

NetworkWorld

THE CONNECTED ENTERPRISE

INSIDER

➔ Cloud Computing

CLOUD-INSPIRED DATA CENTER TRANSFORMATION



INSIDE

- 2 **Building a Private Cloud:
Get Ready for a Bumpy Ride**
- 4 **Private-Cloud Technologies:
5 Things You Need to Know**
- 5 **How to Build a Private Cloud**
- 7 **Paving the Way for the Flat
Network**
- 8 **4 Trends Shape the
New Data Center**
- 10 **What Data Center Hardware
Should You Keep?**

Building a Private Cloud: Get Ready for a Bumpy Ride

Check your traditional data center mind-set at the door or be prepared to fail

BY BILL CLAYBROOK, COMPUTERWORLD

When cloud computing became a topic of discussion a few years ago, public clouds received the bulk of the attention, mostly due to the high-profile nature of public-cloud announcements from some of the industry's biggest names, including Google and Amazon. But now that the talk has turned into implementation, some IT shops have begun steering away from public clouds because of the security risks; data is outside the corporate firewall and is basically out of their control.

Tom Bittman, vice president at Gartner, said in a blog post that based on his poll of IT managers, security and privacy are of more concern than the next three public cloud problems combined. He also wrote that 75% of those polled said that they would be pursuing a private cloud strategy by 2012, and 75% said that they would invest more in private clouds than in public clouds through 2012.

Frank Gillett, an analyst at Forrester Research, agrees that IT's emphasis is more on private clouds these days. He says that IT managers "are not interested in going outside" the firewall.

Still, as Bittman's blog post points out, private clouds have their share of challenges, too; in his poll, management issues and figuring out operational processes were identified as the biggest headaches. And, of course, an on-premises private cloud need to be built internally by IT, so time frame and learning curve, as well as budget, need to be part of the equation.

Indeed, transitioning from a traditional data center – even one with some servers virtualized – to a private cloud architecture is no easy task, particularly given that the entire data center won't be cloud-enabled, at least not right away.

While we generally think of a private cloud as being inside a company's firewall, a private

Cloud computing is transforming data centers

As today's legacy data centers morph into tomorrow's private cloud infrastructures, enterprise IT executives will have a bevy of architectural decisions to make.

These decisions will come at all levels of the data center, ushered in on the tide of virtualization slowly transforming server, storage and data network infrastructures.

Ultimately, server, storage and network resources all will be abstracted from their underlying hardware and aggregated in shared pools for dynamic, on-demand provisioning.

In these articles, *Network World* and its sister publications *CIO* and *Computerworld* present the latest thinking and explore the latest trends in how data centers are evolving toward the cloud ideal.

cloud can also be off-premises – hosted by a third party, in other words – and still remain under the control of the company's IT organization. But in this article we are talking only about on-premises private clouds.

Also, despite all the hype you might hear, no single vendor today provides all of the software required to build and manage a real private cloud – that is, one with server virtualization, storage virtualization, network virtualization, and resource automation and orchestration. Look for vendors to increasingly create their own definitions of private cloud to fit their product sets.

Virtualization is only part of the picture

Many IT managers equate a private cloud with virtualization. What they describe is usually virtual infrastructure, meaning that "you can treat your servers, storage and networks as a single pool of resources that workloads can request on demand," explains Tony Iams, vice president and senior analyst at Ideas International, a research firm with emphasis on enterprise IT infrastructures.

But virtualization and the cloud are not the same thing; to be considered a cloud, the architecture must be set up to provide both

orchestration and automation on top of the virtualization layer.

Orchestration is the coordinated delivery of many types of resources, such as processors, storage and networks, to provide an integrated provisioning process; resources can be delivered in minutes rather than days or weeks. In other words, a single command or request causes a number of actions to occur, possibly in a specific sequence, to coordinate the provisioning request.

The whole point of a private cloud is to allow IT managers to reduce costs and provide so-called agile provisioning rather than just making management of the infrastructure more convenient. A private cloud with virtualization underpinnings turns the technology infrastructure into a pool of resources that can be provisioned on demand with minimal manual intervention.

Without a focus on delivering IT services, it's unlikely you'll attain the full benefit of private clouds.

Comparisons to traditional data centers

In a traditional data center setup, "every time you add a server, somebody has to walk to a firewall console, set up firewall rules, attach

the server to a VLAN, set up load balancing” and do many other tasks, explains Jeff Deacon, cloud computing principal at Verizon Business. But a private cloud needs minimal human intervention other than bringing in new computers or storage to keep up with demand. In a cloud environment, there is one console that lets operators set parameters to automate the entire process, rather than requiring IT personnel to log into different consoles for security, networking and server OS functions.

Another of the main differences between private clouds and traditional data centers involves IT processes. Private clouds may require a re-architecture of how data is used, and processes may have to be rewritten.

