451 Alliance Fall 2014 *Global Digital Infrastructure Trends* Articles

I. Green Datacenter Metrics: Do You Measure Up?

Datacenter availability and economic pressures have long been hurdles to datacenter efficiency, but new government initiatives may be changing the playing field.

Datacenters now account for an estimated 2% of all US electrical consumption and carry a significant carbon footprint. Yet until recently, it's been private industry or the individual states that have driven 'greenification' forward, far more than the federal government.

Data Going 'Green'?

In September 2014, Barack Obama proposed cutting carbon pollution from power plants by 30%, but his plan represents the first time a U.S. president has attempted to regulate carbon emissions.

Moreover, the Environmental Protection Agency isn't expected to finalize the actual rule until June 2015, and individual states have until 2017 to submit their reduction plans.

Thus, current greenhouse gas initiatives are mostly being driven at the state level. Beginning with California Assembly Bill 32 – a cap-and-trade program introduced in 2006 – close to 20 states have acted to create regional environmental partnerships.

For example, California and Quebec have launched the Western Regional Climate Action Initiative, with the first joint cap-and-trade auction this month. The carbon allowances sold can be used in either California or Quebec, with reserve prices set at $11.34.

Datacenter operators everywhere should be keeping an eye on initiatives like these, if for no other reason than these programs can increase electricity costs – and thereby increase the short-term cost of doing business.

Metrics Driving Sustainability

In an industry obsessed with metrics, there are several that datacenter operators should be using to evaluate and improve their datacenter's energy and environmental efficiency. Here's a quick rundown of some of the most helpful metrics used to keep operating costs down.

Power usage effectiveness (PUE) is the universally accepted standard, and has more recently been accompanied with *carbon usage effectiveness* (CUE) and *water usage effectiveness* (WUE). By combining all three of these, facility managers are able to get a comprehensive snapshot of their datacenter's overall efficiency and sustainability.

PUE is the measurement of a datacenter's total energy used, divided by the energy specifically used by its IT equipment.

CUE is specific to a datacenter's carbon footprint, and measures its greenhouse gas emissions in relation to its IT energy consumption.

WUE is a site-based metric that assesses the volume of water used for a datacenter's operations divided by the energy consumed by its IT equipment. This includes water used for humidification, on-site energy production and cooling, and related support systems.

Of course, there are a host of other metrics that datacenter operators need to use at any given time. But with government environmental regulations in the works that are likely to raise energy costs, minding these metrics is a must to keep operating costs down.

Above and beyond responding to government regulations, operating an environmentally efficient facility simply makes good business sense.

II. A New Era at Oracle?

The September announcement that Oracle CEO Larry Ellison is stepping down could mark the beginning of a new era – but will it have any impact on the Oracle customer experience?
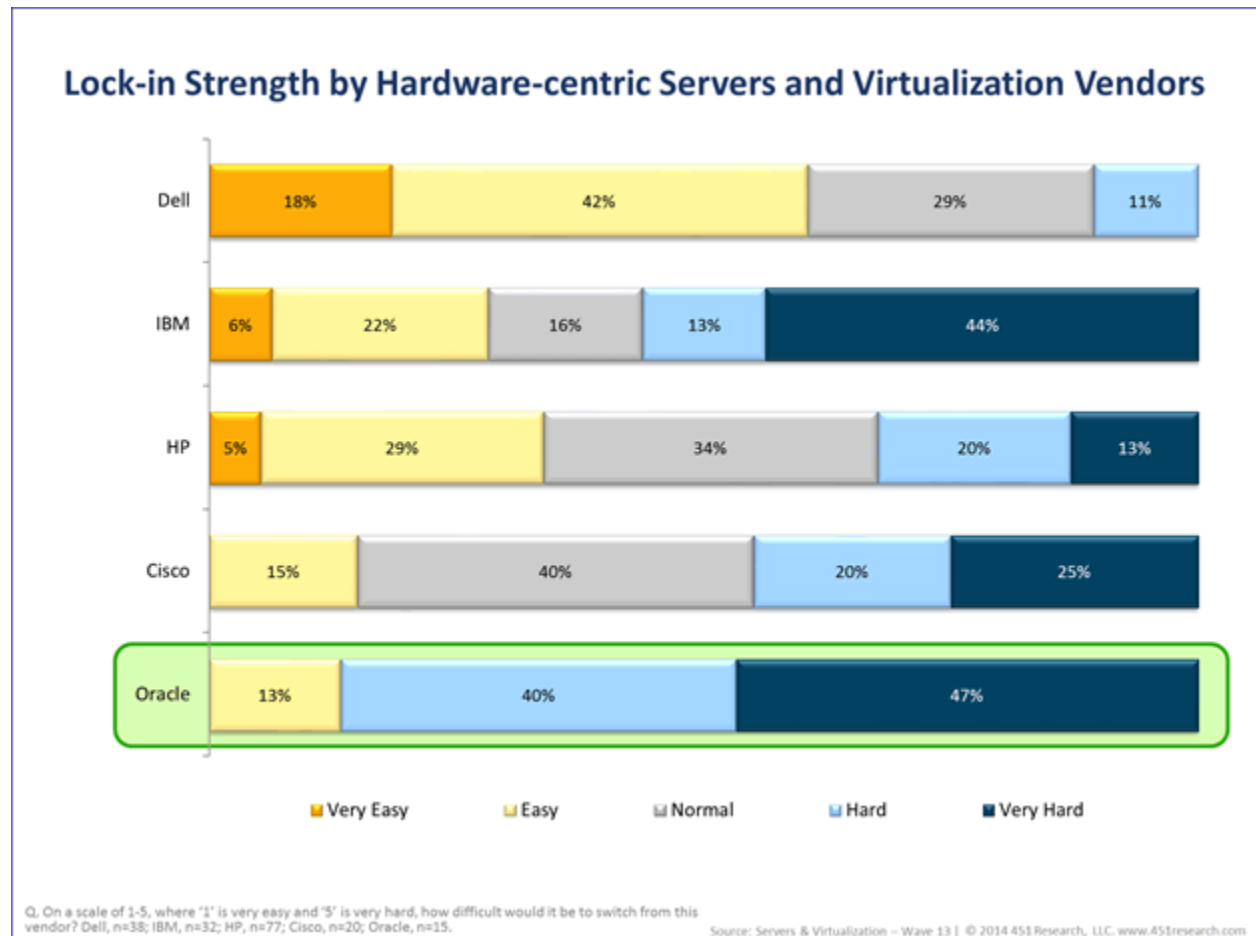
Co-presidents Safra Catz and Mark Hurd are taking the helm as co-CEOs, but Ellison will remain on board as executive chairman and CTO. While the triumvirate will likely stick to the status quo early on, new CEOs Catz and Hurd face multiple existing problems with Oracle's hardware business.

