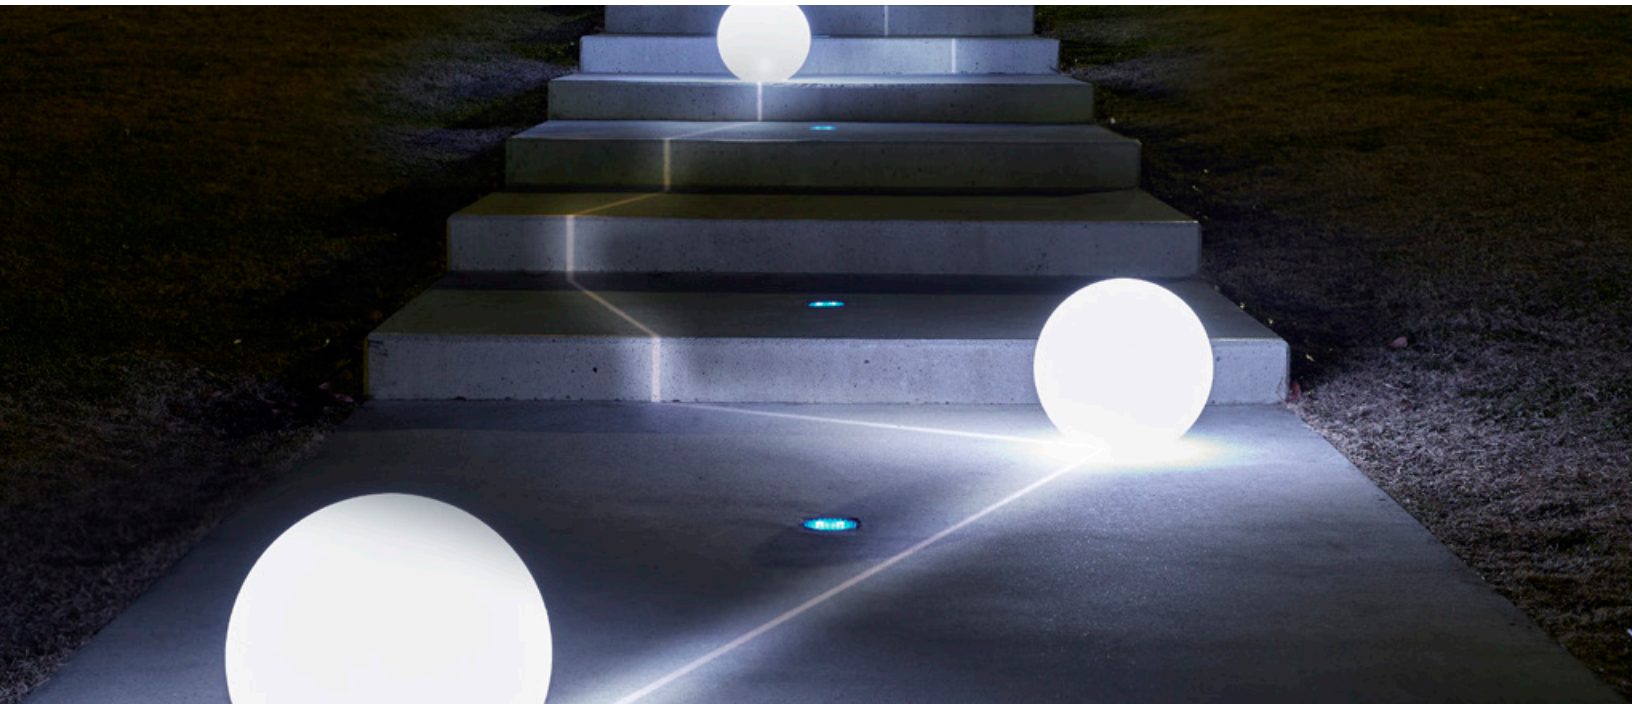

First and foremost, ask for credentials when evaluating a cloud service provider (CSP). What level of trust and reputation does the provider have in the market? How will it protect valuable data and personal information?

"It's important to ask these questions and have the CSP describe their security operational controls, such as how they handle security breaches and how threats are addressed, as well as how certain insider threats are identified and countered," advises Thomas Hogan, sales specialist for BT Cloud Compute. Additionally, organizations should deploy identity access management to control the security credentials in the cloud and manage who has access to what information.Hill agrees: "Without careful oversight, it is all too easy for someone in an organization to misunderstand the responsibilities and assume that the cloud provider is doing more than they really are." For example, if a CSP states that it has achieved PCI compliance, does that mean that your applications are automatically PCI compliant? Or is the scope of the compliance limited to the payments made by customers to the CSP? "Strong IT governance by knowledgeable individuals is essential, or the organization should engage a third party with the expertise to review the issues," Hill says.

Without careful oversight, it is all too easy for someone in an organization to misunderstand the responsibilities and assume that the cloud provider is doing more than they really are.

"If your organization is required to keep its data within a geo-location due to regulatory issues, you should make the CSP describe how it will ring-fence or guarantee data will not cross borders," adds Hogan, "It should also address access methods, encryption techniques, and all authentication processes needed to access data."

In terms of the responsibilities the CSP is willing to provide, the organization needs a mechanism to determine how well the service provider is implementing the security controls, Hill says. "This is typically done by a combination of testing and relying on independent security audits under a compliance program," he notes. "In some cases, an organization may not be satisfied by a compliance statement, and it may require that it perform its own audit."

This tends to be more practical when using a small cloud provider. Amazon, Microsoft, and Google generally don't allow customers to perform their own audits, he points out: "Customers of those providers usually have to be satisfied by compliance certifications and some form of testing that they can perform."

In some cases, depending on the sensitivity of the data and the nature of the customer relationship, an organization may want or even need to assume some of the responsibilities the CSP is willing to provide, says Hill. For example, an organization might determine that it needs to encrypt its data at rest. Many cloud services provide some level of cryptographic key management. But an organization might decide that the cloud provider should not be able to decrypt the data.

"In that case, the organization will need to assume all aspects of key management or use a third party to perform the key management," he says. "If an organization wants the ability to see any subpoenas served and control the response to them, then encrypting the data with keys under its own control is a critical control."

Take a multi-pronged approach to security

A multi-layered approach to cybersecurity provides flexibility and choice in selecting the right security solutions for the computing environment. In addition to foundational safeguards, there are many feature-based protections that mitigate security risk.

At a foundational layer, organizations should ensure that CSPs have the right certifications for specific needs, including ISO 27001, PCI-DSS1, and SSAE. "This ensures the right adherence to security processes and procedures," says Hogan.

And don't just ask what certifications the CSP has. "Probably more importantly is when the certifications are available to the buyer," meaning how recent they are, adds Jim Hurley, principal analyst at ISG Research.