For example, today many IT organizations have to contend with sets of requirements that must be met in the provisioning process for budget; discussions with the storage, network and server groups; and tons of paperwork. This type of process is in stark contrast to the streamlined, short-duration provisioning done in clouds. Time to provision may go from weeks in the traditional data center to minutes in a cloud.

You may also have to re-architect deployment of legacy applications to take advantage of private clouds. Many legacy applications are running on mainframes and proprietary Unix platforms. Most virtualized environments, including private clouds, are geared to run on x86-based systems.

Also, in a virtualized environment, we generally don’t know exactly where an application is running at any given time. Because most legacy applications are tied to a specific platform, running them in a private cloud will often require re-architecting them.

Divorcing applications from the hardware is a hallmark of clouds, including private clouds. In a traditional data center, those 10 servers over there might be running billing applications and those five over there running CRM apps. With a private cloud, however, it’s not known ahead of time which servers are running which specific applications. The applications run on whichever servers have free cycles at the time the apps need to run.

Private clouds involve two groups: the IT (data center) operations folks and the business users who want to run applications. A private cloud gives business users the opportunity to quickly provision a server and run an application when they want to – without human intervention.

The IT operations folks have to make sure that sufficient resources are available for the type of on-demand computing that business

users have heard is available with public clouds, and that usually means that the wait for user-requested resources is minutes, not days. Anything short of this, and users will be unhappy.

This is what private clouds are all about: providing the on-demand elasticity of public clouds but doing it within the company’s firewall.

Another difference is that some IT managers or business users may expect private clouds to act like public clouds. In a public cloud, the public cloud provider’s IT operations group is responsible for the compute infrastructure, and the enterprise’s business application groups manage and monitor their own applications, under agreement, on the public cloud. If the private cloud is expected to operate in a similar manner that means the IT group may need to give up its traditional application-management role.

The steps for transition

The first thing is to broaden out beyond server virtualization. At this point, a lot of users are looking at virtualization for purposes of

availability. So look at those aspects of virtual infrastructure that improve availability as the next steps toward a private cloud.

Today, people are integrating storage with virtualization and are beginning to understand the impact of broad virtualization of resources, Iams of Ideas International says. “When we get to the stage where virtualization of servers is the rule rather than the exception and most workloads are virtualized, this is the stage in which virtualization gets woven into the operational process,” he says.

When you get to this stage, you have to rethink what this does to your storage processes, Iams says. “For example, how does virtualization affect backup and recovery?”

Iams outlines the following steps for creating a private cloud:

- Virtualize your storage and try to achieve the same flexibility with storage that you already have with virtualized servers.
- Coordinate server virtualization and storage virtualization with management tools such as Windows Azure Storage

Hurdles involved

Building your own private cloud involves some challenges, including these:

- **Budget.** Private clouds can be expensive, so you need to do your due diligence and figure out what the upper and lower bounds for your ROI will be.
- **Integrating with public clouds.** Build your private cloud so that you can move to a hybrid model if public cloud services are required. This involves many factors, including security and making sure you can run your workloads in both places.
- **Scaling.** Private cloud computing services usually don’t have the economies of scale that large public cloud providers provide.
- **Reconfiguring on the fly.** You may have to tear down servers and other infrastructure as it is working to move it into the private cloud. This could create huge problems.
- **Legacy hardware.** Leave your oldest servers behind – you should not try to repurpose any servers that require manual configuration with a private cloud, since it would be impossible to apply automation/orchestration management to these older machines.
- **Technology obsolescence.** The complexity and speed of technology change will be hard for any IT organization to handle, especially the smaller ones. Once you make an investment in a private-cloud technology stack, you need to protect that investment and make sure you stay up to date with new releases of software components.
- **Fear of change.** Your IT team may not be familiar with private clouds, and there will be a learning curve. There may also be new operational processes and old processes that need to be reworked. Turn this into a growth opportunity for your people – stress of doing and learning all this may be mitigated by helping your folks keep in mind that these are important new skills in today’s business environment.

Management and VMware vStorage.

- Virtualize your network infrastructure and, again, coordinate that with your management tools.

Your infrastructure has been fully virtualized when you have server virtualization, storage virtualization and network virtualization. The crossover point from a virtual infrastructure to private cloud comes when you have the management tools that treat all three types of resources – servers, storage and networks – as a single pool that can be allocated on demand.

Of course, all this is from a technology-centric point of view. Iams says that there is a parallel track that relates to the transition from an organizational perspective, including people, processes, governance, policy and funding. One key question: What does a private cloud structure do to budgets and financial flow within an organization?

Public clouds require users to pay only for what they use. Because a private cloud does not provide users with a fixed amount

of capacity like they may have had with a traditional data center, chargeback is almost certain to be an integral part of private cloud environments.