Since acquiring Sun Microsystems in 2010, Oracle has often promised to turn its troubled hardware division around, but has yet to show results. Its hardware business as a percentage of overall revenue has been steadily declining – from 20% of total annual revenue in FY2011 to 14% in FY2014.

Customer feedback from a recent 451 Research Servers and Virtualization study suggests the Oracle decline is at least partially due to the company's perceived "We are Oracle, take it or leave it" attitude.

As one customer puts it, "A lot of it boils down to business support. Oracle practically forces everyone to buy their support, and then charges an arm and a leg for it." Yet another complains, "Oracle forces you to use their products with ridiculous licensing models. The product is solid, but I feel like we are being held hostage by a bully. They want to force us into a relationship that only benefits them."

One point in Oracle's favor is that it enjoys strong vendor lock-in. A huge majority of existing customers (87%) consider it 'hard' or 'very hard' to make a switch from Oracle, which should make them more likely to maintain the status quo. For example, one respondent noted "It'd be hard for us to switch from Oracle hardware, since we have a lot invested in our relationship."

Lock-in Strength by Hardware-centric Servers and Virtualization Vendors

Vendor	Very Easy	Easy	Normal	Hard	Very Hard
Dell	18%	42%	29%	11%	
IBM	6%	22%	16%	13%	44%
HP	5%	29%	34%	20%	13%
Cisco		15%	40%	20%	25%
Oracle		13%		40%	47%

■ Very Easy   ⊔ Easy   ⊔ Normal   ■ Hard   ■ Very Hard

Q. On a scale of 1-5, where '1' is very easy and '5' is very hard, how difficult would it be to switch from this vendor? Dell, n=38; IBM, n=32; HP, n=77; Cisco, n=20; Oracle, n=15.

Source: Servers & Virtualization – Wave 13 | © 2014 451 Research, LLC. www.451research.com
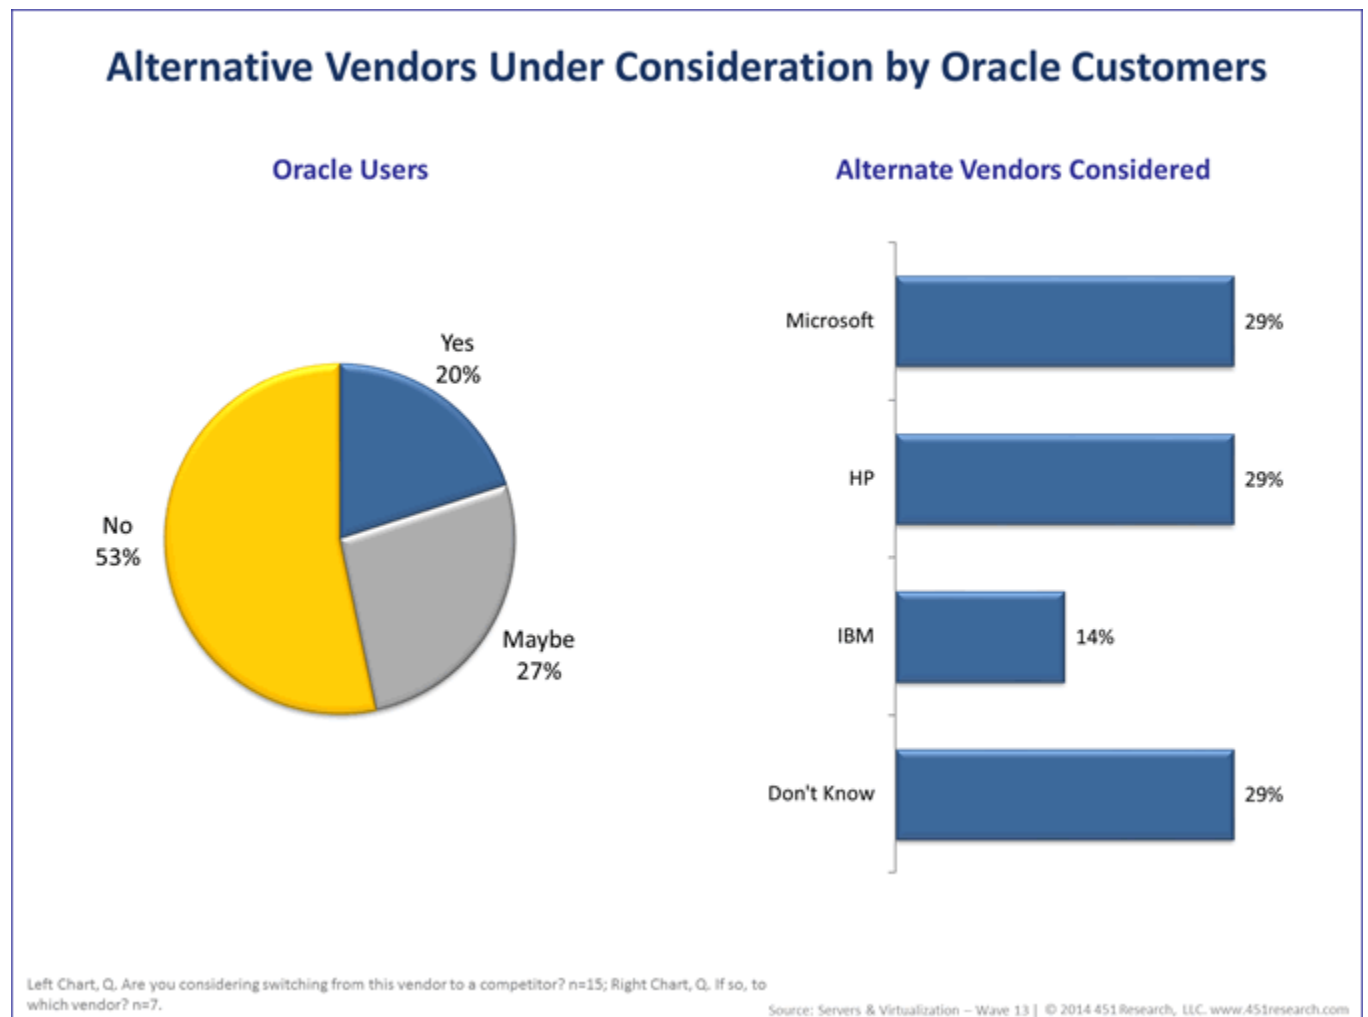
Regardless of its high lock-in strength, Oracle still registered as the most vulnerable vendor in the 451 Research study, with nearly one in two of its current customers (47%) reporting they're considering switching to competitors.