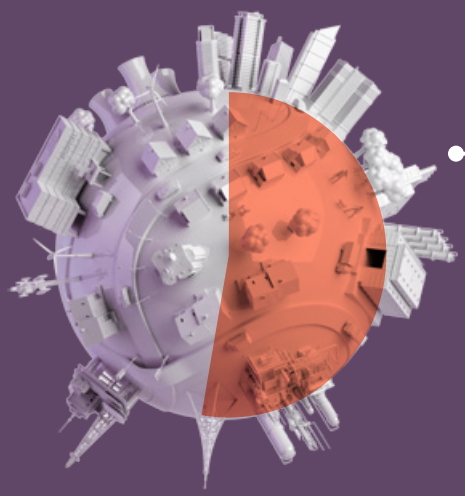

53%

of respondents have, or are considering, bringing their workloads back on premises from public cloud providers.

Source: IDC
('Pay-per-Use Models in IaaS Survey,' July 2016).

In terms of feature-based protections, Hurley recommends organizations look at some of the newer intrusion deception techniques that can be deployed either in the cloud or on-premises. “Intrusion detection is old stuff that doesn’t work,” he says.

Deception protection involves putting up a fake screen that mirrors the existing environment so that anyone with ill intent will only see a decoy or a mirage. “That decoy, when touched, would trigger an immediate notification either back to the cloud services provider or operating center in the organization if they have one,” Hurley says. “So you wouldn’t be wondering six months later if you were comprised.”

Deception fabrics are relatively nascent technology, however. Hurley says it will take some time for organizations to understand and use them effectively as they migrate from other tools and prescriptive measures.

Shadow IT is still lurking

Shadow IT has been around for many years, and reining in cloud apps developed outside of IT continues to be an issue for organizations. If not mitigated properly, it can cause problems for any organization, regardless of industry. Shadow IT problems often arise when businesses feel pressure to digitally transform their organization in order to stay competitive in the marketplace.

In a survey for the 2016 Cloud Security Alliance (CSA) Mitigating Risk for Cloud Applications report, 62 percent of respondents said their companies have written policies discouraging use of unsanctioned apps, but few have technical controls in place. Thirty-eight percent block unsanctioned apps outright, while 29 percent use a proxy or firewall to redirect users.

The majority of security professionals remain as concerned today about shadow IT as they were last year (49 percent), the survey also found. Another large portion are more concerned than last year (30 percent), while a smaller percentage are less concerned or were never concerned (13 and 8 percent, respectively).

Hogan maintains that IT can mitigate the risks of shadow IT by embracing it. "IT can work with the business to build an appropriate security and compliance framework to address any lingering concerns," he says.

In the cloud realm, one way to secure data could be the use of cloud access security brokers (CASBs) to improve visibility and control over both unsanctioned and sanctioned apps. The role of a CASB is to monitor data activity and enforce policies across multiple cloud apps, the CSA report notes.

For 32 percent of respondents, the most important use case is data loss prevention. Already, 60 percent of security professionals say they have deployed or plan to deploy a CASB. Gartner is projecting CASB deployments will grow rapidly in the next few years, reaching 85 percent of large enterprises by 2020.

Know where your data is and maintain control

Enterprise data security is expected to keep IT on its toes as organizations increasingly operate a combination of legacy systems, converged technologies, and public and private clouds. Recent ISG research found that 40 percent of workloads are now in a hybrid environment, Hurley says. "The plan going forward would seem to indicate…that will increase to 60 percent of workloads by 2018 and hold steady through 2020 at that level," he notes.

Data is more distributed than ever with the explosion of both structured and unstructured data, thanks to cloud computing and big data. This makes it even more compelling to keep tabs on your data. Yet, it's an area that clearly requires improvement. Only 10 percent of respondents to Vormetric's 2016 Data Threat Report survey claimed little or no knowledge of the location of their sensitive data, yet nearly half (47 percent) of all respondents said they have "some idea" where their sensitive data is located. Perhaps most troubling is less than half (43 percent) claimed to have "complete knowledge" of where their sensitive data is located.

"Strong IT governance by knowledgeable individuals is essential," stresses Hill, "or the organization should engage a third party with the expertise to review the issues." ■

05 NOW TO NEXT

Keeping your data safe: Lessons for leaders

Understand where and how your data is stored.

Watch out for shadow IT.

Multiple types of security make for more secure environments.

06

Pricing hybrid IT

By Lynn Greiner

- Paying for unused capacity
- Does your business need storage or compute?
- Make sure that staff is on board
- It's not an instant solution

How to finance hybrid IT

Budget time is rarely fun in the enterprise IT world. Squeezing as much as possible out of a usually constrained budget when there's so much to do is enough to drive the most seasoned professional to despair. Yet, it must be done—and done well—to maintain resources that allow the business to flourish.

In years gone by, the juggling act between capital and operational expenses (CapEx and OpEx) was a constant. With all resources residing in the data center, IT had to invest scarce capital dollars just to keep the ship afloat, while still spending significant amounts on operations.

Then came the public cloud and the opportunity to shift expenses to the OpEx side of the ledger, which many chief financial officers prefer. However, that came with its own challenges, including the potential for skyrocketing operational costs as well as compliance and security requirements that public cloud providers could not always meet. These complications provided the impetus for the hybrid cloud.

In an article in IEEE's Cloud Computing journal, cloud economics guru Joe Weinman notes that hybrid clouds "can offer economic benefits, even when—in fact, particularly when—the unit cost of public cloud services and resources is higher than that of private dedicated resources, a scenario that some reasonably sized, well-run IT shops can and have achieved."

