United Business Media

Data Center In A Box
Liquid delivers a total package p.34

Event Management
Prism provides high-end features on a budget p.39

Size Matters
But there's more to choosing hardware p.42

Network Computing

A SPECIAL ALL-DIGITAL ISSUE

For IT By IT

November 2009

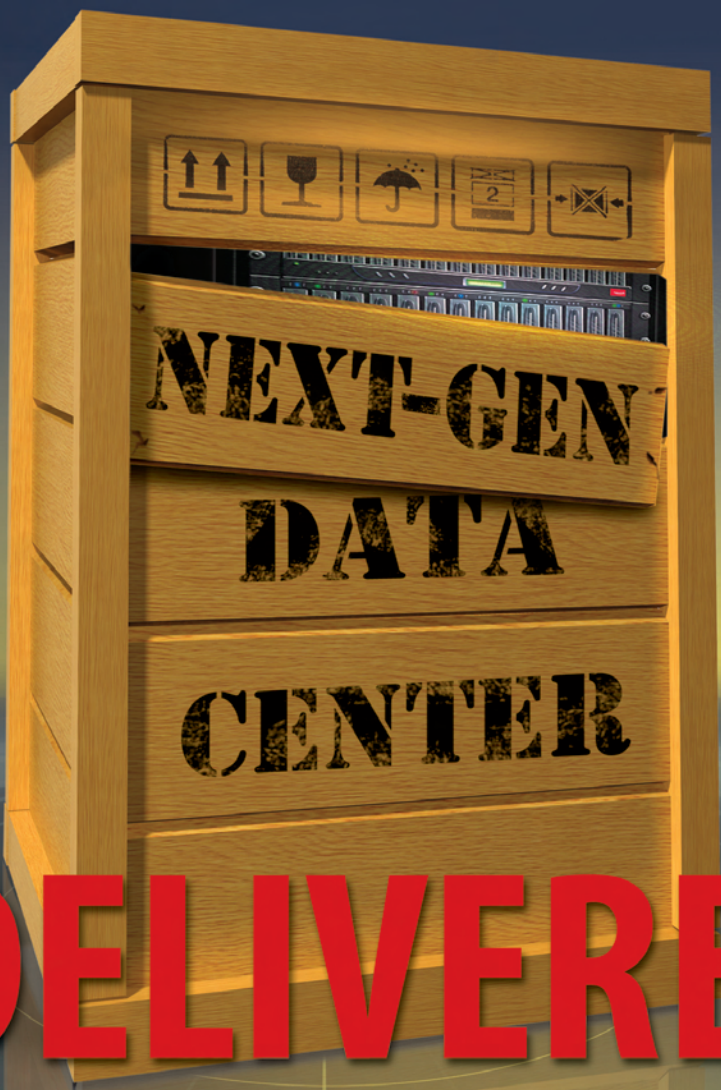

DELIVERED

Unified systems combine servers, networking, storage, and management and put IT back in control. p.12

Green Dreams
High efficiency is good for business, provides stellar PR —and is within reach p.26

Network Computing
For IT, By IT

CONTE

>> *Network Computing*'s second all-digital issue

Cover Story

NEXT-GEN DATA CENTER, DELIVERED

The unified computing platforms of today promise to consolidate everything and anything into a single chassis. And the vendors with the biggest R&D budgets are succeeding at this task. p. 12

ENERGY-EFFICIENT DATA CENTERS

Was your data center designed in the last five years? If not, it's not as efficient as it could be. Brush up on the basics of effective designs. p. 26

INSIGHTS

News Bits And Bytes

HP trades punches with Cisco; Juniper has a new software platform, hardware, and partners; businesses stumble on encryption; Q&A with HP's Paul Congdon. p. 8

REVIEWS

A True Data Center In A Box

Early adopters of the LiquidIQ virtual server say it's reduced setup and reconfigure time, and let them cut back on high-priced staff. p. 34

EventTracker Adds Much-Needed Monitoring

The latest version of Prism's security incident and event management offering makes up for old shortcomings. p. 39

<< Previous

N T S

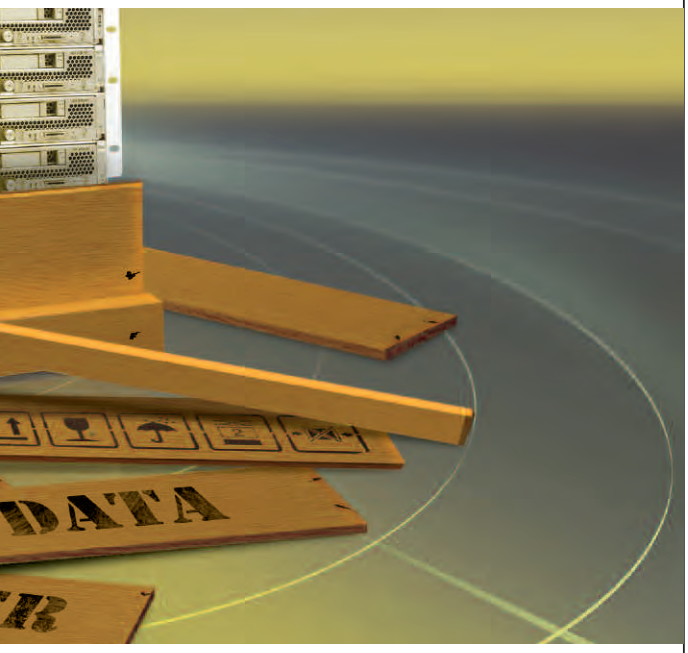

COLUMNS

Preamble
Utility computing can be useful—but take time for a serious talk with the vendor about potential problems. p. 4

Checksum
As virtualization matures and servers pack in more memory, commodity hardware is more attractive. p. 42

NetworkComputing.com

AROUND THE WEB

Open Source You Can Use
Two ways to browse the Web, and one great way to find everything scattered across all your storage media.
informationweek.com/nwc/opensource1109

Cloud Computing's Potential
Will companies move their core business apps to the cloud? It's one of the great unanswered questions, and one reason NetSuite is so interesting to watch.
informationweek.com/nwc/netsuite

IN DEPTH REPORTS

Coming Spectrum Crisis?
CIOs need to identify network-centric trouble spots today, as new wireless opportunities arise and well before anything resembling a "crisis" hits.
informationweek.com/nwc/spectrum

Security Software As A Service
This report offers advice on how to cut through the hype and claims by SaaS security vendors so you can get the best fit for your business.
informationweek.com/nwc/mssp

Tools For The Modern IT Organization
To boost agility, many financial services technology organizations are exploring tools like cloud computing and SaaS. informationweek.com/nwc/modernit

Next-Gen Data Center, Delivered p. 12
Energy-Efficient Data Centers p. 26
Table Of Contents p. 2

Don't Swap One Management Problem For Another

Mike Fratto

Utility computing systems are cool, but before committing to one, make sure you know what you're getting into

When I ran *Network Computing*'s lab at Syracuse University, the network and servers were critical production systems that had to be kept running. If something broke, we had to fix it or get tech support to fix it. This was a product testing lab with a highly dynamic environment, and we were often our own worst enemy in terms of system stability.

We did what we could to keep things efficient, like making server images and restoring them after testing. But no matter what we did, 20% of our time was spent on installs and another 10% to 20% on maintenance. I considered adding more automation, but that seemed likely to just shift the management burden from hardware to software.

One of my goals was to never open a rack. Every time someone opened a door, the chance of knocking something loose loomed, and more important, cabling runs become a mess of orphaned cables, mocking us whenever we strung new cable. You know what I mean.

The utility computing systems that Randy George talks about in "Next-Generation Data Center: Delivered" (see p. 12) and the LiquidIQ system Joe Hernick tests in "A True Data Center In A Box" (see p. 34) would have been just the ticket. A fully racked and cabled system where I could dynamically provision servers, networking, and I/O with a multiterabyte SAN would have significantly reduced test-bed setup time. Sure, we cut provisioning time from days to hours with server imaging software and a flat, consistent network platform, but too often we had to spend time in our noisy, cold data center, running cables and accessing consoles.

To make a system like Liquid's work, you need orchestration software that automates all of the fiddly tasks to provision a new server. It includes multiple integrated systems that provide server management, application deployment, runbook automation, configuration

Next >>

DESIGNED FOR EFFICIENCY.

BUILT TO REDUCE TOTAL ECONOMIC IMPACT.

Achieve greater efficiency and help reduce total economic impact with the latest Dell™ PowerEdge™ M-Series Blade Servers. Built to consume less energy and space, these new servers can help:

- STREAMLINE MANAGEMENT
- REDUCE ENERGY USE AND TOTAL COOLING REQUIRED
- SPEED DEPLOYMENT
- MINIMIZE SERVER SPRAWL
- ENHANCE I/O VIRTUALIZATION

Powered by the Intel® Xeon® processor 5500 series, new PowerEdge servers deliver a range of options with easy configuration, flexible services, and scalable external storage with Dell's PowerVault™, EqualLogic™, or EMC lines of storage arrays.

REDUCE DATA CENTER COMPLEXITY, SPRAWL, AND COST WHILE INCREASING EFFICIENCY

LEARN MORE AT DELL.COM/EfficientDataCenter

Next >>

management, and network and storage management. Orchestration systems use templates that administrators select and customize. Press the "go" button, and within minutes your servers are deployed, booted, and waiting for you to log in. That works great in a homogeneous environment, but few data centers are homogeneous so typically some part of the orchestration ends up being manual.

Given that reality, my fear is that orchestration systems simply shift the management burden from servers to the orchestration software. Take the case of enterprise network management systems. They're expensive and need a small army of experienced administrators to integrate the components and maintain the fragile system. Orchestration systems could suffer similarly. When working properly, they're great. When they crumble, they could cost you as much downtime as any hardware failure. An orchestration software failure could cascade through all of the integrated management subsystems and possibly corrupt the configuration management data.

We had to rebuild our server inventory more than once because of corruption. At the very least, a failure can leave you in a state where you can't or don't want to make changes outside the orchestration system. Making such changes can make recovery more difficult.

If you must make changes, then have a back-out plan in case something goes wrong. It's all about defining and adhering to change management. Utility computing is useful and cool, but before committing to a system, I'd have a real heart to heart with the software vendor about potential failures and recovery steps, and I'd want demos of both. And be sure to talk to peers from companies using the product in production to get their take.

Mike Fratto is editor of Network Computing. Write to him at mfratto@techweb.com.