Customer Evaluations

Respondents were asked to rank the vendors they use, and Oracle scored below average in nearly every category. Their lowest score was in interoperability, followed by ease of doing business with the company and their sales force.

In terms of Oracle customers considering switching, HP (29%) and IBM (14%) stand to benefit most on the hardware side, as does Microsoft (14%) on the database side.

Alternative Vendors Under Consideration by Oracle Customers

Oracle Users

Yes
20%

No
53%

Maybe
27%

Alternate Vendors Considered

Vendor	%
Microsoft	29%
HP	29%
IBM	14%
Don't Know	29%

Left Chart, Q. Are you considering switching from this vendor to a competitor? n=15; Right Chart, Q. If so, to which vendor? n=7.

Source: Servers & Virtualization – Wave 13 | © 2014 451 Research, LLC. www.451research.com

Bottom Line: With nearly half of its servers and virtualization customers eyeing alternative vendors, Oracle will be hard-pressed to slow its hardware revenue decline. On the plus side, the company enjoys the highest customer lock-in rate among hardware vendors, and

high lock-in strength does enable it to push some of its customers into higher spending brackets.

Nonetheless, while a new era has begun at Oracle, the old hardware problems remain and the unit continues to become a smaller part of Oracle's overall business.

Clearly, Catz and Hurd need a fresh strategy for the unit going forward. Listening to customers and addressing their existing problems and concerns could well prove to be a more effective long-term strategy than relying on lock-in.

III. DCIM to the Rescue: Tackling Comatose Servers

Recently in Global Digital Infrastructure Trends, we discussed the widespread problem of comatose servers – also known as 'zombies' – which run obsolete software and have no function even as they remain in operation and drain power.

While tools exist to combat the problem they aren't widely adopted, and as a consequence, the zombie server problem persists.

Fortunately, there are key components within a datacenter infrastructure management (DCIM) software suite that can shed light on comatose servers and help solve the problem.

DCIM to the Rescue

Here's a quick rundown on some DCIM tools and processes that can be used to maintain a zombie-free environment:

First, server-level power monitoring. When a server is using power, there are higher-level DCIM tools that track that the server exists. DCIM power consumption monitoring shows how much a server is being used, as well as the cooling it requires.

Like server power monitoring, datacenter asset management is an integral part of most DCIM products. Among other things, it shows a server's specs, its physical location in the datacenter, power use and deployment history.

Importantly, the data stored in DCIM should be integrated with an IT asset management system, as DCIM systems do not yet show what applications are being run nor who is responsible for them – vital information in managing a wide array of servers. IT asset management systems are able to identify servers eligible for replacement by showing their age and power consumption.

Why is this important? Because a paradox can be at play here: Servers that are lightly utilized may look like candidates for replacement, but aren't necessarily comatose. They may even be your newest and most powerful servers that are simply being under-deployed.

But with the right DCIM and IT asset management tools in place, it's possible to gather the real-time and historical data needed to identify and shut down comatose servers.

A Solution to Many Problems

Purchasing DCIM solely to eliminate zombie servers would of course be too expensive a proposition. But the value proposition shifts when the elimination of comatose servers is just one part of your overall DCIM management strategy.

Once an organization's DCIM systems are in place, eliminating comatose servers becomes a relatively straightforward process that can result in substantial savings, both in dollars and power consumption.

And with DCIM in place, a host of other new functions and processes also become possible – such as better capacity planning, dynamic datacenter provisioning, automated load sharing, energy chargeback, risk-assessed workloads, cloud workload deployment – and of course, the elimination of zombie servers.

IV. OpenStack and You: Enterprise IT Spotlight

OpenStack promises what enterprises both love and hate about open source projects: the appeal of flexibility, lower cost and no vendor lock-in, counterbalanced by the fear of hidden costs, an unpredictable pace of product development, and a lack of support for the do-it-yourself open source model.

Despite its benefits, deploying OpenStack requires loads of technical expertise – and enterprises need to be doing it for the right reasons when they decide to wade in.

OpenStack: What You Need to Know

From a broad perspective, enterprise cloud adoption continues to gain momentum.

The OpenStack Road to the Cloud

As one participant in a recent 451 Research cloud study put it, "We're moving to converged infrastructure to get us cloud-ready – combining our network, storage, and computing power for purposes of cloud provisioning and infrastructure as a service."

That said, evolving an organization's overall IT infrastructure into cloud-readiness is normally slow going.