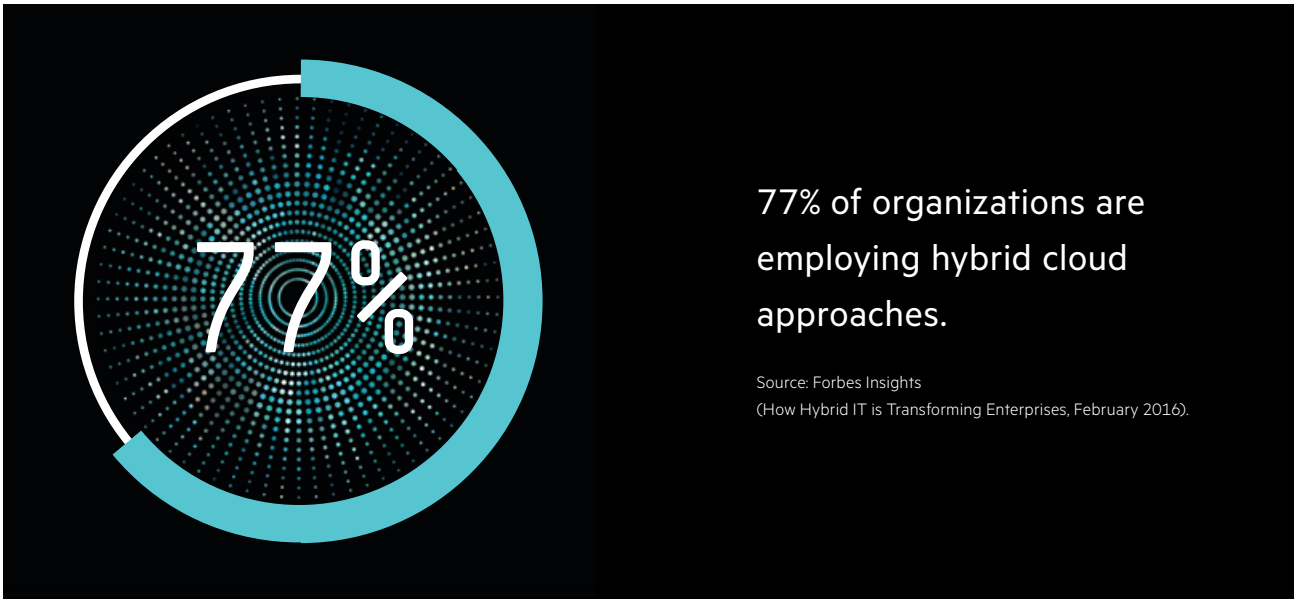

With hybrid cloud, the cost models change. There can still be capital requirements, depending on the state of your in-house infrastructure (it will likely need some software changes, at the very least). Vendors are now being creative in how they price these elements, allowing payments to scale to match business growth. But the public cloud component will be all OpEx.

Capacity on demand

One of the appealing factors of the cloud model is the ability to scale up and down quickly, spinning up and decommissioning resources as required. In theory, that controls costs as well, but as many companies have discovered, it doesn't always work that way. Without strict governance controls, including policies and automation, cloud resources that should have been turned off can remain active and end up adding to costs, to the point that the supposedly cheaper cloud costs significantly more than the in-house resources it replaced.

The cloud platform market will triple by 2020, from $21 billion in 2016 to $64 billion.

Source: Forrester.

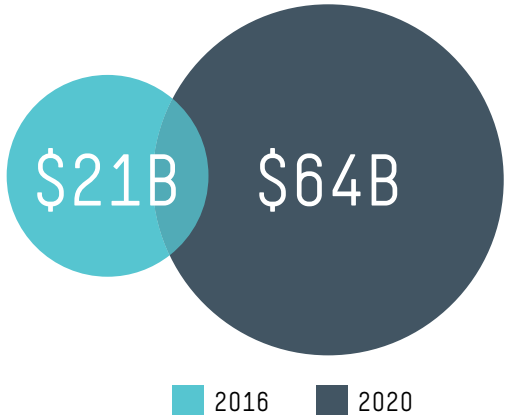

In a hybrid model, while this is also a risk, the public cloud is often used only for surge capacity to supplement in-house resources, at least initially. As cloud-native applications enter the mix, however, it's even more important that any financial plan involving hybrid cloud include governance controls to determine how public and private clouds can be used in the most cost-effective manner.

Private clouds can also take advantage of vendors' "pay as you grow" plans, in which users receive additional capacity from the outset and pay for it as they grow into it. One of the biggest annoyances for corporate IT is the length of time it takes to purchase and install additional capacity when required. Users want it yesterday, but the vendor may quote weeks or months before delivery. That disconnect in expectations often leads business units to put systems into the public cloud that shouldn't be there, simply because they don't have time to wait for the proper resources.

While missing a business opportunity is expensive, a precipitous leap into the public cloud can be even more costly in a number of ways. If regulatory and compliance requirements for an application aren't met, there could be fines. And if security isn't up to scratch, a data breach could cost millions of dollars, and potentially put the company out of business. However, forecasting capacity requirements can be akin to a black art. IT (or the business) may know that it will need additional private cloud resources later in the buying cycle but can't justify the immediate expenditure when purchasing. That can lead to additional expense later.

Vendors have addressed these issues by introducing hybrid payment plans. The customer commits to a certain capacity, making fixed payments as with any equipment lease. However, the vendor also provides buffer capacity with usage metered and charged for separately. If the customer's usage bursts into the buffer, there's a bill. If not, there is no charge for the buffer that month. As the business grows and begins to consistently devour the existing buffer, additional buffer capacity can be installed as required and the amount of committed capacity (and its fixed charge) increased. It's a cloud-like way to expand without putting sensitive data in the public cloud.

For less sensitive data, of course, the application can automatically burst to the public cloud.

Not just compute

While many concentrate on the data center aspects of hybrid cloud financing, you also need to factor in telecommunications costs. Data transfer to and from a public cloud isn't free, and while it's nowhere near as expensive as it used to be, public cloud usage still adds to overall corporate bandwidth requirements. In a 2011 paper published in the journal Information Systems Frontiers, researchers Oleksiy Mazhelis and Pasi Tyrväinen argued that "as the volume of data transferred to/from the public cloud increases, a greater portion of the capacity should be allocated to the private cloud."

Don't forget the people

The shift to hybrid cloud doesn't just impact hardware and software. There's an unavoidable people cost as well. Managing a public cloud vendor requires different skills than those for running a data center. IT staff may need training, and in fact, some may need to be replaced.

Don't expect immediate returns

Moving to a hybrid cloud model doesn't come with instant returns. The upfront costs involved in "cloudifying" existing infrastructure, adapting applications for private or public cloud use, acquiring monitoring and management tools, training, and other expenses have to be considered.

How long will it take? Intel IT built a model of the projected costs for its own hybrid cloud initiatives in 2013, which showed that potential savings would only begin to manifest themselves in year two but would grow steadily after that. At the time, the authors noted, "The hybrid cost savings may vary from those shown, depending on the rate of cost-efficiency improvements in both the public cloud and in Intel's private cloud. Public and private hosting costs are both decreasing rapidly; however, we believe that due to the operational efficiencies we are putting in place, including new technologies, training efficiencies, and use of open source solutions, hybrid costs will decrease as fast as or faster than public cloud costs. It should be noted that the efficiencies public providers have introduced into large-scale computing are a catalyst to help direct and lead the way for all computing—including private clouds."

While Intel's savings may not be achievable by all—its internal IT operation is extremely efficient, keeping its private cloud costs low—its model still indicates that a hybrid cloud can indeed save money while meeting the operational goals of the company. But to do so, the cost structure of the private cloud must be similar to that of the public cloud provider. ■

06 NOW TO NEXT

Financing hybrid IT: Lessons for leaders

Keep a close eye on your operational expenses during cloud deployments.