Virtualization expert Bernard Golden views chargeback as very important because price is an important rationing mechanism – and rationing computing resources will be more important in an environment where obtaining resources is as easy as filing out a Web form.

Few, if any, companies go through all of the above steps/stages in parallel. In fact, there is no single “correct” way to transition to a private cloud environment from a traditional data center. A private cloud is in part the logical conclusion of server virtualization where it is extended to storage and networks and then managed with tools that treat servers, storage and networks as a single pool of resources. Automation and orchestration tools are the key to moving from a virtualized infrastructure to a true private cloud.

But one thing is very clear: If your IT organization is not willing to make the full investment for whatever part of its data center is

transitioned to a private cloud, it will not have a cloud that exhibits agile provisioning, elasticity and lower costs per application.

As part of the transition, you need to determine whether your staff has the experience and skills required for a private-cloud environment or whether you need to hire someone who has been involved in building private clouds.

How you get started depends on your existing infrastructure. If you already have server virtualization, you have a definite advantage over those who do not. Most important: Do not rush out and buy a ton of software from vendors, especially from a single vendor, without a plan in place.

Original article at Computerworld.com

Claybrook is an analyst with more than 30 years of experience in the computer industry, specializing in Linux, open source, virtualization and cloud computing. He is president of New River Marketing Research in Concord, Mass., and holds a Ph.D. in computer science.

Private-Cloud Technologies: Five Things You Need to Know

Many CIOs are contemplating establishing private-cloud technologies for their company. Here are five things to consider before you make the move.

BY DAVID F. CARR, CIO

Your cloud may not be a cloud. Private cloud computing goes beyond rebranding virtualization. Virtualization is an enabling technology, but public cloud applications are attractive because of their flexibility and pay-as-you-go nature. What are you doing to bring those characteristics within your firewall? If your infrastructure is virtual, but you're managing virtual machines the same way you've always managed servers, maybe your cloud isn't really a cloud.

You must be able to respond to fluctuating capacity needs. Your private cloud should provide for rapid provisioning and re-provisioning of capacity. Public cloud services accomplish this with large server farms.

Within a corporate network, you won't be able to justify having thousands of servers sitting idle, but you should have a way of quickly reallocating the capacity you have when user needs demand it.

You need charge-back mechanisms. You should be able to measure and charge for use of cloud servers and services. With cloud services, you only pay for what you use, and you pay more when you use more. Firms that lack a strong central services organization with charge-back mechanisms are likely to be organizationally incapable of implementing a fair charge-back system, says Brad Tagg, a cloud consultant and former IBM Distinguished Engineer.

You must decide who controls access, you or the users. You will have to decide where to draw the line with self-provisioning versus centralized control. Anyone with a

credit card can sign up for an account on a public cloud service in a matter of minutes. What would the equivalent look like within your organization? Do you want employees or departments provisioning their own user accounts, or assigning themselves big blocks of storage?

Now's the time to plan for future options. How you architect your private cloud affects how you expand. For example, some service providers are adopting VMware standards while Amazon Web Services use open-source Xen virtualization technology. Kevin Smilie, a partner at the consulting firm TPI, says establishing the right service-management discipline will put you in a better position “when you need a burst of capability beyond your own firewall.”

Original article at CIO.com

How to Build a Private Cloud

Expert advice on how to approach an on-premises cloud, from conception to implementation

BY BETH SCHULTZ, NETWORK WORLD

If you're nervous about running your business applications on a public cloud, many experts recommend that you take a spin around a private cloud first.

But building and managing a cloud within your data center is not just another infrastructure project, says Joe Tobolski, director of cloud computing at Accenture.

"A number of technology companies are portraying this as something you can go out and buy – sprinkle a little cloud-ulator powder on your data center and you have an internal cloud," he says. "That couldn't be further from the truth."

An internal, on-premises private cloud is what leading IT organizations have been working toward for years. It begins with data center consolidation, rationalization of OS, hardware and software platforms, and virtualization up and down the stack – servers, storage and network, Tobolski says.

Elasticity and pay-as-you-go pricing are guiding principles, which imply standardization, automation and commoditization of IT, he adds.

And it goes way beyond infrastructure and provisioning resources, Tobolski adds. "It's about the application build and the user's experience with IT, too."

Despite all the hype, we're at a very early stage when it comes to internal clouds. According to Forrester Research, only 5% of large enterprises globally are even capable of running an internal cloud, with maybe half of those actually having one, says James Staten, principal analyst with the firm.

But if you're interested in exploring private cloud computing, here's what you need to know.

First steps: Standardization, automation, shared resources

Forrester's three tenets for building an

internal cloud are similar to Accenture's precepts for next-generation IT.

To build an on-premises cloud, you must have standardized – and documented – procedures for operating, deploying and maintaining that cloud environment, Staten says.

Most enterprises are not nearly standardized enough, although companies moving down the IT Information Library (ITIL) path for IT service management are closer to this objective than others, he adds.