In this regard, cloud platforms that feature OpenStack modular architecture have a number of helpful components including:

- **Compute (Nova) –** a cloud computing fabric controller that's the main piece of an IaaS system. It automates pools of computing resources using any one of a number of possible hypervisors, such as KVM, Xen, Hyper-V and Linux LXC.
- **Object storage (Swift) –** a scalable storage system whereby objects and files are written to multiple disk drives spread throughout servers in the datacenter – and OpenStack software manages data replication and integrity across the cluster.
- **Block storage (Cinder) –** in contrast to Swift, Cinder provides block-level storage devices for OpenStack computing – managing the creation, attaching and detaching of block devices to servers.
- **Networking (Neutron) –** an OpenStack system for managing networks and IP addresses to ensure the network doesn't become a bottleneck to cloud deployment. Neutron gives IT users self-service capability.
- **Dashboard (Horizon) –** provides administrators and users with a graphical interface to access, provision and automate cloud resources for third-party services like billing, monitoring and other management tools.

When thinking of wading into an OpenStack deployment, here are some key considerations to keep in mind:

Know where you are in the maturity cycle. In a recent 451 Research study, we asked the following question:

What phase are you in evolving your internal infrastructure (i.e., cloud-enabling) environment?

One in four (26%) respondents said they were in the automation phase, 11% in the orchestration phase and 17% in the private cloud phase.

So while many organizations understand the need for consolidation and virtualization, the survey data shows that the more advanced phases of cloud enablement – automation, orchestration and private cloud – remain a work in progress.

A second takeaway is that if you are still in the process of getting your virtualization work done, it's probably too early to jump into OpenStack.

Viable use cases. Currently, the most common enterprise use cases we see for OpenStack are on-site private clouds for DevOps (i.e., agile software development and IT operations) and 'big data' technologies such as Hadoop and Cassandra.

If these types of projects are on your organization's radar for 2015, then it makes sense to get an early start on a cloud computing platform like OpenStack.

That said, while OpenStack has been launching new releases every six months, it's important to note that not all of its platform components are evolving at the same pace. Also, OpenStack implementations require a great deal of technical expertise to deploy. Organizations need to have the right expertise in place in order to successfully install and manage OpenStack.

V. Zombie Warfare – The Fight Against Comatose Servers

The IT industry has a problem on its hands: zombies and more zombies. But the "walking dead" in this case are comatose servers that run obsolete software or otherwise have no function – and yet remain in operation and continually drain power.

Comatose servers are a huge waste of resources, whether it be in energy, maintenance or space. And while it may sound like an easy problem to tackle, for many organizations it is definitely not.

Are the Walking Dead taking over YOUR datacenter?!

How come datacenters are able to house large numbers of servers that do very little or nothing? In any other industry this kind of inefficiency would be considered absurd – comparable to an airline flying planes with mannequins in the last 10 rows rather than paying customers.

Of course, no one intends to end up with comatose servers. They were originally installed for good reason. But when applications are retired, virtualized elsewhere or moved to the cloud, no one necessarily arranges for the removal or redeployment of those servers.

Yet even though the technology to locate and prevent comatose servers currently exists, many organizations simply don't make use of it.

How big is the problem?

The Uptime Institute – a division of The 451 Group – has tracked the problem of comatose servers and found it to be widespread. In its 2014 survey of datacenter professionals, nearly one-quarter of respondents estimated that 10% or more of their company's servers were comatose.

Moreover, Uptime analysts estimate that at any given time 20% of racked servers are doing no real work, which, on an industry-wide basis, means billions of dollars wasted.

One example – in just two years, Barclays Bank was able to save $10m by decommissioning more than 9,000 servers, freeing up nearly 600 racks and 2.5MW of power. Similarly, AOL estimates that in 2013 it saved $5.05m in power, software licenses and maintenance, simply by removing comatose servers.

Why the lack of response?

The software for identifying comatose servers has been around for years. For example, 1E's *Night Watchman Server* software measures server power consumption and compares it against usage patterns. That said, the product hasn't been widely successful and is no longer sold as a stand-alone.

Another product, *EnergyCenter*, measures server power use and maps it against server workloads to identify mismatches. Schneider Electric has acquired this technology and integrated it into its DCIM (datacenter information management) software product *StruxureWare* for Data Centers.

But while Schneider will certainly agree that its tools are helpful for identifying comatose servers, it's not a key marketing message for the product. This raises the question, why isn't there a bigger market for products that identify comatose servers?

First, to effectively identify and remove zombie servers requires complicated sets of tools, which makes the cost seem high and the effort disproportionate to the challenge. Many organizations do not perceive their zombie problem to be large enough to justify this investment – and hope they can solve the problem using manual methods.

Second, there's the issue of diffused responsibility. Many companies leave the issue of zombie servers to their IT teams, not the facilities teams that manage space, infrastructure and power. While the entire organization wins from removing zombie equipment, the biggest beneficiaries – facilities teams – often have limited responsibility.

Finally, there's the argument that virtualization and cloud technologies are enabling the separation of physical and virtual layers, shifting these kinds of issues to cloud and capacity management tools, where they can supposedly be more easily solved.

Over the long term, that may prove to be the case, but for now such solutions remain elusive and the problem of comatose servers continues. Nonetheless, with the right tools and processes in place, the battle against zombie servers can be won.

VI. Driving Organizational Change – the Movement to Cross-Functional Teams

Has your company decided to pursue a "service-oriented" approach to IT operations? If so, you're not alone.

In a recent 451 Research survey focused on servers and virtualization, nearly half of respondents (47%) say their company has experienced changes within their server administration teams over the past 12 months.

When we asked the same question in a recent storage survey, only 33% reported changes in their team structures, and for networking it was just 31%.

So why are such a large percentage of server administration teams undergoing change?

First and foremost, widespread virtualization has brought greater efficiency to the area of server administration. In response, IT leaders are trying to maximize the benefits through greater use of cross-functional teams.