Have a plan and goals to determine the effectiveness of cloud and hybrid cloud transitions. If those goals aren't being met, find out why.

Don't be afraid to de-cloud applications and services if they prove more expensive to operate in the new paradigm.

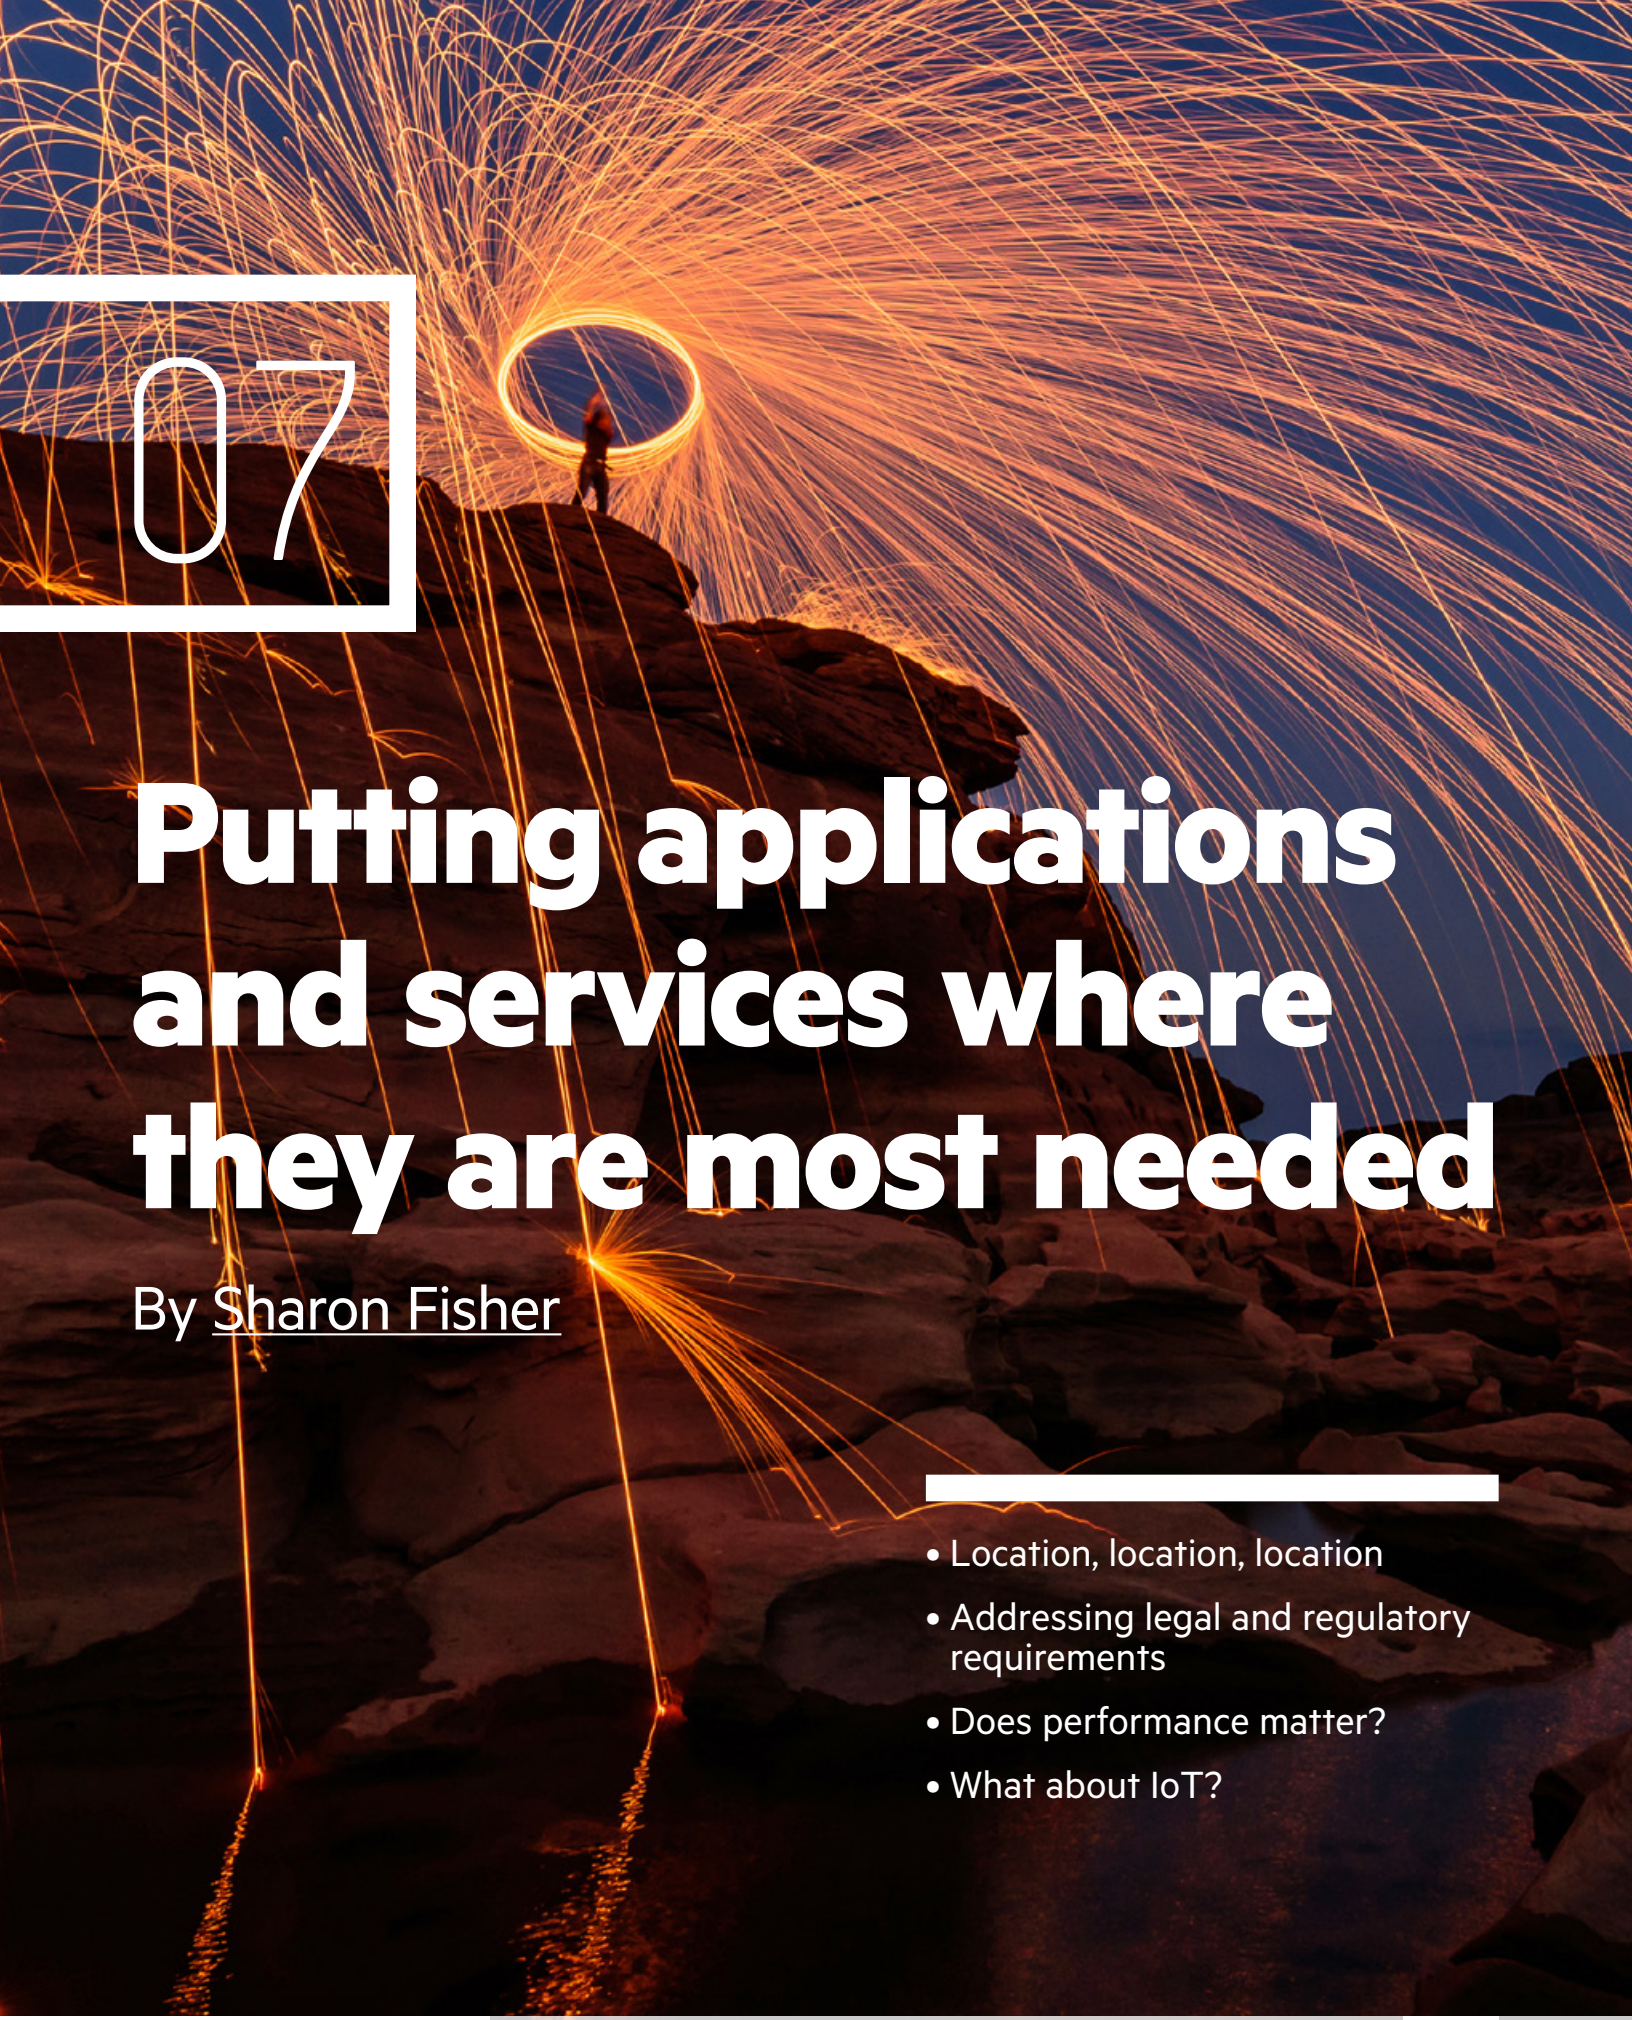

07

Putting applications and services where they are most needed

By Sharon Fisher

- Location, location, location
- Addressing legal and regulatory requirements
- Does performance matter?
- What about IoT?

Edge data centers bring processing closer to home

Edge data centers move data processing closer to where the data is, reducing the amount of data that traverses the Internet. The word edge usually connotes something dangerous and risky. But edge data centers can also mean better, safer IT for the enterprise.

Providers argue about the definition, and one could say it's nothing more than a marketing term. "We actually trademarked the term 'edge data center' three years ago," says Phill Lawson-Shanks, chief architect and vice president of innovation at EdgeConneX, a Herndon, Va., data center provider. "We marketed the bejesus out of it." Other edge data center vendors include 365 Data Centers, Netrality, and Zayo Group.

The definition is also a moving target. "Edge is defining itself as we speak. It's going to take some time to really define what 'edge' means," says Srdan Mutabdzija, future offer manager at Schneider Electric, a provider of physical infrastructure for data centers with North American headquarters in Andover, Mass. It can include technologies such as a single-rack micro data center, a regional data center that does processing and sends results to a hyperscale data center, a distributed data center, or a data center in a box. Because of the fuzziness of the definition, it can be hard to figure out how much the market is growing, though Mutabdzija estimates 10 to 20 percent annual growth.