Standardized operating procedures that allow efficiency and consistency are critical for the next foundational layer, which is

developers asking for "three VMs of this size, a storage volume of this size and this much bandwidth," Staten says. Self-service for end users seeking resources from the internal company cloud would be "I need a Share-Point volume or a file share."

Thirdly, building an internal cloud means sharing resources – "and that usually knocks the rest of the companies off the list," he says.

This is not about technology. "It's organizational – marketing doesn't want to share servers with HR, and finance won't share with anybody. When you're of that mindset, it's hard to operate a cloud. Clouds are highly inefficient when resources aren't shared," Staten says.

Faced with that challenge, IT Director Marcos Athanasoulis has come up with a creative way to get participants comfortable with the idea of sharing resources on the Linux-based cloud infrastructure he oversees at Harvard Medical School (HMS) in Boston. It's a contributed hardware approach, he says.

At HMS, which Athanasoulis calls the land of 1,000 CIOs, IT faces a bit of a unique challenge. It doesn't have the authority to tell a lab what technology to use. It has some constraints in place, but if a lab wants to deploy its own infrastructure, it can. So when HMS approached the cloud concept four years ago, it did so wanting "a model where we could have capacity available in a shared way that the school paid for and subsidized so that folks with small needs could come in and get what they needed to get their research done but also be attractive to those labs that would have wanted to build their own high-performance computing or cloud environments if we didn't offer a suitable alternative."

With this approach, if a lab bought 100 nodes in the cloud, it got guaranteed access to that capacity. But if that capacity was idle, others' workloads could run on it, Athanasoulis says.

"The reason why people use our cloud... is because it provides compelling value to them – and that's not a bad place for IT to be."

-- Marcos Athanasoulis, IT Director, Harvard Medical School

automation. "You have to be trusting of and a big-time user of automation technology," Staten says. "That's usually a big hurdle for most companies."

Automating deployment is probably the best place to start because that enables self-service capabilities. And for a private cloud, this isn't Amazon-style in which any developer can deploy virtual machines (VM) at will. "That's chaos in a corporation and completely unrealistic," Staten says.

Rather, for a private cloud, self-service means that an enterprise has established an automated workflow whereby resource requests go through an approvals process.

Once approved, the cloud platform automatically deploys the specified environment. More often, private cloud self-service is about

"We told them - you own this hardware but if you let us integrate into the cloud, we'll manage it for you and keep it updated and patched. But if you don't like how this cloud is working, you can take it away." He adds, "That turned out to be a good selling point, and not once [in four years] has anybody left the cloud."

To support the contributed hardware approach, HMS uses Platform Computing's Platform LSF workload automation software, Athanasoulis says. "The tool gives us the ability to set up queues and suspend jobs that are on the contributed hardware nodes, so that the people who own the hardware get guaranteed access and that suspended jobs get restored."

Don't proceed until you understand your services

If clouds are inefficient when resources aren't shared, they can be outright pointless if services aren't considered before all else. IBM, for example, begins every potential cloud engagement with an assessment of the different types of workloads and the risk, benefit and cost of moving each to different cloud models, says Fausto Bernardini, director IT strategy and architecture, cloud portfolio services, at IBM.

Whether a workload has affinity with a private, public or hybrid model depends on a number of attributes, including such key ones as compliance and security but others, too, such as latency and interdependencies of components in applications, he says.

Many enterprises think about building a private cloud from a product perspective before they consider services and service requirements - and that's the exact opposite of where to start, says Tom Bittman, vice president and distinguished analyst at Gartner.

"If you're really going to build a private cloud, you need to know what your services are, and what the [service-level agreements], costs and road maps are for each of those. This is really about understanding whether the services are going toward the cloud computing style or not," he says.

Common services with relatively static interfaces, even if your business is highly reliant on them, are those you should be considering for cloud-style computing, private or public, Bittman says. E-mail is one example.

"I may use it a lot, but it's not intertwined with the inner workings of my company. It's the kind of service moving in the direction of interface and independence - I don't want it to be integrated tightly with the company. I

want to make it as separate as possible, easy to use, available from self-service interface," Bittman says. "And if I've customized this type of service over time, I've got to undo that and make it as standard as possible."

Conversely, services that define a business and are constantly the focus of innovative initiatives are not cloud contenders, Bittman says. "The goal for these services is intimacy and integration, and they are never going to the cloud. They may use cloud functions at a low level, like for raw compute, but the interface to the company isn't going to be a cloud model."

Only once you understand which services are right for the cloud and how long it might take you to get them to a public-readiness state will you be ready to build a business case and start to look at building a private cloud from a technology perspective, he says.

The final tiers: Service management and access management

Toward that end, Gartner has defined four tiers of components for building a private cloud.