New Efficiencies from Cross-Functional Teams

To learn more about cross-functional server teams, 451 Research analysts asked survey respondents in companies deploying them to share their experiences. Some reported that they are witnessing improvements in the time it takes to deploy new technologies, while others are using them to transform into a service-oriented IT model.
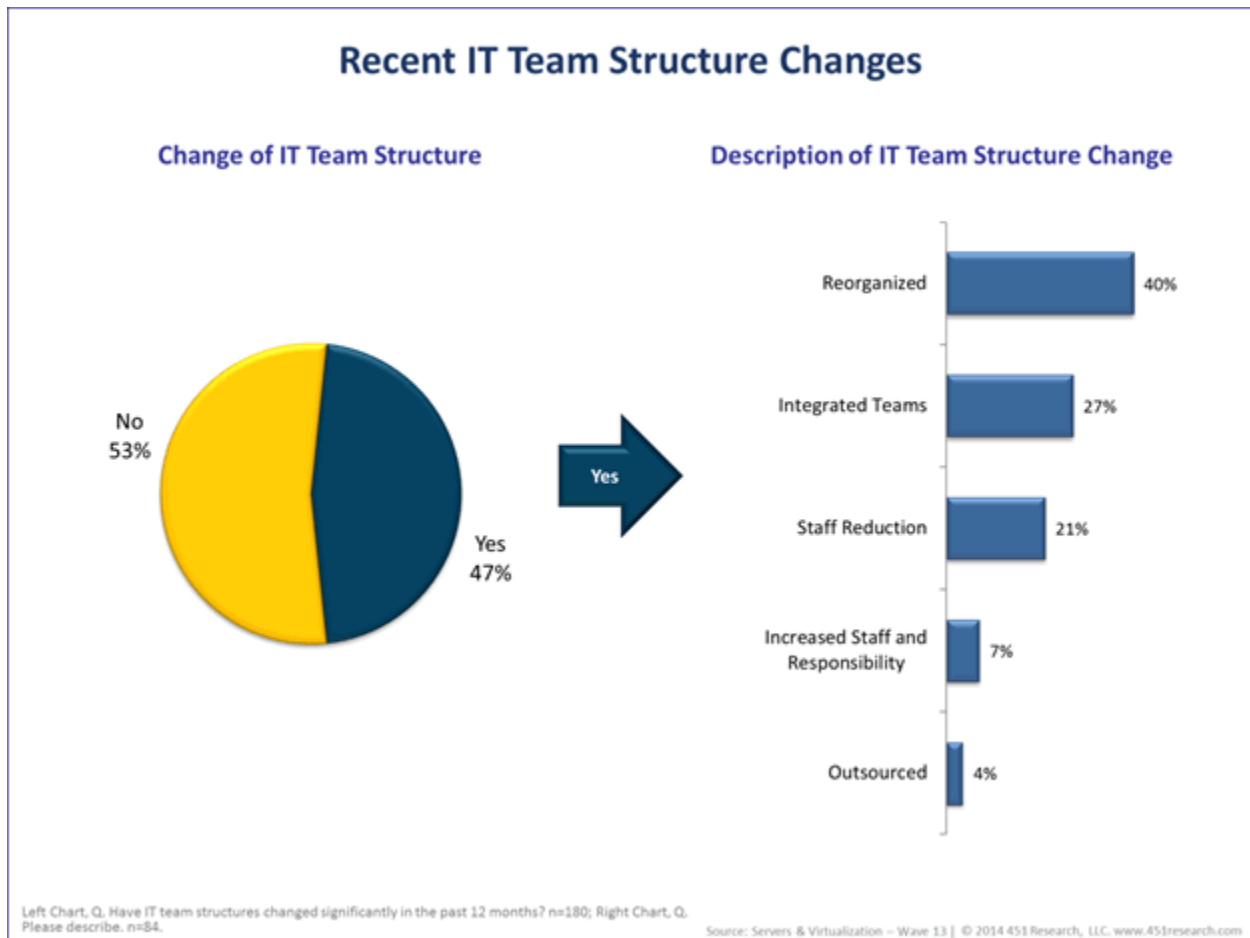
A handful of comments provide some insight:

- "The teams are more service-related rather than technology stacked. It is much more cross-functional."
- "We have a distributed model and have cross-pollinated our teams. Teams are architectural as well as operational. Allows us to rapidly deploy technology."
- "We have to work to the requirements of our colleagues on the networking side – a little difficult because our networking team is outsourced. Same with the

requirements from our storage team. If we could take on some of these roles, we could complete our projects quicker."

- "We have changed from a siloed org to a services-type model."
- "We don't classify as server and storage anymore; people have responsibilities across functional areas."

When we asked study participants to list the changes within their server teams over the past 12 months, reorganization (40%) was by far the most common, followed by team integration (27%) and staff reduction (21%).

Recent IT Team Structure Changes

Change of IT Team Structure

No
53%

Yes
47%

Yes →

Description of IT Team Structure Change

Category	Percentage
Reorganized	40%
Integrated Teams	27%
Staff Reduction	21%
Increased Staff and Responsibility	7%
Outsourced	4%

Left Chart, Q. Have IT team structures changed significantly in the past 12 months? n=180; Right Chart, Q. Please describe. n=84.

Source: Servers & Virtualization – Wave 13 | © 2014 451 Research, LLC. www.451research.com

Change Management

Change management is particularly important when deploying cross-functional teams, where changes in employee roles can create uncertainties that undermine the benefits of new team structures. For example:

- "On the people side, people get scared when we talk about automation. Right now we focus on scripts. We are staging things to be automated. We have to overcome the fears of people – there are different teams involved."

- "In the architecture group, we are automating things for other teams of people, so there are challenges in directing others. Technology is the easy part. The organizational, people and political issues are more difficult."

Automation and a more service-oriented approach imply removing some of the processes traditionally done by server administration staff, which can cause an adverse perception of job security. Open communication is therefore necessary to manage employee expectations and the cultural changes catalyzed by cross-functional teams.