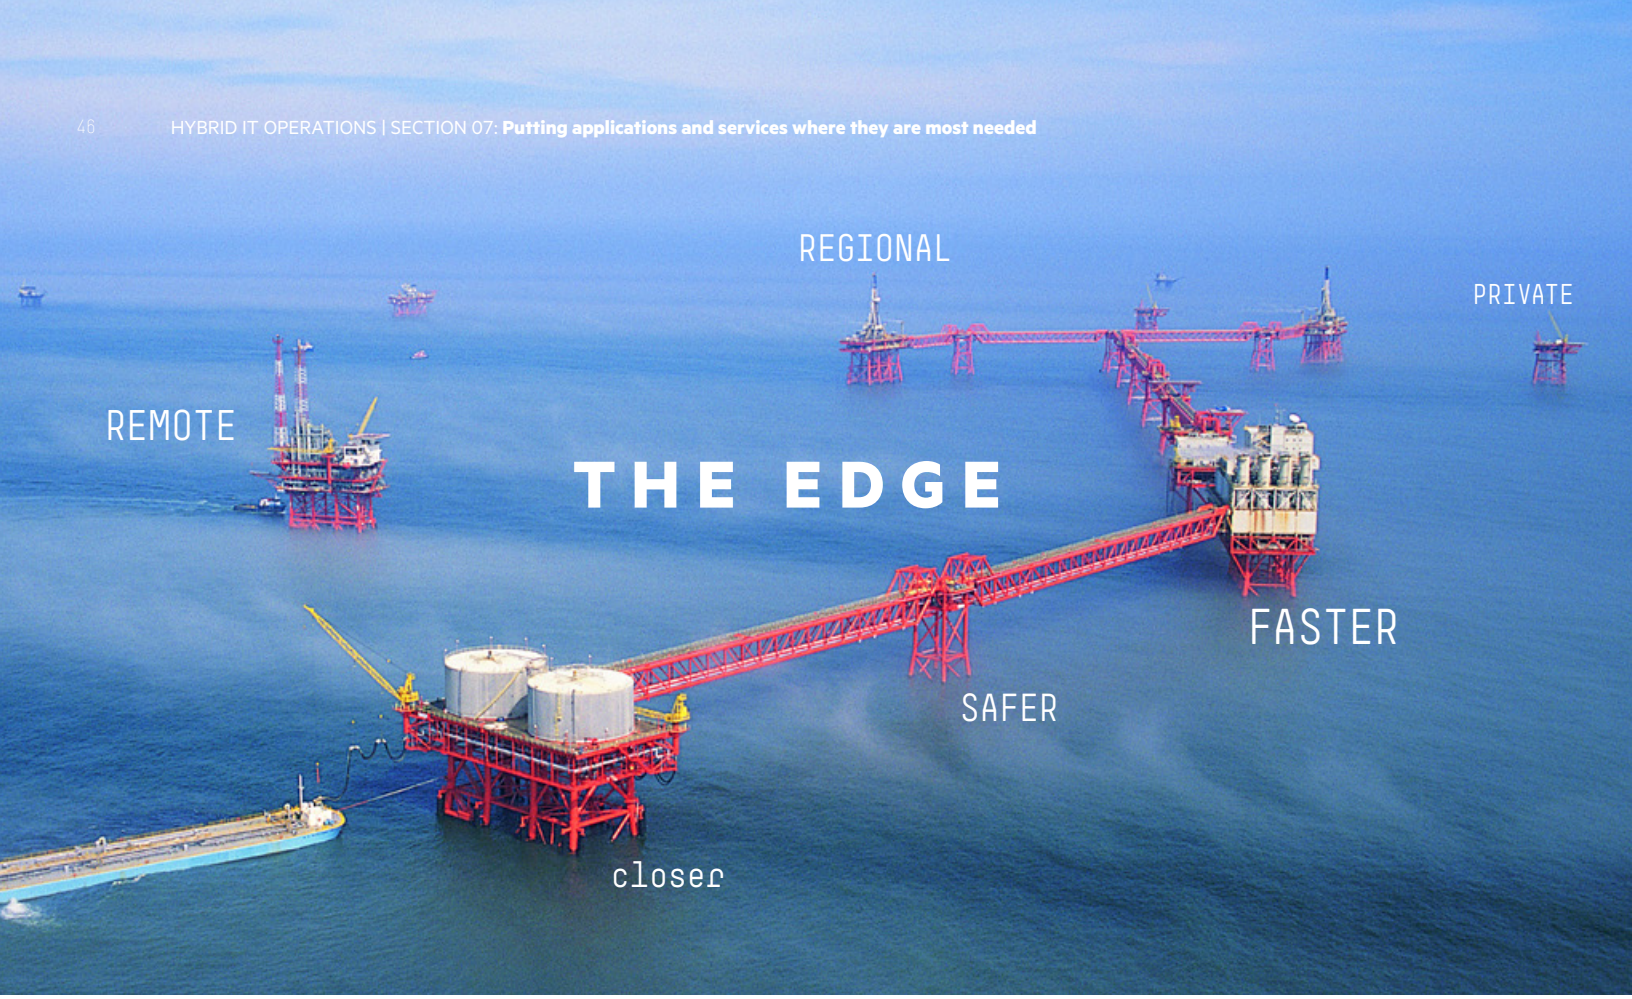

Never mind what you call it. What does it do?

"People are putting the IT power where they need it," says Jennifer Cooke, research director for data center management at IDC. End users can process their data at a service provider or through a colocation provider—any service provider that offers resources where they are needed. Edge data centers can also be situated at a remote location, or a location where a company wants to use private resources rather than public resources due to regulatory and privacy requirements.

Particularly in Europe, some countries have data sovereignty regulations that restrict the ability of companies to move personal customer data out of the country. In one ongoing legal case, the U.S. Department of Justice is claiming access to user data stored in Ireland because the vendor, Microsoft, is a U.S. company—an effort that Microsoft, the European Union, and the Irish government are all resisting. The push for more access to data by U.S. law enforcement officials is running up against an increased emphasis on user privacy in Europe. Britain's recent decision to leave the European Union adds another layer of complication.

EdgeConneX first got involved in edge data centers by negotiating fiber telecommunications access to buildings and their tenants. "We know where the Internet truly is," boasts Lawson-Shanks, claiming the company has rights in 46,000 buildings worldwide. "Network guys don't share data,

but they share it with us. We have a proprietary database with every piece of fiber in the country."

What brought edge data centers into the mainstream? In a word, Netflix. About four years ago, Netflix exploded on the scene and started "destroying everyone's backbone," as Lawson-Shanks puts it, through its massive load of video data. At the same time, Comcast developed a cloud-based DVR service, based in Denver, and tested it out in the Chicago and Atlanta markets. It worked great, until it became too popular and put too much of a load on the Internet.

Enter EdgeConneX. Comcast asked if some of its buildings could be repurposed into data centers, with an over-the-top video ecosystem to offload traffic. "We built 23 of them in 18 months," Lawson-Shanks says. "Now we have 30." And wherever they're built, they service major Internet sites such as Facebook, Netflix, Google, and Microsoft—75 percent of the Internet, he claims.

Similarly, Comcast needed to move to edge data centers when it implemented its Xfinity video-on-demand service. To deliver the required performance, the company rolled out approximately 75 identical data centers around a year and a half ago along the East Coast, according to Cooke.

> The push for more access to data by U.S. law enforcement officials is running up against an increased emphasis on user privacy in Europe.

(Related reading: Edge vs. central IT: Where do my apps and services belong?)

The latency problem

Flying packages to Memphis and then flying them to their destinations works for FedEx. Why doesn't it work for data? First of all, Lawson-Shanks says, using an analogy, FedEx has more capacity in its aircraft than the Internet currently does. Second, "while a late package can be an inconvenience, it's not as big a problem because of the way Internet protocols are designed."