At the bottom sits the resource tier comprising infrastructure, platforms or software. Raw virtualization comes to mind immediately, but VMs aren't the only option - as long as you've got a mechanism for turning resources into a pool you're on the way, Bittman says. Rapid re-provisioning technology is another option, for example.

Above the resource pool sits the resource management tier. "This is where I manage that pool in an automated manner," says Bittman, noting that for VMware environments, this is about using VMware Distributed Resource Scheduler.

"These two levels are fairly mature," Bittman says. "You can find products for these available in the market, although there's not a lot of competition yet at the resource management tier."

Next comes the service management tier. "This is where there's more magic required," he says. "I need something that lets me do service governance, something that lets me convert pools of resources into service levels. In the end, I need to be able to present to the user some kind of service-level interface that says 'performance' or 'availability' and have this services management tier for delivering on that."

As you think about building your private cloud, understand that the gap between need and product availability is pretty big, Bittman says. "VMware, for example, does a really good job of allowing you to manage your virtualization pool, but it doesn't know

anything about services. VMware's vCenter AppSpeed is one early attempt to get started on this," he adds.

"What we really need is a good service governor, and that doesn't exist yet," says Bittman.

Sitting atop it all is the access management tier, which is all about the user self-service interface. "It presents a service catalog, and gives users all the knobs to turn and lets you manage subscribers," Bittman says. "The interface has to be tied in some way to costing and chargeback, or at least metering - it ties to the service management tier at that level."

Chargeback is a particularly thorny challenge for private cloud builders, but one that they can't ignore for long. "It's tricky from a technology perspective - what do I charge based on? But also from political and cultural perspectives," Bittman says. "But frankly, if I'm going to move to cloud computing I'm going to move to a chargeback model so that's going to be one of the barriers that needs to be broken anyways."

In the end, it's about the business

And while cloud-builders need to think in terms of elasticity, automation, self-service and chargeback, they shouldn't be too rigid about the distinctions at this stage of cloud's evolution, Bittman says. "We will see a lot of organizations doing pure cloud and a lot doing pure non-cloud, and a whole lot of stuff somewhere in the middle. What it all really comes down to is, 'Is there benefit?'"

Wentworth-Douglass Hospital, in Dover, N.H., for example, is building what it calls a private cloud using a vBlock system from Cisco, EMC and VMware. But it's doing so more with an eye toward abstraction of servers and not so much on the idea of self-provisioning or software-as-a-service (SaaS), says Scott Heffner, network operations manager for the hospital.

"Maybe we'll get to SaaS eventually, and we are doing as much automation as we can, but I'm introducing concepts slowly to the organization because the cloud model is so advanced that to get the whole organization to conceive of and understand it right off the bat is too much," he says.

As HMS' Athanasoulis says, "The reason why people use our cloud ... is because it provides compelling value to them - and that's not a bad place for IT to be."

Original article at Network World

Schultz is a longtime IT writer and editor in Chicago.

Paving the Way for the Flat Network

Is it time for data center networks to shed a tier?

BY BETH SCHULTZ, NETWORK WORLD

All of a sudden, data center networking has gone from ho-hum to hot.

The status quo – that hierarchical, three-tier network that has dominated the data center since the late 1990s – is being threatened, as enterprises move toward a virtualized, service-oriented, scaled-out, converged infrastructure.

The question facing enterprise IT architects and planners today is whether to stick with three tiers or to flatten out the network in order to improve performance and ultimately save money.

From the vendor perspective, companies like Brocade, HP and Juniper Networks are pushing hard for the flat network. After all, if enterprises decide to go flat, they will need to buy new data center switches and that opens up an opportunity to displace Cisco, which for years has been all but synonymous with the data center network.

“Even companies with Cisco Catalyst switches have to re-qualify their data centers and, to go into flat network mode, they’ll have to look at all brand new equipment. This is a huge opportunity for others to come and take the data center network away from Cisco,” says Andre Kindness, senior analyst for enterprise networking at Forrester Research.

Cisco, of course, has next-generation plans of its own, built around its Nexus switches. And going head-to-head against industry powerhouse Cisco might also backfire against some of the more vulnerable network vendors.

Kindness says the survivability of some of these vendors could even be at stake. As servers, storage and the network converge, if these companies don’t get themselves into the data center network, the rest of their businesses could be affected, he adds.

From the IT side, the allure of the flat network is being able to deploy products purpose-built for automation, convergence and

virtualization. The promise is a much simplified, fabric-based architecture from which enterprises will enjoy dramatic performance improvements as well as streamlined operational chores, and expenses.

“The data center network hasn’t been this exciting in probably 10 years,” Kindness says. And enterprise network planners are sitting up and taking notice.

“This is a brand-new game, with the latest and greatest new technology and everybody has their specialties. The data center network definitely warrants a fresh look. We can’t rely on old assumptions anymore,” says John Turner, director of networks and systems at Brandeis University, in Waltham, Mass.