Similarly, enterprises need to monitor their existing talent pool and plan ahead to ensure their skill sets match the requirements for cross-functional teams. Moreover, new roles and expectations have to be communicated, and accountability measures for team members have to be established.

Efficiency Measures

Finally, when we look at the most common efficiency measures being applied to server and virtualization teams today, the top three measures all relate to time – with time to provision (22%) emerging as the most common measure, along with time to issue resolution (21%) and uptime (17%).

Servers managed per full-time employee (16%) is another efficiency measure that scores in the double digits.

Measuring Servers and Virtualization Team Efficiency

Metric	Percentage
Time to Provision	22%
Time to Issue Resolution	21%
Uptime	17%
Servers Managed per FTE	16%
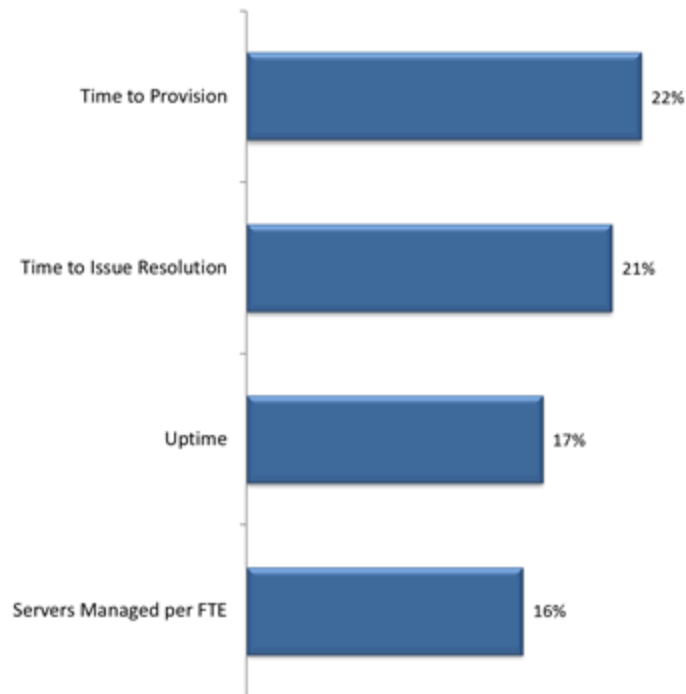
Q. How is the server infrastructure and virtualization team's efficiency measured in your organization? n=179.

Source: Servers & Virtualization – Wave 13 | © 2014 451 Research, LLC. www.451research.com

VII. Threat Intelligence

The term 'threat intelligence' is a popular one in today's security market where nearly every offering seems to come with a side of threat intelligence.

The most valuable threat intelligence offerings uncover data you don't already have in a form your organization can make quick use of. With so many vendors staking claims in this market, it's important for organizations to dig deep

Threat Knowledge is Power

before picking a threat intelligence product they will use internally.

Definitions of threat intelligence vary widely depending on the source, but the product itself tends to have the following characteristics:

- **It's a data-only product – not a combination of security offerings and inputs.** While a firewall consumes threat data that can help generate blocking rules, if the data can't be separated from your firewall, it's not a threat intelligence offering.
- **It's used to analyze and share threat intelligence data only.** Neither enterprise consoles with alerting and reporting nor intrusion-detection systems should be considered threat intelligence products.
- **The user must be able to act on the data as needed.** It can be used, for example, to decide which IP addresses to monitor. Or it can be used in conjunction with forensic software to figure out whether a breach has actually occurred and what else you need to be looking for.
- **Threat intelligence must include data you don't already have.** Monitoring a network and generating alerts doesn't qualify, and simply figuring out that you've been breached is not threat intelligence. But getting external data that tells who may have attacked you and their probable purpose – that's threat intelligence.

Threat intelligence products are designed to take inputs from your existing security products and export data to other security products. They are normally complementary to what an organization already has.

Many threat intelligence vendors aggregate and resell one another's feeds.

VirusTotal's data, for example, is a popular input to numerous threat intelligence offerings. But with so many open source threat intelligence feeds out there, it's important for your organization to do its homework. What one vendor is selling might simply be white-labeled from another. If you're not careful, you can end up buying the same data multiple times.

Players in the Threat Intelligence Space

Nearly every large security vendor now uses 'threat intelligence' to enrich its products and services. Here are some vendor offerings that meet the 451 Research definition of threat intelligence:

- **Threat intelligence feeds** – Well-known players in this area include iSIGHT Partners, Cyveillance, Verisign iDefense, IID, FireEye, Webroot, Anubis Networks, Norse Corporation, Farsight Security, AlienVault, ThreatConnect, CrowdStrike, ThreatTrack, Symantec and Malcovery.
- **Threat intelligence data references** – These vendors make data accessible for querying and reporting through their central consoles. They also accept uploaded malware or logs from customers and match them with their own intelligence. Vendors in this area include ThreatGRID (acquired by Cisco), ThreatStream, Lookingglass and Seculert, as well as many of the vendors listed above.

- **Threat intelligence research and reporting** – Some vendors offer boutique threat intelligence services. For example, Dell SecureWorks researches threats to specific company brands or executives; ZeroFOX monitors customer social media assets; and Bitsight Technologies evaluates the security risk to a customer from its partners.

When evaluating different threat intelligence offerings, the first step is to figure out what kinds of decisions your organization needs to make and how threat intelligence can help you make them. Always think in terms of those decisions, because if you can't use the data to support your organization's internal decision-making, then the intelligence you receive won't prove actionable no matter how good it may be.

VIII. The European Metamorphosis: From Colo to Hybrid Cloud

The adoption of Microsoft's Azure ExpressRoute – a product facilitating private connections between datacenters and the cloud – has resulted in a growing number of partnerships between Microsoft and colocation players that want to diversify their cloud offerings.

This is particularly apparent in Europe, where colo providers such as Equinix, TelecityGroup, Level 3 and BT now offer their customers private connections to the Microsoft Azure cloud.