In contrast, the typical Internet protocol TCP/IP works with email by chopping the message into pieces and sending those pieces off using the best route available at the time, to be reassembled at their destination. But that doesn't work for video and other high-performance applications—it ends up choppy.

Other industries that see value in edge data centers include retail, manufacturing, healthcare, and finance. Focus areas vary by industry. Retail typically runs in the cloud and is mostly about connecting to apps. Industrial companies tend to prioritize real-time process control and latency.

Healthcare and finance companies are more focused on government regulations and security. For these industries, it makes sense to process data locally because it's more secure than pushing it all to the cloud. In healthcare, fines for HIPAA security violations totaled $23 million in 2016, up from $6 million in 2015, and Cooke sees those growing. "It's something that's been on the rise but is not even near where it could be," she warns.

How well does it work?

One EdgeConneX user is Cloudflare, a content delivery provider based in San Francisco. Cloudflare is collocated in a number of EdgeConneX data centers in cities such as Detroit, Salt Lake City, Sacramento, Las Vegas, and Phoenix, according to Nitin Rao, head of the company's infrastructure strategy team. Cloudflare operates in 102 cities in 50 countries and participates in 150 Internet exchanges, supporting 5 million websites. "Our job is not done in Frankfurt or San Jose," Rao adds. "It's just as important to make websites fast in Omaha, Sacramento, and Zagreb."

Partnering with edge data centers helps improve performance for Cloudflare customers in several ways. First, it reduces traffic on the overall Internet, because Kansas City visits are now served by Kansas City, rather than having to travel to Chicago where the data used to have to go. That reduces latency and improves performance.