Next >>

TechWeb Digital Library

techweb.com/digital_library

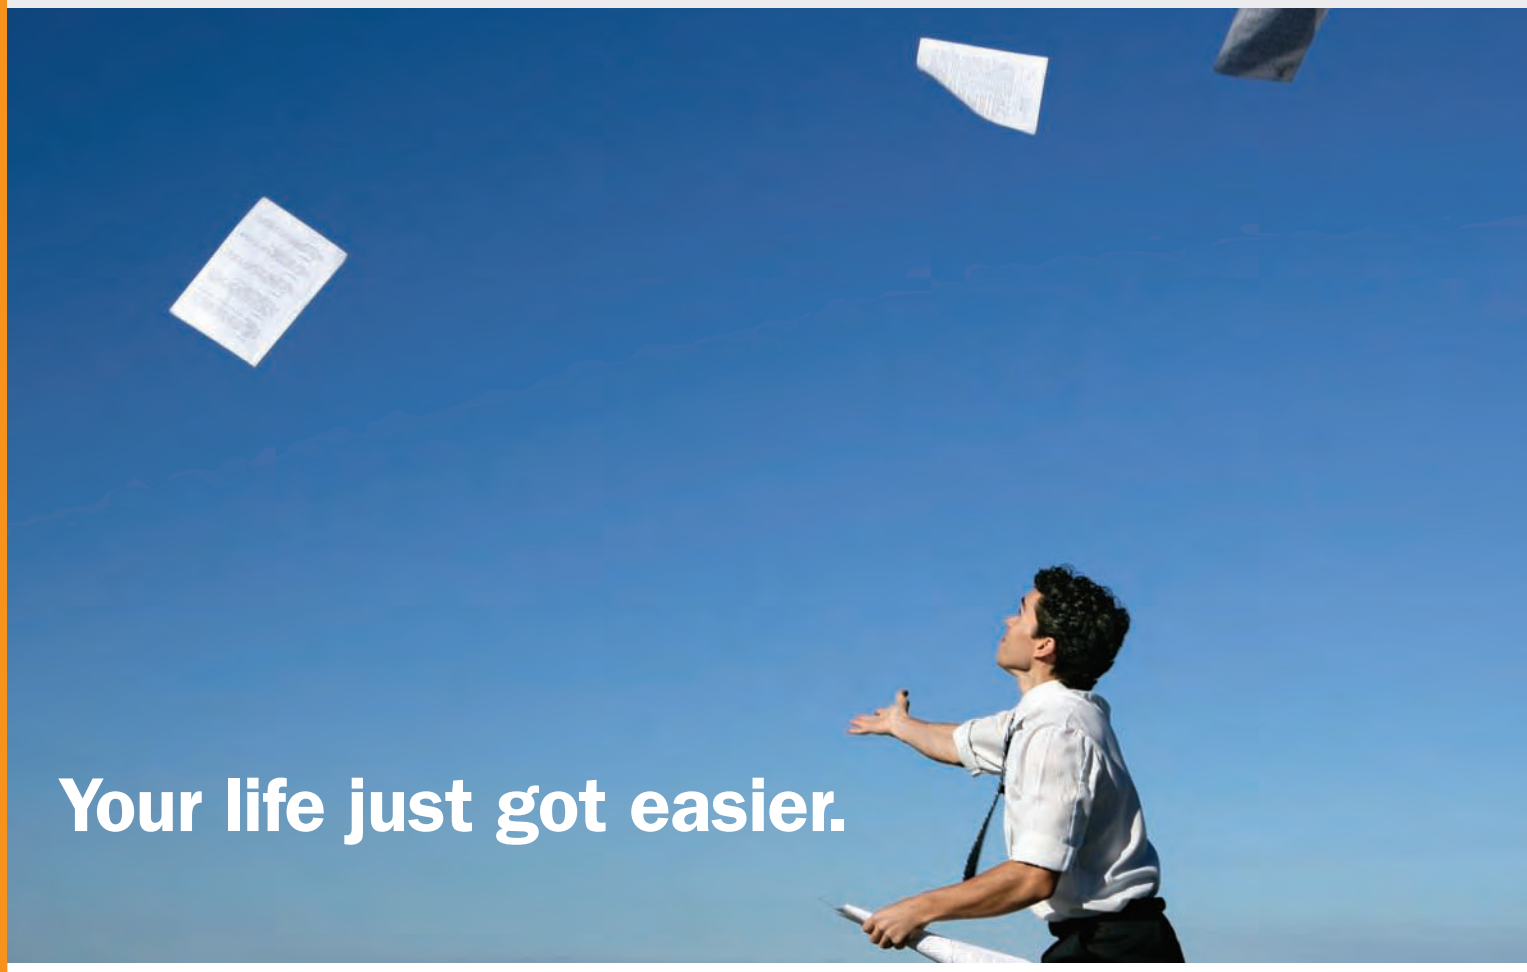

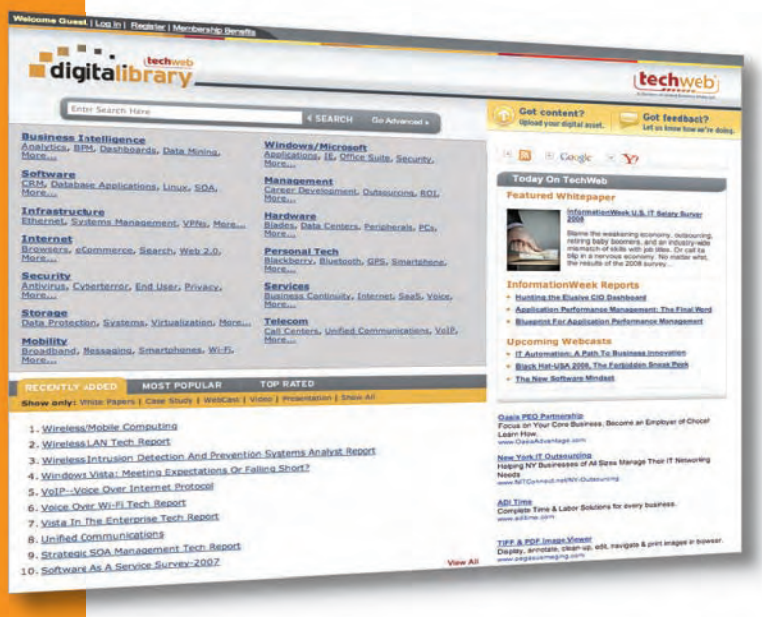

Introducing the **TechWeb Digital Library**, the premier research application for demanding IT decision makers. **We're out to revolutionize the way you research and evaluate business technology.**

Access comprehensive buying intelligence from your peers, technology vendors, independent research houses, leading industry events and TechWeb's business technology editors. Search and organize thousands of data sources according to your criteria. The **TechWeb Digital Library** is your go-to knowledge base to research your options, develop your short list, and make the right buying decisions.

Change the way you choose technology.
techweb.com/digital_library

techweb
A Division of United Business Media LLC

techweb digitalibrary

Next >>

IN THIS ISSUE

Data Center Strategy

HP, Cisco Trade Punches

Once Hewlett-Packard acquires 3Com, in a $2.7 billion deal expected to close by Q2 2010, the company will have a packed portfolio that sets the stage for a heavyweight brawl with Cisco in the enterprise data center.

While Cisco has a sweeping vision in its Data Center 3.0, with the 3Com acquisition, HP becomes a significant threat by being able to do it all, from servers and networking to storage and management and security, via 3Com's TippingPoint division. As with any such assimilation, the devil is in the management. IT groups will want seamless integration of the disparate parts. And, both HP and Cisco will need to watch out for Juniper, which is inking deals with Dell and IBM. ***—Mike Brandenburg***

BY the NUMBERS

{31%}

OF IT PROS SAY their organizations use just enough encryption to comply with regulations. Heartland Payment Systems was PCI compliant when it was breached, and yet it had to pay nearly $32 million in fines. Now Heartland's CEO Bob Carr wants end-to-end encryption beyond what PCI requires.

Data: *InformationWeek Analytics*

Juniper's New Products And Partners

Juniper's staked out its position in the data center switch market with a recent rollout of software, products, and partnerships.

Junos Space is new software that provides integration capabilities with Junos operating system. Junos Pulse combines its SSL VPN, Universal Access Client, and WAN optimization clients for better deployment and management. And Juniper's Trio chipset offers a massive jump in router performance over its previous chipsets.

Partnerships, such as OEM deals with IBM and Dell and a deal with Blade Network Technologies to put the Junos OS on Blade's blade chassis switches, drive at the heart of rivals like Cisco and Hewlett-Packard. Missing from Juniper's plans is storage networking. While companies deploying converged Ethernet are few and far between, Juniper's lack of Fibre Channel over Ethernet could hamper its momentum.

—Mike Fratto

{69%}

OF ORGANIZATIONS AREN'T USING CITRIX XEN SERVER, while 64% are using VMware extensively. But both VMware and Citrix have cause to worry: 40% of IT administrators say they have limited deployments of Microsoft's up-start Hyper-V. Not bad for a product that's barely more than two years old.

Data: *InformationWeek Analytics*

IN THIS ISSUE

A Higher Level Of Data Center Management

As the vice chairman of the IEEE 802.1 working group, Hewlett-Packard ProCurve CTO Paul Congdon is working on the data center bridging standard, which seeks to make Ethernet the single converged fabric in the data center. He recently spoke with sister site *InformationWeek.com* editor in chief Alexander Wolfe about how the bridging standard will help with data center management.

Network Computing: Why are we seeing tighter couplings between networking and server providers?

Congdon: It's due to the convergence of data and storage over Ethernet, but the network is clearly not as standardized as the server component. Look at the storage piece. How well does a Brocade switch talk to a Cisco switch? At the edges, they talk to initiators and targets well, but switch to switch? The Fibre Channel world has never been too great with its interoperability. There's always been a fair amount of lock-in with storage. Nothing changes with FCoE except the physical layer.

Network Computing: Is there tension between vendors wanting to be

INTEL FACES N.Y. ANTITRUST SCRUTINY

New York Attorney General Andrew Cuomo filed a federal antitrust lawsuit against Intel, saying the chipmaker used "bribery and coercion" to prevent computer makers from using rival products that threatened the company's market dominance.

In 2006, Intel paid Dell nearly $2 billion in "rebates," the suit says, more than its net income that year. The lawsuit claims Intel and Dell worked together to market microprocessors and servers at prices below cost in order to deprive Intel rival AMD of sales.

—*Antone Gonsalves*

Forty-four percent of Interop attendees plan to increase spending on technology in 2010, signaling a tech sector recovery, according to a pre-show survey.

a sole solution and working together to develop interoperable standards?