In doing so, they give a glimpse into how colo players are suddenly offering hybrid IT services involving multiple cloud infrastructures – without becoming cloud providers themselves.

Moreover, such partnerships signal a transformation likely to pick up steam, and may serve as a guide for end users as to where the European IT services market is really headed.

Partnering for the Cloud

As enterprise users express growing interest in hybrid IT services and traditional colo providers look to diversify their offerings, expect to see more such partnerships.

Microsoft worked closely with European colo giant Equinix in designing Azure ExpressRoute, and has based its product on the belief that IT services are clearly moving toward a hybrid model.

In partnering with colocation providers, Microsoft promotes its own cloud offering while simultaneously enabling colos (and other service providers using the colo space) to offer more robust services.

For colo users, Microsoft Azure partnerships offer the benefit of private connections, as well as full control over their cloud environments. It's a proposition that works particularly well for

traditional users of Equinix facilities, which tend to want strong SLAs and high levels of security over cost efficiency.

Colocation and Hybrid IT

Equinix has been preparing for the shift toward hybrid cloud for quite a while, and has set up its International Business Exchange (IBX) datacenters to address the expected demand.

IBXs host communities of cloud service providers that utilize both public and private cloud connections – and enable Equinix to accommodate cloud demand while also maintaining its colocation revenue.

TelecityGroup – which, like Equinix, is a European colo market leader – has been slower to adopt such an approach, but is now rolling out its Cloud-IX platform, which allows its clients to connect directly to a variety of cloud providers, at its European datacenters.

Building an ecosystem like the one Equinix already offers can take years, and it will be interesting to see how fast TelecityGroup can bring in service providers. While it already has some key partnerships – such as Amazon Web Services, VMware and, of course, Microsoft Azure ExpressRoute – it will need a lot more to be able to compete on a scale with Equinix.

That said, TelecityGroup may have benefited by sitting back and waiting for some of the tougher technical challenges surrounding cloud ecosystems to first be addressed.

Even Equinix has admitted that cloud Ethernet – the transmission of data from one datacenter to another in a different geography – is still not operating as efficiently as it should be.

In sum, while these hybrid IT offerings aren't yet perfect, both Equinix and Telecity's established partnerships with cloud providers show they're committed to hybrid IT services as the future of their product offerings.

Given their momentum, it's likely that hybrid IT service offerings will soon be dominating the European market.

IX. Bridging the Gap Between IT and Facilities

Most IT operations teams don't spend a lot of time thinking about datacenter infrastructure management software (DCIM). Rather, DCIM tends to be on the

DCSO: Bridging the Gap
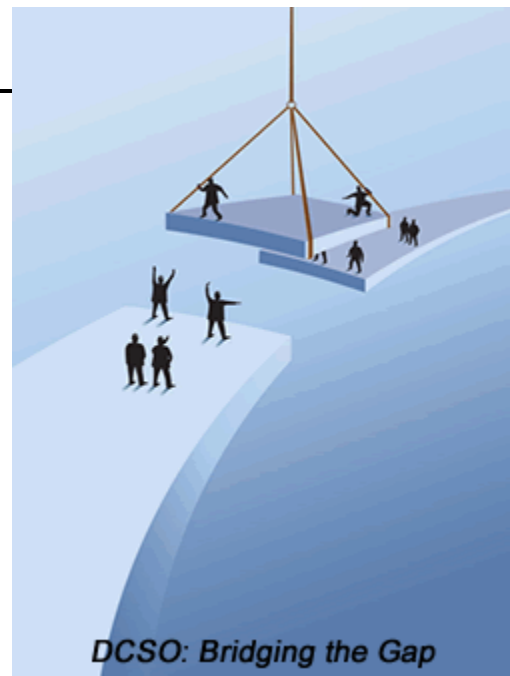
minds of datacenter facilities teams, because traditionally it's been used mostly as a facilities tool.

But that may be starting to change.

The next generation of DCIM software – dubbed Datacenter Service Optimization (DCSO) – is no longer just a data collection and analysis tool. It's become considerably more service-oriented and is now far better able to bridge the gap between the physical and virtual environments.

In other words, IT operations teams are suddenly giving DCIM and DCSO software a good hard look to see if these types of tools can make their operations more efficient.

DCIM, DCSO and the Enterprise

While some DCIM products have been on the market for more than a decade, 451 Research estimates that only 15% of midsized datacenters (3MW and above) currently deploy DCIM software.

So even though DCIM is now a mature market with established players, practically speaking it's still a small market in search of growth opportunities. And one such growth opportunity may lie in the integration of DCIM tools (which monitor power, cooling and other datacenter infrastructure) with those of mainstream IT operations management.

Here are some key facts regarding DCIM and DCSO:

- **DCIM (Datacenter Infrastructure Management Software)** has historically been utilized as a systems management solution for datacenter facilities. DCIM software collects and manages information about physical datacenter assets, resource use and operational status. This data is analyzed and distributed to optimize datacenter performance. A DCIM platform typically includes datacenter power and environment monitoring, asset and capacity management, and predictive analytics and modeling. The focus is on managing the critical physical assets within the datacenter.
- **DCSO (Datacenter Service Optimization Software)** bridges the physical and virtual environments within and beyond the datacenter. While the focus of DCIM is still primarily at the physical level, DCSO extends its capabilities to manage both the physical and virtual assets of the datacenter and across geographically dispersed facilities. DCSO components include core DCIM features plus datacenter service management, energy optimization, business planning and service-based costing.
- **For enterprise IT operators, integrating datacenter facilities with IT systems is the next big frontier.** What makes DCSO important to enterprise IT operators is its growing focus on integrating all types of datacenter facilities with the IT operations and applications that reside there.

  The ability to manage performance and optimize efficiency across all types of datacenters (e.g., private, on-premises, third-party cloud, colo) is a win-win for

everyone involved. For enterprise datacenter owners and cloud/colo providers alike, the extension of DCSO into the traditional IT systems management space lowers cost, drives efficiency and simplifies management.