“Clearly all the switch vendors are racing to keep up with what’s happening in the data center – and nobody has it quite right just yet, Turner says. “Some of these products are literally being released to customers with the paint still wet on them.”

Fun, but no games

Jim Metzler, vice president of IT consulting firm Ashton, Metzler & Associates, says industry trade show Interop is a barometer of how the buzz around flat architectures has put the excitement back into networking. In 2008, Metzler organized a session entitled, “Is There Anything Left to Say about the Local Area Network.” The answer: Not a whole lot.

This spring, Metzler moderated a session he called, “Why Networking Must Fundamentally Change.” “I had 212 people in the audience and vendors lined up angry that they couldn’t get on the panel. Why? – because networking is fun again,” he says.

Perhaps so, but flattening the network is no mere amusement, experts warn.

“Truth is, if we look at data center networking trends from the demand side, meaning how the structure and mission of data centers are changing, then there’s definitely a logic to

rethinking how we look at data center networks,” says Tom Nolle, president of CIMI, a technology assessment firm.

But network executives would be wrong in thinking a flat network is always best for the new data center, he adds.

“What really matters is what’s happening in the data center, such as an increase in inter-processor communications, and how that affects the behavior of traffic and require-

“The data center network hasn’t been this exciting in probably 10 years.”

-- Andre Kindness, senior analyst for enterprise networking at Forrester Research

ments placed on the network,” Nolle says. “If you decide to flatten your network without looking at the entire chain, then you run the risk of either investing prematurely or in the wrong thing,” Nolle says.

Many enterprises, caught up in the marketing furor around re-architecting, will end up undertaking and then aborting flat network projects over the years, Nolle says he fears.

“As we start to promote the notion of flattening the data center network, we’re taking what should be a project-based decision and pulling it out of that context,” he explains. “If a project is expected to cause performance issues, and you’re considering flattening the network to solve those problems, than that’s a reasonable decision. But the notion of going out and flattening the network just because it’s the cool thing to do is like buying a car because you think it’ll make you socially successful.”

Flattening the network will require a major investment in new equipment. Network executives must be able to tell a good financial

story about how flattening the network will increase revenues or reduce costs outside of the data center. That'll be hard for the average enterprise to do, Nolle says.

"The benefit case won't prove out ... and flat network projects will be aborted. Hopefully this will happen early enough in the cycle that relatively little will have been procured," he adds.

Price and performance rule

To flatten or not needs to come down to a price/performance evaluation.

"At end of the day, I don't care if the data center network is two or three or however many tiers. We've got to get our arms around the cost," Metzler says.

"You have to remember why we went to three tiers in the first place," says Robin Layland, head of Layland Consulting, which specializes in network architecture and strategy for

enterprise customers. "We couldn't get the density we needed to with two tiers. It was just too expensive, so we went to three tiers. That was the cheapest way to get the performance we needed. There was no inherent goodness in a three-tier architecture," he adds.

Data centers are changing, but the decision points for network managers remain the same. "If a two-tier network can give you great performance at a better price than a three-tier design, it wins because theoretically it should have lower latency and even better reliability too," Layland says.

The decision doesn't need to be hurried, but does require a careful study of options and setting of expectations, Kindness says.

"This isn't going to happen for two to five years out and even then not as one big changeover at once," he says. "This is going to be an evolution, and no one should expect their vendors - server, storage or networking

- to enable this to happen automatically."

"None of these vendors is going to be able to produce the whole thing end to end perfectly in the beginning," Layland says.

So in the short run, they'll differentiate themselves on how they handle the new networking requirements brought on by server virtualization. "It'd be nice to use all possible connections at least from the servers to the first switch or top of rack, because that's where the problem first shows up and where we have a lot to gain. Plus, the first storage integrations will take place at those top-of-rack switches, too - it won't mix all across the network at first," he adds.

"We'll have to see who can pull this off and how well they handle the execution," Layland says. "This is where you see the most opportunity for proprietary magic."

Original article at Network World

Four Trends Shape the New Data Center

IT execs adapt to a new reality as x86 virtualization transforms the data center forever

BY BETH SCHULTZ, NETWORK WORLD

Thanks to x86 server virtualization and its follow-on technologies, the state-of-the-art enterprise data center looks vastly different than it did even a year ago.

And moving from old school to next-generation isn't just about hardware and software - it's a call for a new way of thinking about the data center, as well.

"Some people are so accustomed to one application, one server and a methodology that locks you in to one way of thinking that they're having a hard time fully understanding the new data center," says Bill Fife, director of technology for Wholesale Electric Supply Co., in Houston.

"But now with thin replication and replays and synchronization to disaster recovery sites, and virtual machines being able to move files from data store to data store and having multiple data stores on the server, and adding network adapters, you really have to

sit back and think about how you want to run your operations and remember that you have options. You're not tied down to any one path. You can go down one road today and change directions tomorrow," Fife says.