Congdon: We've spent the last 20 years decomposing that mainframe into a bunch of bits and parts, and now it's a bunch of virtual parts. In the process, we've given customers product choice. But the sprawl from that decomposition has created a management challenge. So we're trying to collapse back this disaggregated system into something that customers can manage at a higher level.

Network Computing: So is that what the Data Center Bridging group is—a single, converged standards effort under the IEEE?

Congdon: That's absolutely one of the motivations here. We recognize there's an ecosystem of vendors—hypervisors, NIC vendors, switch vendors, management software, storage, and compute—needed. Our position in the Edge Virtual Bridging group is that this distributed architecture gives you the best of both worlds as compared to a unified-computing-type system, which is very centralized and proprietary, requiring traffic to traverse into the core.

Read the full interview at networkcomputing.com/congdon.

WLAN VENDORS CHANGE PARTNERS

It's time to take a closer look at who's paired with whom among the wireless LAN OEM and switching vendors. The Enterprise Mobility group at Motorola recently became the OEM WLAN for Extreme Networks and has also displaced Meru Networks to become Brocade's wireless partner. While these wins are good for Motorola, they're inevitably a concern for customers who have invested in the prior technology.

The lesson for enterprise IT managers is that OEM relationships come and go—so think hard about who you get your products from.

—Mike Brandenburg

IN THIS ISSUE

Next-Gen Data

Moore's Law, server virtualization, and some fancy I/O processing and management software are conspiring to deliver on the promise of blade server computing

With a bit more than a year to go in the decade, we can count on one hand the number of true innovations that have transformed IT through the aughts. We're a skeptical lot, loath to accept new risk, particularly when it means substantially changing how we do our jobs.

Consider for a moment the upward trajectory of software as a service. In the early part of the decade, no one had heard of Salesforce.com, and few business technology organizations would consider buying software under a subscription-based licensing model. Salesforce.com is now a juggernaut, and IT shops have overwhelmingly accepted SaaS as a viable way to license software. Even Microsoft is preparing to offer up Office in the cloud.

Then there's virtualization, which as a technology is as old as the second generation of IBM mainframes. Four decades later, it has transformed the x86 architecture and how we deploy operating systems and applications.

What makes for a transformative technology is a real solution to a burning business problem. With SaaS, enterprise apps with 90% of the functionality of their on-premises kin could be provisioned almost instantaneously with little or no capital expenditure. With server virtualization, IT was able to collapse the data center footprint, save power,

Next >>

By Randy George

Center, Delivered

Next >>

IN THIS ISSUE

provide for fast, cost-effective disaster recovery, and offer near-instant provisioning of additional computing resources as they're needed. The benefits and value of each technology were so apparent, so transformational, that acceptance was just a matter of time.

Today, we are looking squarely at another innovation that you will soon see making its way to your data center. You've probably heard the buzzwords already—unified computing, utility computing, agile computing, cloud computing. Those words sound like a marketing gimmick, but what they really represent is the continued collapse of the traditional data center well beyond just server consolidation. The unified computing platforms of today promise to consolidate everything and anything possible into a single chassis, and month by month, the vendors with the biggest R&D budgets are succeeding at this task.

Tip 1 Take It Slow

Go slow, please. Your results will definitely vary compared with Tutor Perini's herculean accomplishment. While James McGibney's team did much of the consolidation itself, you can safely assume that Tutor Perini got some velvet glove service from Cisco. Bringing a unified computing platform into your environment isn't a whole lot different from plopping a big ol' DEC VAX into your data center back in the day. Of course, a unified computing chassis isn't as big as a VAX, but it's going to become pretty important as you pile on lots of critical network and application services onto one chassis. You wouldn't want to shotgun implement a VAX, and you don't want to shotgun implement your agile computing strategy either.

Data Center In A Box

Consider what it takes to deploy an application. A project manager tasked with building the plumbing must sit down with server administrators, network engineers, and storage engineers for starters. That's bad news for the application, because senior-level admins tend to be smart people with significant cross-functional expertise, and that means they're going to have a lot to say about your application, its impact on the network, and the resources that are going to be required to make it work. Before you know it, you're in negotiations that

Sponsors of Tomorrow.™ intel

IT sees:

Servers that intelligently scale performance.

The CFO sees:

Servers that save energy.

The latest Intel® Xeon® processor analyzes its workload and automatically adjusts to deliver maximum performance when you need it—and big energy savings when you don't. That's smart no matter how you look at it.

Learn more at **intel.com/go/xeon.**

IN THIS ISSUE

Tip 2 Delineate Operational Responsibility

If learning curve theory holds, your server, network, and storage teams are going to need time to adjust to sharing a common administrative interface. And since learning curve theory never fails, you should start to see and feel the operational agility relatively quickly as you implement. All IT disciplines will need to change the way they plan, provision, and deploy services around the new capabilities of the unified computing platform.

make passing healthcare reform look simple.

One customer of a unified computing vendor relates that it was taking almost 18 months to simply provision the servers needed to execute on projects in their environment. While that's an outrageous example of operational lethargy, as a framework for the provisioning challenges IT faces today, it drives home the point that there's a serious agility issue within the enterprise.

When you're through planning, you can start running cables—lots of them. You need storage connectivity redundancy, power redundancy, and out-of-band management. Maybe you need a little more than Gigabit Ethernet bandwidth, and you probably need to support network connectivity to multiple upstream switches for IP failover. If

Dell PAN System Pricing

BLADE	QUANTITY	LIST PRICE	EXTENDED LIST (Quantity x List)
PowerEdge M610 blade (includes a E5540 CPU)	8	$5,397	$43,176
2nd E5540 CPU	16	Included	
48-GB RAM per blade	96	Included	
2 X 73-GB HDD	16	Included	
Interfaces for Virtual I/O	8	Included	
Total cost of 8 blades			**$43,176**
CHASSIS			
Dell/Egenera Datacenter-In-A-Box	1	$99,000	$99,000
Redundant chassis power supplies		Included	
All management software, including HA		Included	
Virtual I/O		Included	
Total cost of chassis			**$99,000**
Total price			**$142,176**

Note: Advanced orchestration, provisioning, and system and I/O management software is included and comparable to that offered by HP and IBM. System requires an external SAN that's not reflected in the price.

Cisco Unified Computing System Pricing

BLADE	QUANTITY	LIST PRICE	EXTENDED LIST (Quantity x List)
B-200 M1 blade server	8	$1,575	$12,600
E5540 CPU	16	$1,099	$17,584
4-GB DIMM	96	$230	$22,080
73-GB HDD	16	$349	$5,584
Menlo Q	8	$749	$5,992
Total cost of 8 blades			**$63,840**
CHASSIS			
Chassis 6U	1	$3,199	$3,199
Power supply unit	4	$499	$1,996
Chassis management capability		Included	
10-GbE fabric extenders	2	$1,999	$3,998
Fibre Channel I/O		Included	
Total cost of chassis			**$9,193**
REQUIRED SWITCHING INFRASTRUCTURE			
Nexus 6100 Fabric Interconnect (8 ports, 2 PSU)	1	$15,497	$15,497
8-port FC GEM	1	$1,699	$1,699
Additional port license key	4	$1,479	$5,916
			$23,112
Total price			**$96,145**

Note: Cisco UCS Manager is included with each chassis purchase. Advanced orchestration and infrastructure management software needs to be purchased from third party, such as CA or BMC. System requires an external SAN that's not reflected in the price.

you're deploying 1U pizza box servers, that's a lot of cables sticking out of the back of your server. Add in KVM and multiply all of that by a rack full of 1U servers, and you've got a wiring mess requiring a Ph.D. in structured cabling to keep it all straight.

For the last several years, forward-thinking vendors like Cisco, Hewlett-Packard, IBM, and Egenera have been spending lots of R&D dollars trying to figure out a way to not only consolidate servers, but to also consolidate the networking and storage aspects of

IN THIS ISSUE

delivering services to the business. What has emerged will change the way your data center looks in five to 10 years.

Imagine a blade chassis where you only need two patch cables to your distribution or core layer switch. Imagine a world where you virtually slice up the bandwidth available on those cables and make it appear to the operating system as a Fibre Channel interface, an iSCSI interface, a GigE link, a bonded GigE link, or even a 10 GigE link. Imagine a world where your storage administrator pre-provisions LUNs all at once so that you can dynamically access the storage later by referencing the LUNs' World Wide Port Name (WWPN). And image not having to trek all the way to your collocation facility to build the Layer 1 infrastructure needed to add or change network and storage services. Today's "wire once" unified computing platforms, a true data center in a box, isn't a pipe dream. It's here today and it's evolving at breakneck speed.

If all that sounds good to you, it won't mean a thing to your CFO, so let's talk about the things she cares about. Let's start with power consumption and data center floor space. Tudor Perini, a billion-dollar construction company that recently became Cisco's first big UCS customer, collapsed its server footprint from 400 physical servers into 36 Cisco server blades running on a four-UCS chassis, says James McGibney, Tu-

Tip3 Evaluate Each Server's I/O Needs

Getting unified computing gear racked and stacked takes only a couple of hours, but you've got a finite amount of concurrent I/O you can push through each unified chassis. With the current generation of Cisco UCS, 80 Gbit/sec is the max. That's not to say you can't configure more than that virtually, but past 80 Gbit, your quality-of-service policies will dictate any bandwidth throttling that has to happen. As a result, map out your needs based on real-world performance counters. For example, does file server X actually need bonded Ethernet NICs? If you're averaging only 10% utilization on the bonded NICs, you can scratch a virtual interface from your migration plan. Do a similar exercise on the storage side—what's the average utilization on your iSCSI and Fibre Channel links? Can you size down, or where necessary, can you turbo boost I/O to a high-priority transactional application by allocating a 4-Gbit virtual FC interface?

dor Perini's data center operations manager. At the same time, McGibney and his team managed to consolidate five physical data centers across the U.S. into one, in 58 days—a remarkable accomplishment, regardless of how much "help" Cisco provided.

While McGibney wasn't able to provide an actual dollar amount saved—considering facility, power, administrative, and management costs—he estimates the operational IT budget savings at around 60%. It's tough to measure ROI when you don't have exact capital side numbers, but based on our best guess, you're looking at a return of your capital expenditure in less than a year, assuming typical four-year replacement cycles. That's a payback period that should get that CFO on board.