What should enterprise IT leaders be doing about DCSO?

If you don't operate your own datacenter, aren't planning to run a private cloud, or don't work closely with facilities teams, then DCIM and DCSO doesn't need to be on your radar.

For everyone else, taking a closer look at the technologies and where they're headed is well worth the effort. In coming years, 451 Research expects reseller agreements will be turned into deeper integration partnerships. You can expect the vendors you work with in the IT operations management space to soon be pitching you on broad DCSO solutions.

Here are a handful of action steps organizations should be taking now in preparation for this:

Assess your DCIM/DCSO needs and build appropriate use cases. There are five needs that typically drive enterprise investment in DCIM/DCSO:

(1) Constrained or stranded capacity,

(2) Difficulty tracking assets that reside inside the datacenter,

(3) Problems maintaining availability,

(4) Lack of visibility around energy usage, and

(5) Lack of useful and actionable data.

DCIM and DCSO can help with any and all of these, but some solutions are better than others. The key is for your organization to spend time up-front to figure out which options are best-suited, build your business case, and down the road guide your vendor selection process.

Focusing on collaboration and bridging the datacenter facilities/IT operations divide. DCIM, however well deployed, can't solve every problem in the datacenter. Collaboration between datacenter facility managers and IT operations managers remains absolutely essential – even though they have different priorities, technical skills, preferred suppliers and budgets.

DCSO should be viewed as the crossover product through which IT and facilities share their information and their understanding of shared systems. Establishing a good working relationship with set processes and accountability is a critical part of this.

Identifying best-fit vendors based on your organization's needs – not on vendor features. Not surprisingly, DCIM/DCSO is yet another one of those software categories

where the term means different things to different people. It is highly recommended that organizations first conduct their own detailed feature/function assessment in order to eliminate ill-fitting solutions.

If properly done, you'll narrow the vendor field and ensure that any use of DCIM/DCSO is based on your organization's real-world needs and its own business objectives.

X. The Rise of 'Everything as a Service' – What it Means for OSS/BSS

The OSS/BSS market is evolving quickly, with a growing list of new products and business models accelerating the pace of change.

From the enterprise standpoint, the gold standard lies in identifying and working with OSS/BSS providers that can deliver an enriched customer experience.

The subsequent pressure on providers in the era of 'everything as a service' has left them little choice but to reinvent their operating models and software tools for a range of enterprise OSS/BSS functions – everything from orders to billing to service fulfillment.

OSS/BSS – The Changing Landscape

When we speak of Operational Support Systems (OSS) we refer to the software applications, and occasionally hardware, that support back-office activities. Business Support Systems (BSS), on the other hand, refers to software applications that support customer-facing activities like billing and order management.

The OSS/BSS transformation is well under way, with multiple technology and equipment vendors jockeying for the lead. But whether running in the cloud or on customers' premises, next-generation OSS/BSS architecture is about providing new services and capabilities that operate seamlessly between diverse systems.

New technologies and services, however, bring with them new requirements – and pressure enterprise end users and OSS/BSS providers to undergo major changes to address them.

In the short term, many take a phased approach, but the ultimate goal is likely to be nothing less than a complete OSS/BSS transformation at the organizational and business-function levels.

Jumping on the Bandwagon

A number of big OSS/BSS providers are jumping on the M&A bandwagon, and we expect further consolidation in the next 12-18 months as changing OSS/BSS requirements and the need for new capabilities ripple across the industry:

- Ericsson recently acquired software monetization vendor MetraTech to simplify a range of business functions such as the configuration of billing and settlement management.
- NEC-NetCracker's acquisition of Convergys, CSG's acquisition of Intec, Redknee's purchase of Nokia's BSS business, and Cisco's acquisition of CompTel's UK subsidiary (formerly Axiom Systems) are all M&A activities that are having an impact on OSS/BSS.

So it's good news that providers are putting their financial skin in the game. They're investing in tools and platforms in order to simplify management and improve customer satisfaction.

That said, we're still in the early phases of an OSS/BSS transformation that will take years to complete.

OSS/BSS and the Cloud

With the global economy moving ahead, flexible on-demand cloud service models continue to replace traditional capital-intensive licensing models. But as SaaS usage grows, enterprises face the problem of software from multiple third parties and the huge need to integrate disparate services.

This is where next-generation cloud marketplaces are leaping in, providing more intuitive design and functionalities for such things as pay-per-use licenses, usage reporting and analysis, and customized tier-based management.

Cloud services now routinely provide back-end management portals that track the use of multiple services and create catalogs to appeal to different targeted audiences. In the process, enterprise customers can now manage multiple services and subscriptions – everything from billing and helpdesk ticketing to usage analysis – simply by accessing their user dashboards.

How are OSS/BSS providers embracing the cloud marketplace model? Primarily by leveraging core components of their existing product offerings:

- **Ensim**, for example, has added a storefront component to provide platforms for cloud-marketplace enablement for service providers.
- **Parallels' Automation for Cloud Marketplace** is another example, with Parallels Automation the enabler of the storefront, and subsystems like Parallels Business Automation and Parallels Operations Automation providing core BSS/OSS functions.
- **RES IT Store** from RES Software is that company's latest workspace management tool, providing a storefront-like user experience while working in conjunction with the two existing software components – Workspace Manager and Automation Manager for IT automation.

So there's plenty of enthusiasm around cloud marketplaces and OSS/BSS – some providers now refer to them as Cloud Service Brokerages, or CSBs. But whether cloud or non-cloud, OSS/BSS providers are in an arms race for faster and easier service delivery and a one-stop shop for OSS/BSS transformation.

This is the first of two articles on the OSS/BSS transformation. In an upcoming *Global Digital Infrastructure Trends* Newsletter we'll take a close-up look at the vendor landscape, along with key trends shaping the future of OSS/BSS in the cloud era.