Here are four of the major trends in today's data center:

Trend No. 1: I/O virtualization

At Wholesale Electric Supply, Fife is capitalizing on the ability to virtualize I/O, one of the latest of several significant technology trends shaping the new data center.

I/O virtualization, also known as I/O aggregation, splits interconnections across either 10-gigabit InfiniBand or Ethernet links. Xsigo Systems' virtual I/O Director uses the former and Cisco's Nexus 5000 and 7000 switches the latter, for example.

"In either case, you connect this pipe and then you can get as many virtual Ethernet and Fibre Channel connections as you want out of it," says Logan Harbaugh, an independent analyst and member of the Network World

Lab Alliance. "The architectures are similar, as there's a limit to how much they can vary and still provide some level of functionality."

I/O virtualization simplifies the hardware scenario in the data center rather considerably, reducing the number of connections running to each device while increasing flexibility. Take into consideration VMware's best practices recommendation that you assign 1G port per virtual machine (VM). With newer 24-core servers, you could theoretically run at least 24 and maybe as many as 50 VMs on a single piece of hardware, which in turn would mean needing 50 1G ports, Harbaugh says.

Realistically, even if you could get six four-port Ethernet boards, you'd still only be able to support 24 VMs. "The nice thing about I/O virtualization is that everything shares the one InfiniBand or 10G Ethernet connection as lots of 1G pipes."

At Wholesale Electric, Fife is using Xsigo's virtual I/O Director to decouple processing, storage and I/O. "By doing so we've essentially built our own cloud because we can assign processor, RAM, disk and I/O on an

as-needed basis, and then, when they're no longer needed, get rid of it all and do something else," he says. "There are no rigid guidelines within which we have to operate. We can be extremely flexible."

Trend No. 2: Data and storage convergence

Today's data centers typically have distinct data and storage networks, and nobody much likes that situation. "As soon as people can recombine those two networks, that's what they're going to do," says Joel Snyder, senior partner with consulting firm Opus One and another member of the Network World Lab Alliance.

"My belief and, yes, hope is that we'll get rid of pure Fibre Channel and go to Fibre Channel over Ethernet [FCoE] - but I still see people buying a lot of Fibre Channel because they're told it's the way to go, even though our tests actually show that the network often isn't the bottleneck," he says. "What you can do with Fibre Channel you can do with 10G Ethernet and get equivalent or better performance, even if that's not the belief of SAN buyers and vendors."

This is early days for FCoE, but plenty of folks are looking at the technology, says David Newman, president of Network Test, an independent test firm, and Network World Lab Alliance member. If nothing more, the technology has cost in its favor, he says.

"Besides the capital cost of the equipment, there's the operational expense issue. People who run plain old Ethernet cost less than people who know Fibre Channel," Newman says. "On economic grounds, it'll be cheaper to provision FCoE than running separate infrastructures."

Today, Brocade and Cisco have FCoE-capable switches that fully support all prioritizations and new mechanisms on Ethernet for delivering Fibre Channel-like service levels, and other vendors are coming into the fray, as well. So building a working, end-to-end FCoE network that handles data and storage is possible today - at least using the same vendor's gear, Newman says. Interoperability is unproven as yet.

Scott Engel, director of IT infrastructure, Transplace, a third-party logistics provider in Dallas, identifies FCoE as one of the two biggest networking and infrastructure changes coming to the company's data center over the next year. The other is 10G to the servers, he says.

Indeed, Newman says, the real tipping

point in the data center will happen over the next 12 to 18 months when 10G replaces 1G Ethernet on server motherboards. "That'll have all sorts of follow-on effects, enabling data-storage convergence is just one," he says.

Watch for this year to be the first with "appreciable numbers" of 40G switch ports shipping, Newman says. Fatter network pipes will be needed to accommodate the higher-speed server connections.

Trend No. 3: Faster processors, greater consolidation

By now, most enterprises have server consolidation stories to share, spun around a virtualization theme. They tell of impressive physical-to-virtual server ratios, often in the double digits. But consolidation in the data center is just beginning, some say.

The maturity and comfort levels around virtualization are growing, which means enterprises are showing the willingness to put more and more VMs on a single system, says Steve Sibley, an IBM Power Systems manager. Within the year, he adds, the Power 750 will support up to 320 VMs on a single server, the Power 770 and 780 up to 640 VMs and plans for up to 1,000 VMs.

The ability to support higher numbers of VMs per physical server comes on the back of faster processors, of course. In IBM's case, the company recently introduced the Power7, an eight-core chip that delivers four times the virtualization capability, scalability and performance than its predecessor, Sibley says. The high-end Power7-based Power 780 and 770 servers will come with up to 64 Power7 cores, for example.