Tip 4 Test Virtual I/O Connectivity

This is where you should spend most of your time. If you're using Cisco UCS, then you're also committing to using the Nexus Fiber Interconnect switch, because the Layer 3 control protocol between the Cisco UCS blade chassis and the Nexus Fiber Interconnect switch is proprietary. With IBM, for example, Open Fabric Manager is able to communicate with a multitude of switching vendors and technologies. So the key takeaway as you move from a physical to a virtual I/O world is to be aware of the limitations and compatibility that your unified computing solution provides you based on your existing network and storage infrastructure. While in pilot, now's the time to learn the ins and outs of creating and destroying various types of virtual I/O interfaces, as well as automating their provisioning and learning how to map them during failover scenarios.

Strategy For Success

If space/power utilization and business agility are a concern for you, then the question isn't if a unified computing solution is in your future, but when. The question then becomes which vendor, management solution, and migration strategy to employ to get there safely.

Given the big names in unified computing, it's safe to say that if you make an investment in Cisco, HP, IBM, or Dell, they'll be around to service your solution 10 years from now. So the answer to this multivariable equation now depends on what your total infrastructure management needs are. Perhaps a flow chart would help here.

Decision: Do you require a vendor that can deliver a hardware solution and advanced orchestration software capable of automating complex business rules and scripts all under the same roof? If yes, then HP

IN THIS ISSUE

and IBM may be the best choice for you.

However, many organizations are happy with the alerting and fault tolerance they get from VMware's ESX combined with their network management system of choice. Not every organization needs a team of HP or IBM engineers to orchestrate the reconstruction of the matrix based on a set of events that happen in real time. In that case, the hardware and provisioning innovation in the Cisco UCS chassis could be everything you've been waiting for.

Tip 5: Know The Impact Of Aggregating I/O And Apps

Let's assume you're the data center manager in charge of a Tudor Perini-like data center consolidation. Since you're going from five data centers to one, it's safe to say that traffic patterns will change within your network. So how will the consolidation of internal applications within a single site impact the user experience? How will it impact network topology? How will it impact the design of your network at the access, distribution, and core? These things can be simulated with the proper tools. It's imperative to take the time to analyze the impact of centralizing resources into a single site.

Major Players

So perhaps by now you're envisioning a utility computing platform as some sort of monstrosity that resembles a petaflop supercomputer like IBM's Roadrunner. Don't worry, today's unified computing platforms won't take up an entire computer room. In fact, a single rack will do just fine, because part of the value proposition of unified computing is data center consolidation, not expansion.

If you stood a standard blade server chassis right next to a unified computing chassis, you wouldn't be able to tell the difference. However, if you walk around and look at the back, you'll discover right away that the new name of the game is virtual I/O and massively scalable RAM support.

Cisco's the newest utility computing player of the bunch, and of late they're also the leading innovator. Several innovations in the Cisco UCS Blade Server Chassis make it quite interesting from an engineering perspective. First, Cisco is the first vendor to offer "stateless server computing." Loosely translated, that means a Cisco UCS server blade has

no embedded BIOS or firmware like a conventional server blade does. So without a UCS chassis, a Cisco server blade can most aptly be described as "a thing."

What's interesting about Cisco server blades is that the BIOS and firmware are abstracted from the physical hardware and are managed as attributes of a "service profile." In fact, every hardware component of a traditional server is abstracted and can be described to the UCS manager via an XML file. Here's why that's cool: Suppose you have a standard Web server build that you've certified to run on a certain hardware platform. Those hardware characteristics can be packaged into a template and deployed over and over again almost instantaneously. Let's further suppose that your standard Web server build needs two GigE interfaces and a Fibre Channel HBA for connectivity to a SAN. With Cisco UCS, all of those hardware resources can be packaged as a service profile and are deployable as a template at will. Assuming your storage engineer is able to pre-provision a pool of LUNs, you can populate

HP BladeSystem Matrix Pricing

BLADE	QUANTITY	LIST PRICE	EXTENDED LIST (Quantity x List)
HP BL460 G6 blade	8	$1,575	$12,600
E5540 CPU	16	$1,099	$17,584
4-GB DIMM	96	$230	$22,080
73-GB HDD	16	$349	$5,584
HP Qlogic 2462 4-Gb FC	8	$749	$5,992
Total cost of 8 blades			**$63,840**
CHASSIS			
HP C7000 chassis	1	$6,619	$6,619
Power supply units	6	$349	$2,094
HP C7000 VCEM	1	$7,000	$7,000
10-GbE VC Ethernet Flex-10	2	$12,199	$24,398
HP BLc VC 4-Gb Fibre Channel module	2	$9,499	$18,998
Total cost of chassis			**$59,109**
Total price			**$122,949**

Note: Advanced orchestration, provisioning, and system and I/O management software is included and comparable to that offered by Dell and IBM. System requires an external SAN that's not reflected in the price.

IN THIS ISSUE

your service profiles with a pool of WWPNs corresponding to those LUNs, and they can be automatically attached to newly deployed virtual servers, or can be deployed to an individual UCS server blade if necessary.

Cisco is just now delivering its second-generation server blade, code named Ventura. Jointly developed with Intel, Ventura makes use of a custom ASIC that allows it to address up to 384 GB of RAM on a dual CPU socket blade. That's a big deal, because one of the problems with traditional two-socket Intel Xeon 5500 series blades is that the amount of addressable RAM is limited to 144 GB. With Cisco's Extended Memory Technology, the ASIC essentially maps four DIMM slots into one at full speed. As a result, each Ventura blade has 48 DIMM slots, compared with the typical 12 DIMM slots found on standard two-socket blades. With 48 DIMM slots on each Ventura blade, IT has the option of buying cheaper RAM for its virtual server needs, or it can pile 8-GB sticks into each slot to maximize the number of VMs per blade. Packing that much RAM on a blade is a feature unique to Cisco.

IBM BladeCenter HS22 Pricing

BLADE	QUANTITY	LIST PRICE	EXTENDED LIST (Quantity x List)
BladeCenter HS22 (includes an E5540 CPU)	8	$2,895	$23,160
2nd E5540 CPU	8	$1,175	$9,400
4-GB DIMM	96	$235	$22,560
73-GB HDD	16	$359	$5,744
Virtual Fabric Adapter	8	$699	$5,592
Total cost of 8 blades			**$66,456**
CHASSIS			
IBM BladeCenter H	1	$4,725	$4,725
Power supply unit (optional)	1	$1,099	$1,099
Chassis Management capability		Included	
BNT virtual fabric 10-Gb switch (1 required, 2 optional)	2	$11,199	$22,398
Total cost of chassis			**$28,222**
Total price			**$94,678**

Note: Note: Advanced orchestration, provisioning, and system and I/O management software is included and comparable to that offered by Dell and HP. System requires an external SAN that's not reflected in the price.

Virtual I/O is also key to Cisco's UCS offering, and it's baked right into the chassis and Manager software, so no blade slots are taken to do virtual I/O. Unfortunately, in order to do virtual I/O on a Cisco UCS chassis, you must terminate all of that virtual I/O using a Cisco Nexus Fabric Interconnect Switch, because the Layer 3 control protocol that encapsulates the virtual I/O is proprietary. IBM and HP support multivendor virtual I/O scenarios.

But Cisco's the new kid on the block, and HP has been the x86 server market leader for years. Cisco has made no secret that it thinks HP is playing a shell game with respect to the way it decouples vital management software from the server purchase. But the profit margins and competition in the commodity server market are tight, so it stands to reason that HP puts an emphasis on value-added management software. Conversely, as a new entrant to the server market, Cisco now has the luxury of charging more for a complete package with unique, albeit proprietary, capabilities, and of attacking HP for offering management software a la carte.

What makes Cisco UCS worth considering is its agility with hardware provisioning and its integrated ability to discover and manage additional UCS boxes at no additional cost. A further value add is the limited visibility you get into your Cisco switching and routing infrastructure through the UCS Manager. But UCS has a ways to go to catch HP, BMC, and CA offerings. UCS can't manage the actual workload of each individual virtual server running in your environment. Cisco

Tip 6 Test Disaster Recovery Plan

If you're Cisco customer Tutor Perini, one of the advantages of a distributed data center is distributed facility risk. But if you're collocated with a Tier 1 colo, uptime and availability are smaller concerns compared with how you'll handle disaster recovery within the unified computing solution itself. You're much more likely to see a blade failure than a UPS explosion at your colo facility. The disaster recovery options from a server, a blade, and an I/O perspective are numerous with unified computing solutions. Take your time and test various scenarios. One benefit of unified computing is that disaster recovery testing is relatively easy to do, even remotely, because you can simply disable various blades through the management interface and see what happens. Out-of-band management further helps minimize your remote management risk, so clearly the time to shake out all of your disaster recovery procedures is before you go live.

IN THIS ISSUE

doesn't offer change management tools, orchestration tools, or tools that assist with application awareness and health, at least not yet.

On Standby

Oh, and then there's that company called IBM, you know, the original big iron computing pioneer. In fact, IBM's UCS strategy looks a lot like Cisco's in that the IBM BladeCenter offers integrated I/O virtualization through its BladeCenter Open Fabric Manager product. And like UCS, BladeCenter's Advanced Management Module provides an impressive array of hardware provisioning capabilities, including the ability to pre-provision server addressing info, WWPN for SAN connectivity, and the like, and automatically deploy those addresses to specific server blades. What's also impressive is BladeCenter's ability to orchestrate a failover routine that will fire up a standby server blade in the same chassis or a different one in the event of blade failure. The standby blade automatically inherits all of the characteristics of the failed blade, including port and I/O mappings.

Last but certainly not least is the Dell/Egenra relationship, which is positioned to compete with the rest of the pack through tight integration of Egenra's Processing Area Network Orchestration (PAN) Management software on Dell blade hardware. While Dell has a leg up on Cisco in terms of the number of installed customers, it has a couple of interesting competitive challenges as the unified computing sector develops. First, Dell is OEMing Egenra's technology, with Dell providing Level 1 support and Egenra providing Level 2 and 3 support. With IBM, HP, and Cisco, you get support under the same roof, and there's no OEM relationship to worry about souring. So while Dell is obviously a massive player in the x86 server market, and Egenra is a proven player in orchestration and I/O virtualization, it will be interesting to see how Dell approaches the changing competitive landscape.

Write to Randy George at rgeorge@nwc.com.

IN THIS ISSUE

Energy-Efficient Data

By Mike Fratto

Remember those green aspirations you had? How you'd do your bit to save the planet by re-engineering your IT operations for power efficiency? Don't worry, no one else remembers, either. But the fact remains, most data centers are still power-sucking cost centers and represent an opportunity for IT to save the company some cash.