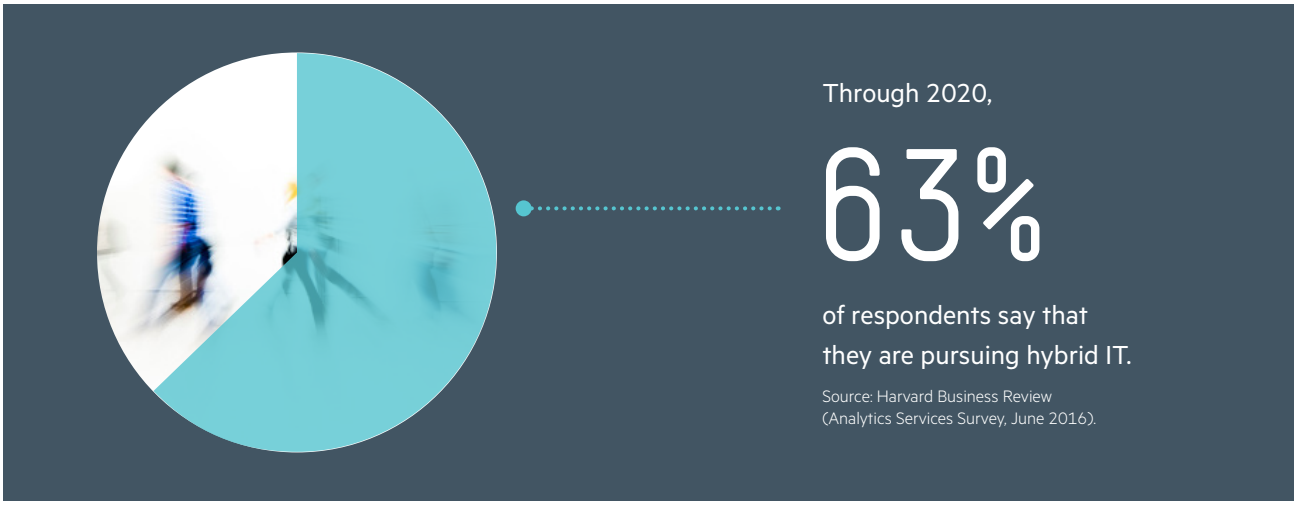

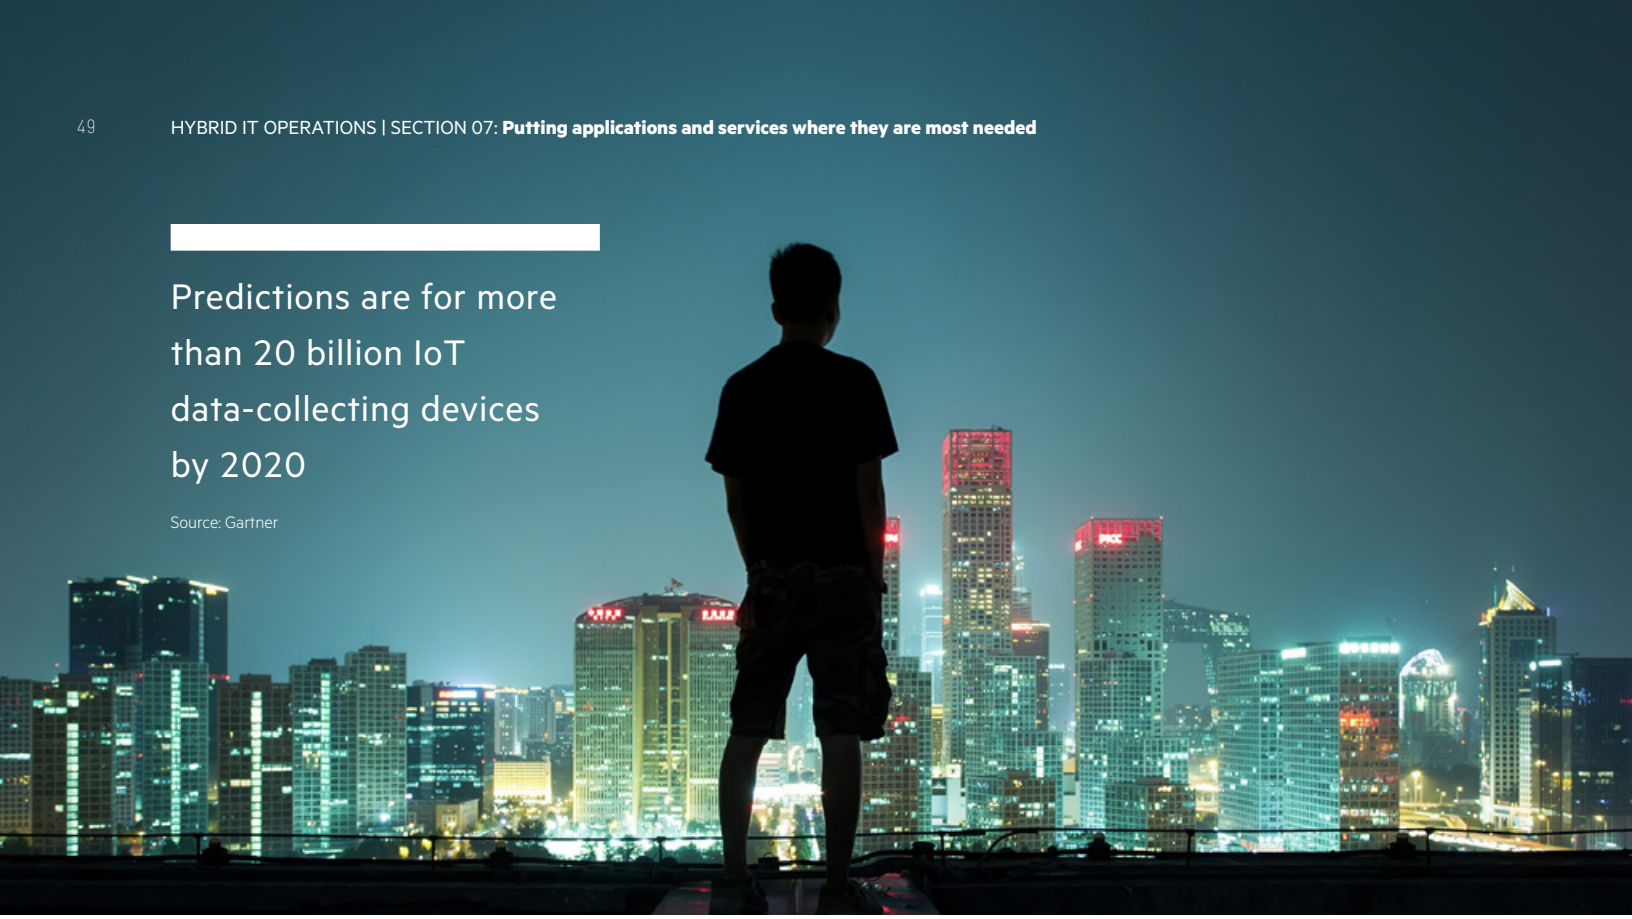

Predictions are for more than 20 billion IoT data-collecting devices by 2020

Source: Gartner

In addition, reducing the distance data travels improves security by reducing the number of potential attack points. Finally, it saves money. “Paradoxically, for a network like ours, it could be more cost-effective to build a bigger network than a smaller network,” says Rao.

How does Cloudflare choose a data center to partner with? To begin with, the company looks for fairly high security standards, noting that Cloudflare itself is compliant with PCI Data Security Standard v3.1.

After that, as the saying goes about real estate, it’s location, location, location. The most important feature of a data center is that it reaches the most local eyeballs. “Typically there’s one building in every city where every major ISP is, and it’s a no-brainer to show up there,” Rao says. “Network engineers know the addresses of the most important data center for every city by heart. It’s a clear winner-takes-all market.”

While Rao doesn’t have specific numbers for how much time or data his company has saved by using edge data centers, he offers a rule of thumb. Cloudflare is responsible for about 5 million websites, or 10 percent of the Internet in the U.S. Moving to an edge data center from a major interconnection point such as Chicago could reduce latency by 20 milliseconds. “If you can speed up a meaningful portion of 10 percent of the Internet by 20 milliseconds, that’s significant,” he says.

Time savings in countries in Eastern Europe and the Middle East can be even more significant,

Rao says. “In Eastern Europe, until you have a data center in that country, traffic invariably falls to Frankfurt,” he says. For example, once his company set up a data center in Bucharest, many Romanian sites offered faster performance.

IoT will redefine the edge once again

While Netflix brought edge data centers into the mainstream, the next big push is likely to be the Internet of Things (IoT). Predictions are for more than 21 billion IoT data-collecting devices by

> “On-premises has traditionally been the default choice for collaboration and communications products”
>
> Source: Gartner

2020, and it’s not going to be realistic to ship all the data from billions of devices to the cloud, process it, and then send it all back. Instead, organizations will process data closer to the sensors and send just the analysis or conclusions to the cloud.

It remains unclear how edge data centers will evolve going forward. No matter what happens, data processing and analysis will become more distributed. The result could be a redefinition of the data center itself. ■

07 NOW TO NEXT

Delivering the edge: Lessons for leaders

Local data centers can improve the performance and availability of applications and services for users

Regulatory concerns can impact where information can be located.

Security remains a top issue when choosing providers.

CONTRIBUTORS

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.

Andy DeBernardis

Andy DeBernardis is currently part of the worldwide cloud solutions marketing team at HP Enterprise. Andy has been with HPE for 10+ years in worldwide roles including managing HP worldwide cloud sales enablement, HP software cloud sales enablement, HP software partner marketing, and HP software vertical solutions marketing.

Paul Ferrill

Paul Ferrill has written hundreds of articles for publications such as Datamation, Federal Computer Week, Information Week, InfoWorld, Network Computing, Network World, and PC Magazine. He serves as CTO for Avionics Test and Analysis Corporation, a woman-owned small business specializing in providing engineering expertise to help solve complex data analysis problems. His third book is for Microsoft Press and coauthored with his oldest son entitled "Training Guide: Designing and Implementing a Server Infrastructure."

Sharon Fisher

Sharon Fisher has worked in the computer trade press for more than 20 years. She is currently on contract as a content strategist for the Economist Intelligence Unit. She has worked on staff for Computerworld, CommunicationsWeek, and InfoWorld, and freelance for publications such as Byte, Network World, Information Week, and PC Week. She also served as a Gartner analyst for seven years, and has worked for Hewlett-Packard and Interex. She is the author of several books, including Riding the Internet Highway.

Lynn Greiner

Lynn Greiner is a freelance technology writer, and the odd combination of a geek with a business degree. She spent many years working in corporate IT, both on the ground and in management, and is ITIL certified.

Scott Koegler

Scott Koegler is a technology journalist with 20 years experience writing about business, computing and technology topics. He was CIO for 3 mid-sized companies for a total of 15 years and that experience has provided an important perspective for his journalistic contributions. His work with developers, marketing, business processes, and C-level executives has allowed him to focus on the intersection of business and technology. Scott publishes ec-bp.com, a supply chain industry newsletter.

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.

Amy Newman

Amy is a B2B technology writer and editor with over 15 years of experience following and analyzing IT infrastructure trends. She coauthored "Practical Virtualization Solutions: Virtualization from the Trenches" published by Prentice Hall Pearson Education in 2009.

Esther Shein

Esther Shein is a veteran freelance technology writer and editor. Her work has appeared in numerous publications including Inc., Computerworld, InformationWeek, CIO, Communications of the ACM, TechTarget, Forbes and The Boston Globe. She writes features and news articles as well as thought leadership whitepapers, customer stories and marketing materials for industry leaders and innovative, early stage startups.

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

Visit insights.hpe.com

enterprise.nxt

Hewlett Packard
Enterprise

a00006639enw

SHARE THIS WITH YOUR NETWORK