Intel, too, has readied an eight-core chip, the Nehalem-EX.

"If you start at the chip level, the ability to deliver more performance per processor core but also pack four times as many cores onto a single chip gives a vast amount of new capacity and capability to put more virtual servers onto a single platform without sacrificing performance or capability of the overall system," Sibley says. "That design point is enabling systems or offerings that give clients the ability to consolidate even more than they used to on a single platform at much cheaper prices than ever before."

Trend 4: Infrastructure optimization

Will your data center strategy one day include

a semi tractor-trailer full of hands-off gear parked in some spot selected for optimal cooling and power supply?

Dan Kusnetzky, vice president of research operations at The 451 Group, says he can imagine so - at least as one potential alternative to building out new or extending existing data centers. "Software routes around failures, and maybe you'd replace that truck with a new one every three years or so," he says.

The data center-in-a-box concept is one that bears watching, agrees Doug Oathout, vice president of converged infrastructure at HP. Companies already are using data centers like pods or trailers outside their facilities, optimizing server, storage, networking, cooling and power distribution resources for that size container, he says. "Now we see the performance-optimization trend moving inside the data center."

This is not to say the data center is going to turn into a parking lot full of semis. But enterprises that run out of space, electricity, cooling and capacity today can take the container concept and move that type of asset inside the data center, Oathout says. "We're not talking about the container itself, but the concept, being able to say 'I need eight racks of servers, four racks of storage, a rack and half of networking, and here's the power and cooling it will consume,' and optimize that way."

Piecing together a data center section by section is far less costly than the traditional go-for-broke approach, and delivering power and cooling a section at a time is far more efficient than moving it across a long distance, Oathout says.

"There's so much more waste when you build a data center to the ultimate capacity vs. building it to what it needs to do, so you could almost call this a retrofitting trend," Oathout adds. "I'm going to optimize what I've got, doing it with localized power, cooling and energy for the specific work I want to get done in this environment. Then I take the next step, with multiple pods, instantiations or building blocks within the data center. It's mindboggling how much more efficient that is compared to building a monolithic data center that has mega watts and 100,000 square feet of space yet is incapable of supporting the equipment you need for your next workload."

Original article at Network World

Cloud Computing:

What data center hardware should you keep?

As your enterprise adopts cloud solutions, it can be hard to determine what systems should stay and what should go.

BY VANESSA ALVAREZ, CIO

Enterprises today have so much to think about when designing their overall data center strategies. It's critical that a holistic approach be taken when laying out a data center design. The challenge for many CIOs and their IT organizations today is working with their partner to determine the best approach. First, what to do with the existing infrastructure in place?

With virtualization, many organizations rid themselves of physical servers, despite the significant dollar investment, and went to virtualization. Although many continue to do virtualization in a more phased approach, the understanding was that long term, virtualization not only provided organizations with cost savings, but also better performance, agility and flexibility, to name a few benefits. The point is, virtualization was not just about the cost savings, but about the value adds that it brought to the overall IT environment and ultimately, the business.

Similar challenges are being faced today. While enterprises are making the move to deploying cloud computing models in their data centers and leveraging external public clouds, it becomes difficult to determine what to keep and what to let go, particularly in terms of physical infrastructure.

True, driving cost out of IT is the goal, but what about the investment already made? It comes down to the old "rip and replace" story that has plagued so many emerging technologies in the past. Should an organization really consider getting rid of its infrastructure? What happens to all that investment made? Worse, what happens to the IT talent that managed it? It has always been a taboo for any IT executive to even consider this option.

However, in the past few months, I've heard more and more stories about CIOs and IT executives making the decision to take out particular on-premises solutions and instead go with an as-a-service solution, whether it be hosted or managed services. In interviewing a number of executives, it started to make sense. The cost of maintaining and managing old infrastructure and the human resources that go along with doing that, in the long run, actually becomes more expensive than just pulling it all out and moving to a more cost-efficient deployment model. Their only challenge was making the case to the CEO.

However, CEOs today expect their CIOs and COOs to work together and make the necessary IT decisions that will enable further growth for the business. In several interviews, the executive indicated that the numbers worked out, so the CEO approved moving forward and



starting to take out much of that on-premises infrastructure.

Each case is different; in many cases, it's still a difficult decision to make. Ultimately, most infrastructure costs are major investments made within organizations and CEOs still are looking to extract as much value as possible.

The challenge is being able to determine this: What is the POTENTIAL value of moving from on-premises to off-premises? Does the cost of managing and maintaining on-premises (including human/capital resources) outweigh the cost of simply taking it out and moving to cloud? In most instances, particularly with respect to legacy infrastructure, it's surprising to see that it makes sense.

Original article at CIO.com

Alvarez is an industry analyst with Frost & Sullivan focusing on monitoring and analyzing emerging trends, technologies and market dynamics in the area of enterprise infrastructure in North America.