A 2007 Environmental Protection Agency report states that about half of electrical costs in data centers goes to powering environmentals like lights, fans, and compressors. That jibes with what Paul Jacobson, principal of data center consultancy Reliable Resources, is seeing. He says organizations typically spend between 45% and 55% of annual electrical expenses on cooling their data centers, making them prime targets for optimization and cost reduction.

Dean Nelson, senior director of global data center services at eBay, says organizations built data centers "2N+1" to meet availability goals. "There wasn't a focus on the efficiencies of the data center because one second of downtime was a lot of money," Nelson says. "If you put in a good, efficient design, you will save money. It's a no brainer."

Building an efficient data center starts with collecting data and reporting on power and cooling usage on a minute, hourly, or daily basis. Collection of environmental data is still in its infancy, with few standards. In January, Cisco launched its Energywise partnership program with the goal of positioning Cisco Catalyst switches as data collection points, which would forward the data to a repository and control points, and in turn could turn off the switches via Power Over Ethernet ports. Cisco initially aimed Energywise at the LAN and not the data center, though a

Next >>

Unless your data center was designed in the last five years, it's not as efficient as it could be. Here are the basics on efficient designs.

Center Designs

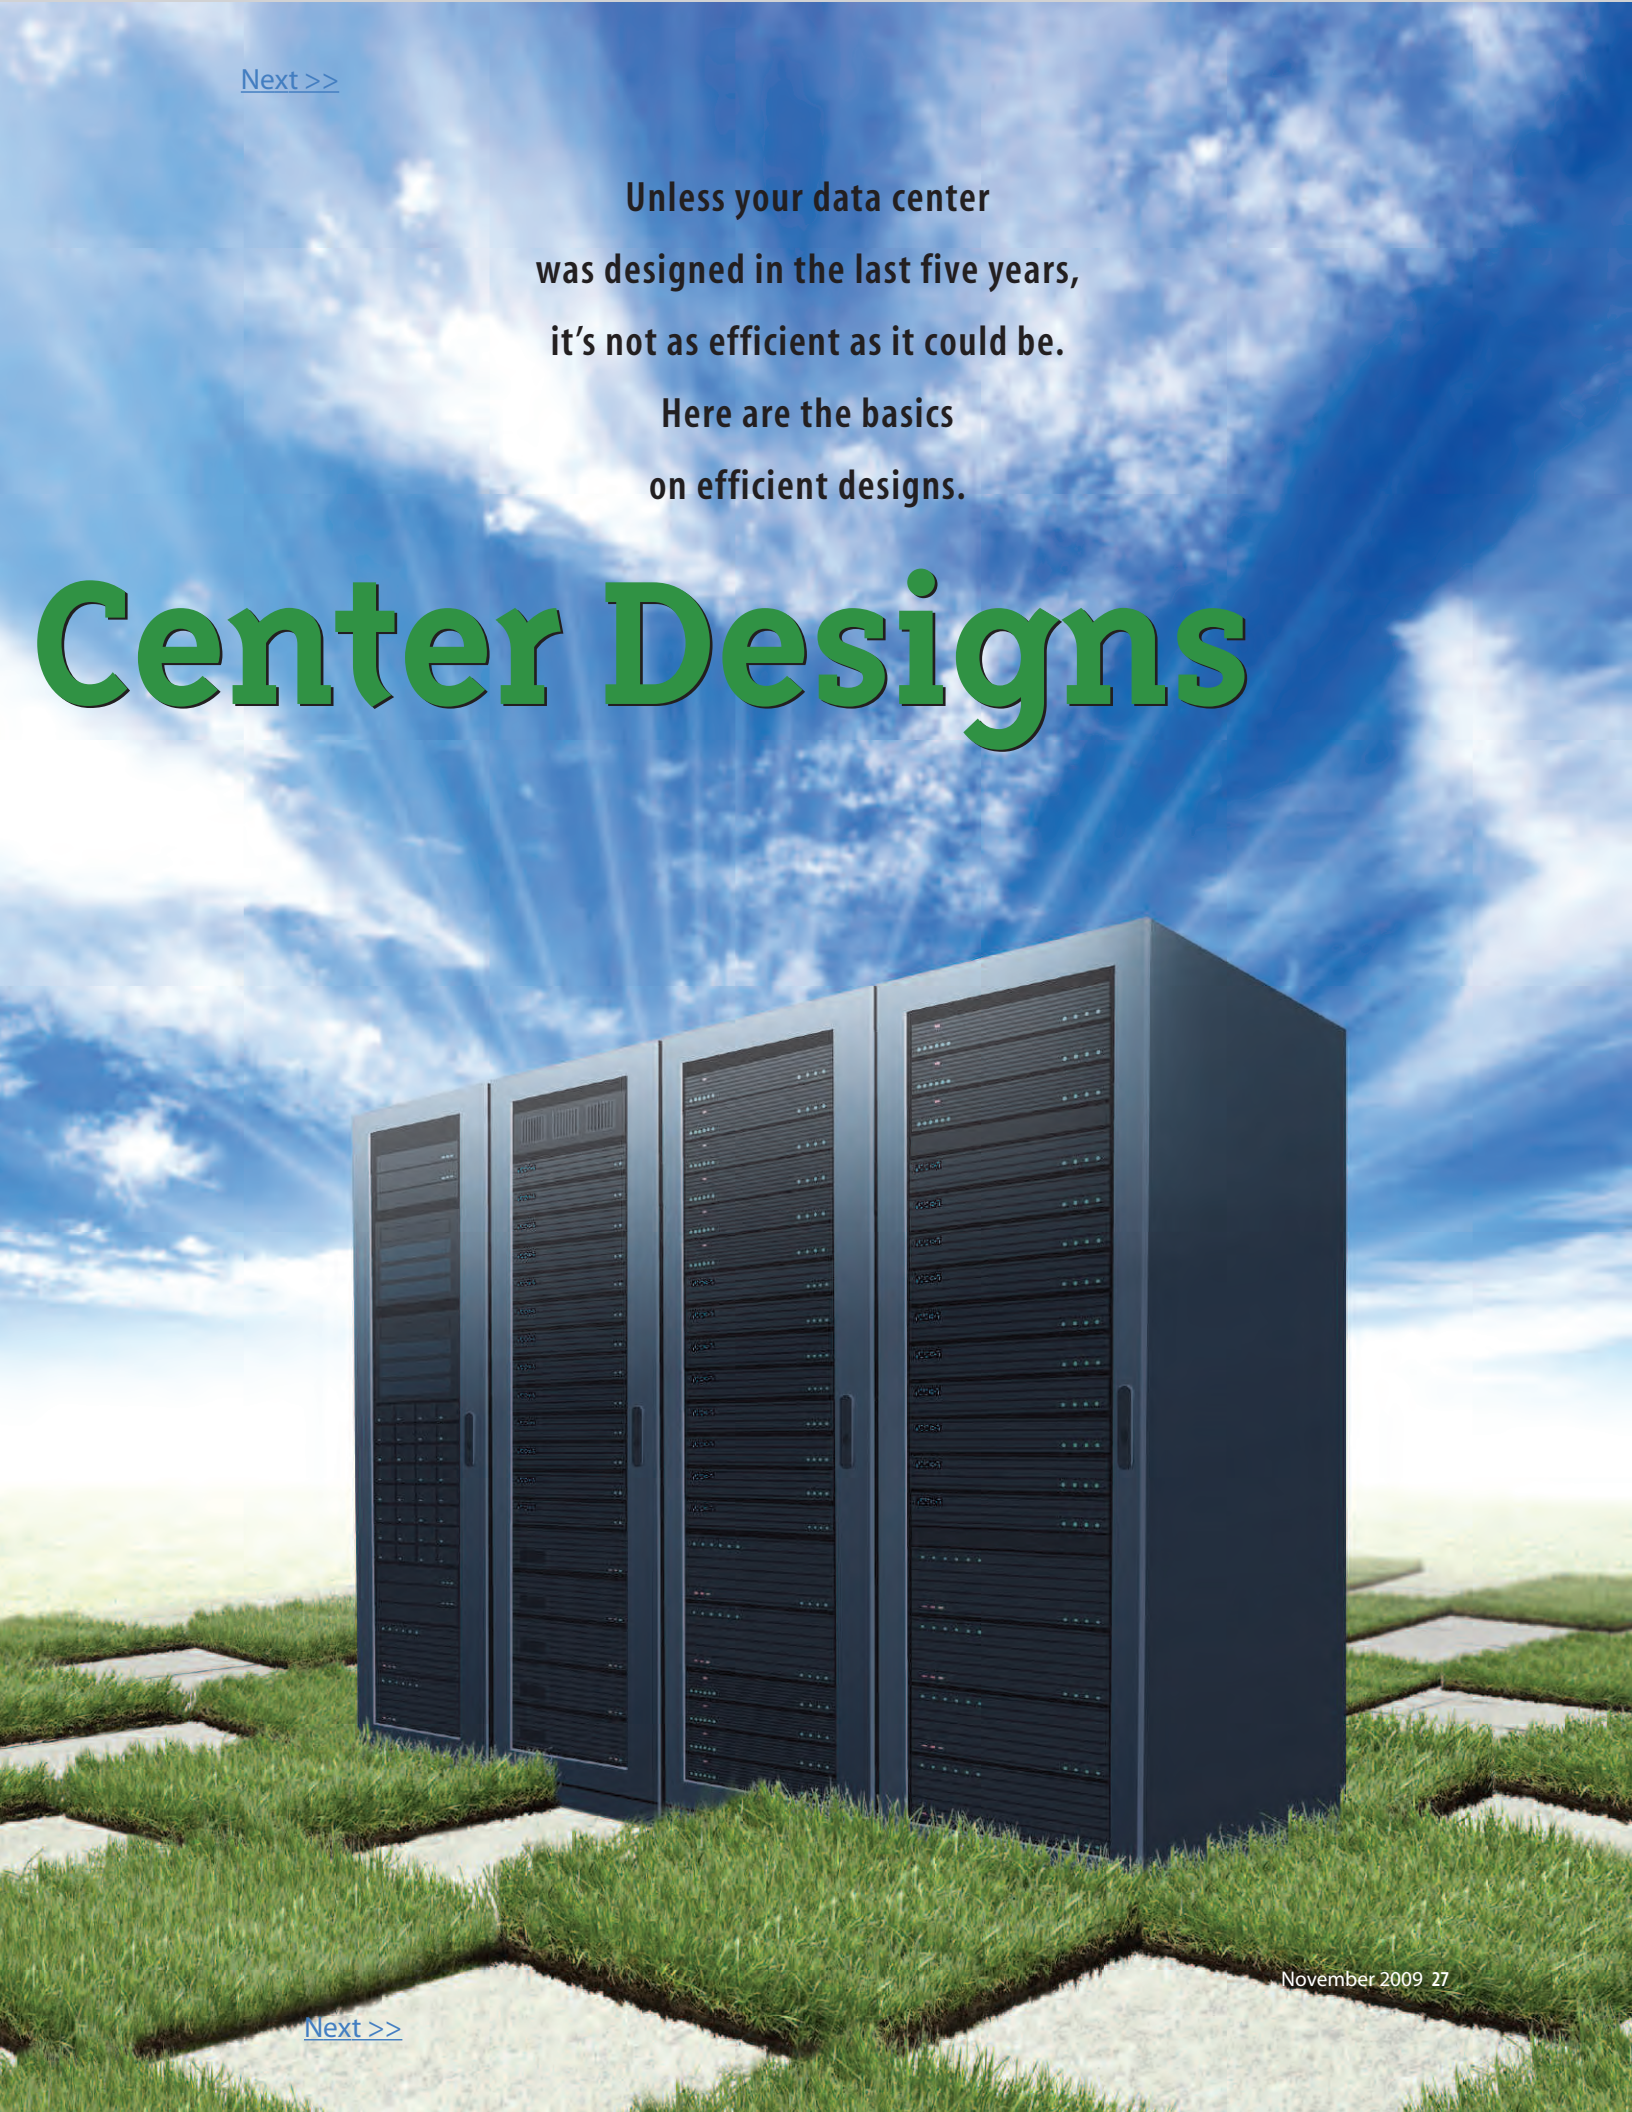

Next >>

IN THIS ISSUE

goal is to plug into building management systems to control lighting and air conditioning.

Syracuse University is building a new data center and its IT team will monitor power usage down to the individual plugs so that it can report on server and application consumption, says CIO Christopher Sedore. But he wants to go further. Syracuse makes extensive use of virtualization, and Sedore would like to be able to measure how much power an application is using whether it's hosted on a physical or virtual machine.

Sedore says there are few relevant measures of power usage. Syracuse has looked at power usage effectiveness (PUE) and data center infrastructure efficiency (DCiE), but those measures either weren't detailed enough or "didn't tell us anything useful about our power usage," he says.

PUE is a ratio that divides total data center power by the power used for IT equipment, including servers, switches, SANs, etc. So a PUE of 3.0 means that for every watt of power used for IT equipment, there's another 2 watts used for non-IT facilities like cooling and lighting. DCiE, the reciprocal of PUE (IT equipment power divided by total facility power), shows the percentage of power used for IT. A data center with a PUE of 3.0 has a DCiE of 33%. A data center with a PUE under 2.0 is considered to be very power efficient.

Like with any number, however, it's not enough to say a data center is green or not. Sedore points out some shortcomings of PUE and DCiE, such as not accounting for reuse of waste heat, or not accounting for the fact that external power supplies and fans aren't factored into facility power, adversely affecting the ratings.

Who's Got The Power?

One of the conventional wisdoms is that DC power is more efficient than AC power because DC distribution goes through fewer conversions. A paper from Lawrence Berkeley National Laboratory shows that

DC distribution systems are about 7% more efficient than "best in class" AC UPSs and 28% more efficient than AC distribution systems typically found in data centers. A task group of the Green Grid, an IT industry consortium that's looking to standardize on energy efficiency metrics, processes, and technologies, takes issue with parts of Lawrence Berkeley's study. The task force argues that well designed DC and AC systems using current technology can actually be within 5% to 7% efficiency of each other.

Before you replace your existing power distribution plant, you can do other, less radical things to reduce power costs. A typical server consumes 60% to 90% of its peak system power even at low utilization, according to the EPA. Consolidating low-utilization servers into high-utilized servers will cut power usage.

"Consumption and power costs are going up and it isn't sustainable," notes eBay's Nelson. "Organizations are looking at efficiencies. Instead of using 2 megawatts to run a data center, I can drop that to 1.5 megawatts to do the exact same work. Best practices are bringing down cooling costs and changing IT loads through virtualization, more efficient hardware, and better power management."

Based on data from the Department of Energy, the annual electricity cost to power eight individual servers is between $2,113 and $4,524 in Arizona and between $2,823 and $6,045 in Alaska. The annual electricity cost for a chassis with eight similarly sized blades is $1,118 in Arizona and $3, 985 in Alaska. Blade servers at 100% utilization use about 35% less power than comparable standalone servers.

That's significant savings. The list price difference between a chassis ($99,515) and similarly configured eight-rack-mount server ($102,016) is roughly equivalent. So you'll start to see savings in year two. Bear in

IN THIS ISSUE

mind the savings are cumulative as you add more systems, and there are other benefits, like easier management.

Turn It Off

Dynamic power management, where servers are turned on and off based on demand, is a promising field. Many of the technologies used for cloud computing, like virtualization and orchestration, are designed to dynamically add application servers to server pools as demand dictates. The goal is better application performance, but a side effect is reduced power usage.

Sedore is examining dynamic power management for Syracuse University, but he isn't sold on the idea. "We serve a large and demanding population," he says. "We need to ensure our service-level agreements are met, first and foremost. We may be able to save, for example, $5,000 a year with power management, but one downtime event will cost us much more. We will be looking at power management, but we won't deploy it until it's rock solid."

Electrical Costs Compared

RACK-MOUNT SERVERS	Watts/hour	Annual power costs (low) $0.0824 per kw/h	Annual power costs (high) $0.1764 per kw/h
Idle	366	$264.19	$565.57
100% utilized	489	$352.97	$755.63
8 RACK-MOUNT SERVERS			
Idle	2,928	$2,113.50	$4,524.53
100% utilized	3,912	$2,823.78	$6,045.07
CHASSIS W/8 BLADES			
Idle	1,549	$1,118.11	$2,393.61
100% utilized	2,579	$1,861.58	$3,985.24

Note: Rack server is dual-core 2.66-GHz server with 16 GB of RAM, a PCI-E card, two hard drives, and dual redundant power supplies. Blades are dual-core 2.66-GHz server with 16 GB of RAM, a PCI-E card, two hard drives, and dual redundant power supplies.

Data: Typical pricing from publicly available sources

Chill Out

All that server, storage, and networking gear generates heat. Keeping the equipment cool makes it run more reliable and extends the life of the components. A number of factors influence data center cooling designs, including computing density, total cooling capacity, data center floor layout, and where indoor and outdoor components can be located, notes Jacobson, the Reliable Resources principal. "The decisions for a dense urban or rural location can be very different," he says.

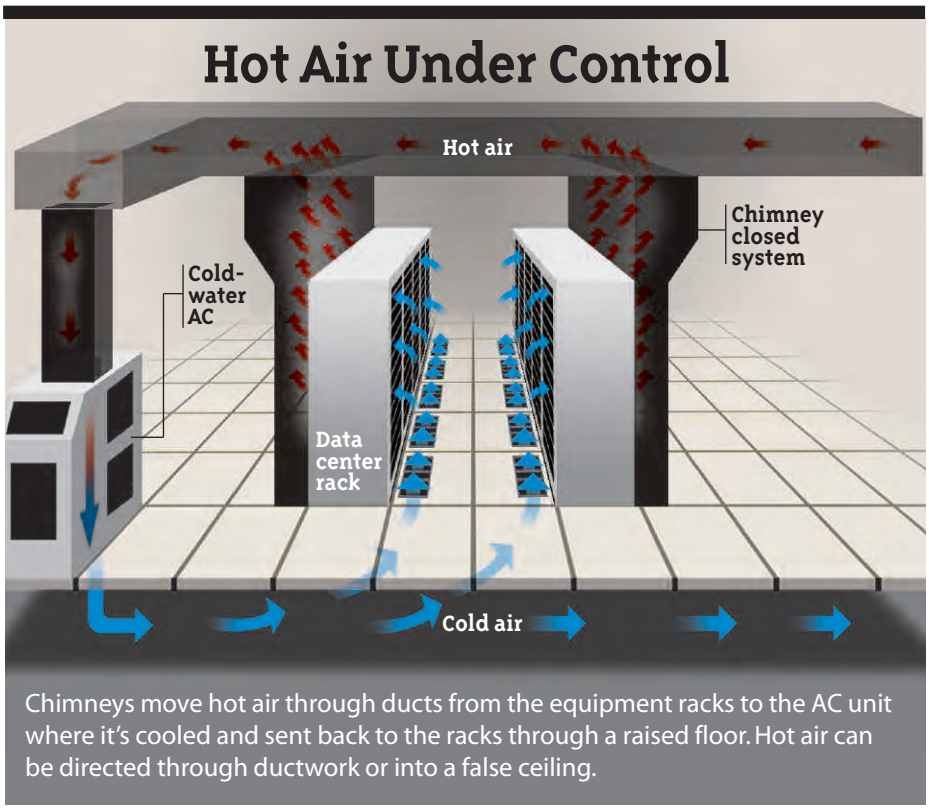

Chimneys move hot air through ducts from the equipment racks to the AC unit where it's cooled and sent back to the racks through a raised floor. Hot air can be directed through ductwork or into a false ceiling.

Many legacy data centers use a strategy that Jacobson calls chaos cooling, where CRAC units are positioned at the perimeter of the room, blowing cold air under the floor in the hopes that cold air will get to the front of the racks and hot air will be circulated back to the exchanger. That approach may work in low-density data centers, but in most cases, cooling is over-supplied because the cool air mixes with warmer air on the way to the racks. By the time it arrives, the air temperature could be 20 degrees higher than when it left the supply. That means if you want to deliver 72 degrees to the rack, the air needs to be chilled to 52. This approach leads to the all-too-common meat locker effect. If you need to wear a coat in your data center, you can bet it's highly inefficient.

Because eBay replaces data center equipment every two years, high heat sources and low heat sources sit side by side, Nelson says. "I need to match the cooling with the heat load at any location," he says. Hot aisle/cold aisle arrangements are the first step. You can attach doors or curtains on the hot or cold aisles to contain air flow. Chatsworth Products, which makes cabinet and enclosure systems for data centers, maintains that customers have reduced their cooling costs by as much as 20% just by using inexpensive Plexiglas barriers between hot and cold aisles.

The next step is to bring containment down to the rack level (see graphic, above). Hot air in the back of a rack can range from 95 to 130 degrees. Di-

IN THIS ISSUE

recting the heat up to return-air plenums or even to a false ceiling using a chimney can improve cooling efficiency dramatically. Directing the heat back to the exchanger can make the chiller three times more efficient, says Chatsworth, technology marketing manager Ian Seaton. Direct expansion AC units—those that use refrigerants—don't offer the same benefit, and you can damage the AC unit with large temperature differences, so check with your AC vendor before directing hot air right back to the exchanger.

Since that heat is moved out of the room and back to the chiller, you can raise the outlet temperature of your AC unit, gaining 1% to 1.5% efficiency for every 1 degree increase in outlet temperature, because you no longer have to compensate for hot and cold air mixing from the AC unit to the rack.

An alternative to centralized air conditioning is in-row cooling or even in-rack cooling products, offered by vendors such as APC, IBM, and Liebert. This strategy is known as "close coupling" because the cooling equipment is close to the source of the heat and can respond to heat demands in a more targeted fashion. These cooling units, similar to full-room AC systems, use either chilled water or direct expansion. Consultant Jacobson says close coupled cooling can be more efficient because each unit is moving less air, the fans can be better quality (more efficient), and hot spots don't affect neighboring equipment.

> Syracuse University's Sedore is considering dynamic power management but won't deploy it "until it's rock solid."

Free, As In Air

Depending on where your data center is located, you may be able to use free air to cool your data center for part of its operating cycle. Air-side economizers take cold air from outside and deliver it to your air distribution system. The air-side economizer and your existing AC units work together to maintain the data center temperature. When the outside air is too hot, the gates are closed and the internal AC units kick in; when the tempera-

Geographic Pricing Variation

Average retail price of electricity for commercial customers in July

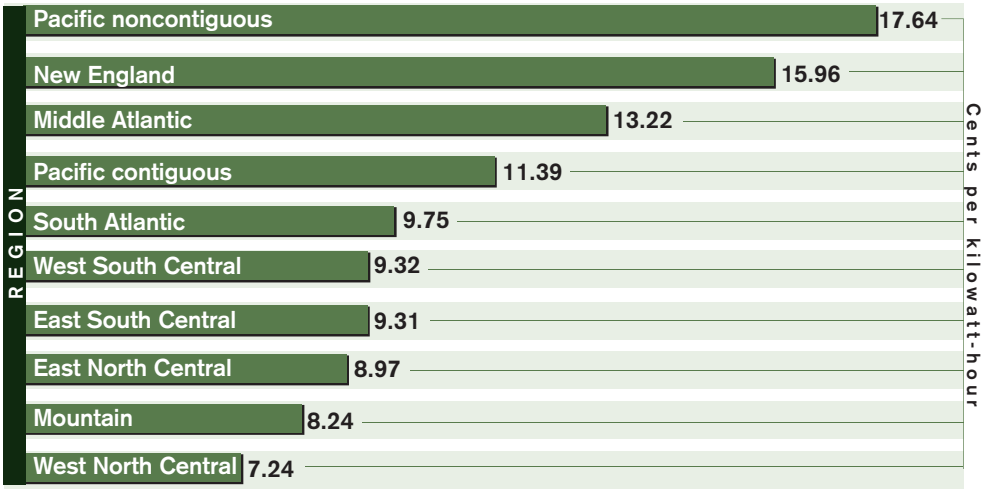

Data: Department of Energy, Energy Information Administration

ture drops, the AC units turn off and cool air is pumped in while hot air is pumped out.

There are, of course, a number of considerations to take into account before you adopt this approach. First and foremost, how much of the year, on average, does the outside temperature and humidity allow you to deliver outside air to your equipment? A data center runs 24x7x365, which works out to 8,760 hours. In many climates, nighttime temperatures and fall, winter, and spring daytime temperatures dip low enough to make free-air cooling possible. Second, you must factor in your cooling strategy. If you use in-row or in-rack cooling, air-side economizers probably aren't going to be a fit because you'd have to build an entirely new cooling distribution system in tandem with your in-row cooling and then coordinate each chiller with your economizers. Finally, you must determine whether you have the space for the air-side economizer equipment and can run the inlet and exhaust systems through your building.

Environments with a lot of dust and other particulate matter may also not be candidates. "You have to consider the outside air quality because outside air requires filtration, but filter frames leak," Jacobson cautions. For example, if air-side economizers bring in 100 times more fresh air than normal to cool a data center, then you're getting 100 times more unfiltered air through leakage around the filter, he says. That particulate matter can clog server electronics cooling fins, making them less efficient, so pay particular attention to filter maintenance and replacement.

Write to Mike Fratto at *mfratto@techweb.com*

IN THIS ISSUE

LiquidIQ

A True Data Center In A Box

Early adopters say this virtual server has cut the time required for setup and reconfiguration, and let them cut back on high-priced staff

While there's no doubt that virtualization has changed the way data centers are managed, bare metal configuration of systems can still take a lot of time. And in some highly dynamic environments, moves, adds, and changes account for a substantial portion of data center management costs. While Cisco has been beating the drum loudly for a different kind of server, it appears that Liquid Computing may have stolen the beat with its LiquidIQ system.

Early adopters report pretty astounding results, such as cutting operations staff by 50% while reducing environment build times by 90%, without substantial impact on capital costs. That's what a large Redmond-based hosted-partner data center is claiming, thanks to a new LiquidIQ rig. Moving from traditional rack-based hardware to a data-center-in-a-box model has radically cut setup and reconfiguration time for the company's sales partner environments, reducing the need for high-priced network and server operations staff. The main benefits from such a unified computing system won't be as dramatic for shops with stable environments, static infrastructure, or constant production loads. But as fewer of us have those 20th century luxuries, all-in-one offerings like Liquid's are an attractive alternative.

Liquid Computing was formed by ex-Nortel talent with a wealth of communications and mainframe experience, and it shows. Liquid IQ is the company's blade-based, "converged" system comprising up to 20 compute modules, storage, and switching in a forklift-delivered chassis,

By Joe Hernick

tied together with a set of management software. The system stretches the metaphor of virtualization to the hardware layer; administrators can rapidly deploy or reconfigure "logical servers" to blades as business requirements dictate, while network modifications to the switching fabric or logical servers can be made to running environments. The platform can support a wide variety of bare-metal configurations, from Red Hat and Oracle Enterprise Linux to 64-bit Windows 2K8 to VMware's vSphere. LiquidIQ falls under the emerging umbrella of centrally managed, complex systems with integrated computing, storage, and network environments. Hewlett-Packard and Cisco are among the other vendors in this market.

Liquid's gear is well positioned to provide flexible virtualized hardware for large-scale, complex, highly volatile (in the revenue-generating sense) environments—say, a vendor hosting site or a travel industry clearinghouse with dozens of partner-site front ends. Or how about any shop looking to run an internal or external cloud where customer demands necessitate the rapid setup and tear-down of physical servers?

LiquidIQ takes blade computing up a notch, allowing processors, memory, and centralized storage to be quickly allocated as required. Network settings can be modified on the fly; NICs and resources can be assigned to running logical servers, the internal switching fabric has up to 85 Gpbs available for data and SAN traffic. Processors and memory can't be reallocated to running systems as needs arise, but a server can be quickly locked, placed in stasis, taken offline, moved to a compute module blade with a more robust set of resources within the same chassis, and restarted in the time it takes to reboot a server. In a large installation, logical servers

the upshot

CLAIM Liquid Computing's LiquidIQ virtualizes compute, networking, and storage resources in a fully configured rack system. It's a data center in a box that only needs power and a network connection to get started, and it makes provisioning and reprovisioning a snap with a well designed management GUI.

CONTEXT As Cisco, HP, and IBM push toward more flexible and configurable servers and storage, Liquid delivers a self-contained system with simplified management and configuration. It's aimed at greenfield environments where new management apps are acceptable.

CREDIBILITY Liquid delivers on its promise of an automated data center in a box, allocating resources as needed by connecting them together like building blocks. Resources can also be allocated on running systems.

can be moved from one LiquidIQ chassis to another. Organizations can run a mix of bare metal instances and hypervisor hosts.

LiquidIQ abstracts server resources from a flexible hardware pool. Administrators can provision a logical server from available compute modules—physical blades with up to four quad-processor CPUs and 64 GB of memory, along with storage and network resources housed in one or more LiquidIQ chassis. Once everything has been selected using the GUI management console or command line interface, the server can be built as a traditional bare metal server with control of the blade and its resources, or as a virtualization host running guests as needed.

The entry-level Liquid IQ chassis starts in the low six figures, delivering 10 dual-socket, dual-core blades with a base 32 GB of memory each, two redundant switch and management modules, a single AC or DC connection, and 6 terabytes of storage from NetApp, all in a full-height chassis. Liquid includes management software and a year of support, as well as setup, installation, training, and shipping and handling. A partner site interviewed for this report is running 600 Hyper-V instances on

Virtual Data Centers Square Off

	Liquid Computing IQ	HP BladeSystem Matrix
Blades	10 dual-socket, dual-core blades, 32 GB of RAM	10 dual-socket, dual-core HP BL680c G5 blades, 32 GB of RAM
Switches	2 switch modules	Cisco Catalyst 3120 X
Management	2 management modules	2 HP Onboard Administrators
Chassis	Chassis	HP c7000 enclosure with DC module
Management software	LiquidView Management Suite: provisioning, cloning, monitoring, failover	Management software licenses included: HP iLO, HP Insight Recovery, HP Insight Dynamics VSE, HP Insight Orchestration
Storage	6-TB NetApp SAN	None included
List price	$170,000	$165,000

Data: Vendors

a 16-blade quad processor, quad-core LiquidIQ box. When LiquidIQ was compared with Sun gear, the price was comparable across both vendors until networking gear and management software were added in. The cost of a turnkey LiquidIQ system was the same as base Sunfire gear.

We had access to a LiquidIQ chassis for this review. After a brief tour, a spin around the planning guide and basic instructions had us up and running. Liquid pre-configures all systems for clients with varying degrees of customization, from base OS builds to tailored gold masters incorporating site-specific specs. Liquid's engineers work with clients to size and configure the chassis before it's shipped, then spend time on site to ensure integration. The standard package includes pre-build and post-build customization to incorporate client network parameters, porting physical servers to logical servers, and migrating VMs onto hypervisors running on logical servers within the LiquidIQ system.

We began with the "manage logical servers" option in the GUI and within minutes had provisioned a new bare metal server, cloned from a W2K8 R2 master. We connected to the new server instance and started tinkering. Jumping to the command line interface, we checked out the config and began editing network settings to configure switch port and VLAN details, install a NIC on a logical server, and assign its IP address. Eight lines of commands, and our server had a new configured network card mapped to a new VLAN in the switching fabric. The config checkout process was disarmingly simple. We authenticated to the management interface, signed out the config file for the entire LiquidIQ chassis, made our modifications, and signed it back in. That's it; the simplicity of the system is commendable, the potential power residing in the hands of administrators is frightening.

The straightforward and comprehensive approach to environmental control bodes well for disaster recovery. The configuration for every switch, VLAN, and NIC in the system is maintained in a single file. All network settings can be fully restored with the redeployment of a single file, and the ability to roll back to previous configurations can quickly resolve an errant edit.

Logical server wizards let admins dynamically create server instances as needed. LiquidIQ's ease of use is similar to that of vSphere or XenEnterprise, and we had to keep reminding ourselves that we were creating and tearing down physical system resources—entire physical servers—and not VMs. While creation of bare metal servers and hypervisor hosts was clear cut, customers still must rely on enterprise management tools for the care and feeding of live servers and VMs. Liquid's management tools are ideal for working with master images, cloning instances, and provisioning servers—at which point administrators will turn to traditional management toolsets. We'd like to see more integration.

Customers who opt for the bundled NetApp SAN benefit from Liquid's partnership with NetApp. While any iSCSI storage solution can be integrated with a LiquidIQ chassis, the management GUI incorporates storage options, provisioning iSCSI boot parameters, creating volume and LUN, and mapping required storage for new servers as part of the build workflow when paired with NetApp arrays. A site making its first move to SAN storage in parallel with a LiquidIQ implementation will likely opt for the integrated NetApp storage. Organizations with broader storage requirements or those running existing SAN gear from other vendors will look at the benefits of the NetApp system.

The flexibility of Liquid's architecture allows for the creation or reallocation of network resources on the fly. A chassis can hold up to 20 compute modules, with each capable of two gigE ports and four 10-gigabit Ethernet ports. The chassis fabric bandwidth is scalable up to 84 Gbps, and a chassis can max out at 24 gigE and 18 10-Gb external interfaces.

That fully loaded 20-module system comes stocked with 20 quad-socket, quad-core blades sporting the maximum of 64 GB of memory each, six redundant switch modules, two management modules, and 90 TB of NetApp storage for a cool $1.8 million.

***Write to Joe Hernick at** jhernick@nwc.com.*

By Andrew Hay

EventTracker 6.4

Prism Adds Much-Needed Monitoring

The latest version of Prism's security incident and event management offering gets enterprise activity and virtualization monitoring

Prism Microsystems' EventTracker is an entry-level to mid-range security incident and event management (SIEM) product that has lacked features such as integration with vulnerability assessment, configuration management, and identity and access management products. To combat its shortcomings, EventTracker has staked a claim to the niche virtualization event management market by integrating with VMware's hypervisor API. That's an area that industry leaders, such as ArcSight, Q1 Labs, and RSA, have yet to break into.

Prism Microsystems will release an impressive update before the year's end that will make up for some of its historical shortcomings. EventTracker 6.4 includes support for the Microsoft Hyper V and VMware hypervisors at the platform and virtualized guest level. It also adds support for physical hardware monitoring and reporting using the Dell OpenManage systems management application and enhanced enterprise activity monitoring by user, system, process, and IP address.

Server virtualization is one of the most complex operational projects most large organizations are undertaking. Unfortunately for those involved in the monitoring of these servers, traditional techniques don't work. Although it might be easy to discern if one or two virtual machines are created based on the logs of the host platform, it becomes increasingly difficult to track the provisioning of hundreds or thousands of virtual machines across a data center. The dynamic nature of virtualized systems makes it extremely difficult to keep

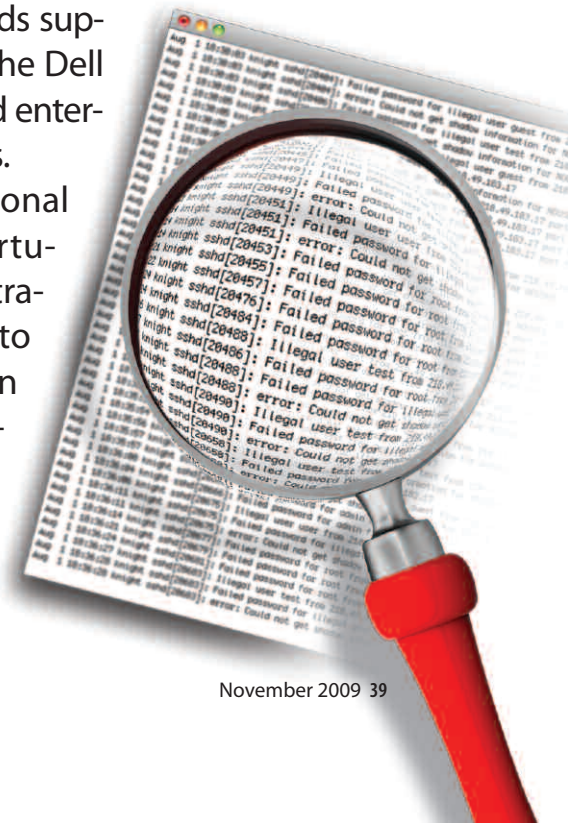

tabs on what's deployed, when it was deployed, and by whom.

EventTracker 6.4 can monitor all three levels of a virtualized deployment: host hardware, hypervisor, and virtualized guest. By interfacing with VMware, using the VMware API and Microsoft Hyper V application logs, this version of EventTracker monitors how virtualization platforms are used. Host and guest configuration changes are recorded, as are relevant status and error messages. They're tied to specific users, the machines they used, and their IP addresses. These logs can assist with security monitoring and auditing of virtualized systems and host platforms.

VMware and Microsoft Hyper V logs are displayed separately, showing product-specific details for each. VMware shows details such as resource usage alarm, user authentication, and permission rule change; Microsoft Hyper V contains product-specific categories such as Hyper V Hypervisor and Hyper V High Availability. Both groups list things such as host and guest configuration changes, status changes, and provisioning.

There's no pre-configured, all-encompassing virtualization group, so you can't report on similar events such as hypervisor logons out of the box. If you only have one virtualization technology, this may not be a deal breaker. However, if you're using multiple virtualization technologies, you'll have to look in several places to get the full picture. For example, a log that details the creation of a virtualized server should be represented within EventTracker as a single group for all hypervisors. A single event group would make for a much more user-friendly experience when building correlation rules for user created alerts, but you'll have to build them yourself. The last thing you want to do when building correlation rules is to have to add tens or hundreds of different event types that might match the data you're looking for.

Eye On Hardware

EventTracker 6.4 also adds support for monitoring physical hardware by interfacing with Dell OpenManage. A multitude of hardware-related events, such as disk errors, power supply messages, and temperature

levels, are tracked. This is a catch-up item that most of the major players, including Q1 Labs, netForensics, and ArcSight, have had for some time. Prism says that it's working on providing support for other hardware monitoring vendors, such as Hewlett-Packard, in future releases. Combining the logs from Dell OpenManage and your virtualization platform, be it VMware or Microsoft Hyper V, could act as an early warning system to a potential catastrophic loss of a virtualized server and its hosted virtualized systems. Alerts could be created that inform administrators of impending hardware or software issues.

EventTracker's enhanced enterprise activity monitoring capabilities include a configurable self-learning mechanism that lets it learn your environment's network traffic and system usage patterns. Based on this analysis, it can then identify new or unusual activity by users, systems, event ID, processes, and IP addresses.

This release also includes catch-up functionality such as per-item activity monitoring. Prism needed this functionality to remain competitive.

I haven't seen other vendors implement the ability to monitor activity by running process with much success. EventTracker 6.4 excels at this. It uses behavior analysis to alert you to the presence of any new process in your environment. This lets it detect the installation of inappropriate software, enhancing your incident response team's ability to identify compromised machines. This enhanced enterprise activity monitoring plus virtualization monitoring at the host and guest level rounds out EventTracker 6.4 and may give Prism an early foothold in the virtualization SIEM market.

the upshot

CLAIM The newest version of Prism's EventTracker can collect, analyze, and report on events in a complex and dynamic data center setting. The ability to monitor server hardware, hypervisors, and OSes provides a complete picture.

CONTEXT Prism is catching up with competitors on hardware monitoring, and EventTracker's process event analysis is second to none. While the virtualization monitoring is tightly integrated, we'd like to see more thorough hypervisor event grouping out of the box.

CREDIBILITY Prism's newest enhancements are good additions for people integrating virtualization into their data centers. We asked about hypervisor event grouping, and Prism says there's no customer demand—seems shortsighted.

Andrew Hay is a Canadian security professional who writes and speaks on privacy, forensics, incident handling, and network security management. Write to us at ahay@nwc.com.

IN THIS ISSUE

Rick Vanover

Right-Size Your Virtual Platform

Two-socket servers offer a great, low-cost consolidation platform. This is yet another case that proves bigger isn't necessarily better.

In virtualization circles, consolidation ratios have become the stuff of bragging rights. A year ago, if you bought a top-of-the-line server and loaded it with memory and connectivity, you might have been able to consolidate 30 or more physical servers onto one tricked-out machine. According to Moore's law, what was last year's tricked-out server should now be mainstream technology. Does that mean everyone should shoot for consolidation ratios of 30 or more? No, and here's why.

To consolidate at such high ratios, most of us would choose the newest four-socket systems. While quite powerful, these servers generally cost more than double their two-socket counterparts. Choosing between the two-socket or four-socket host server is a critical decision for consolidating at modest to very high ratios. The two-socket system can be quite powerful in today's configuration. AMD's and Intel's current systems are built with virtualization awareness, and the performance shows.

What Moore's law does for processor performance, it also does for memory capacity and cost. A few years ago the cost to get server memory north of 64 GB was substantial. Today, 64, 96, and 128 GB are quite affordable. As a result, more, smaller hosts can be provisioned with a lot of memory without breaking the bank. That argues for deploying more two-socket systems with lots of memory in them and not relying on techniques like overcommitted memory.

Whatever your consolidation ratio, the physical server is now a highly critical system and had better be highly available. Virtualization implementations primarily employ one of two high-availability technologies to protect workloads: host-based or virtual

machine-based with Microsoft Hyper-V and VMware's Fault Tolerance. For host-based schemes such as VMware's HA feature, there's a reserved inventory of host memory and compute resources known as admission control. This capacity is set aside to handle a host failure by reserving one host's equivalent resources in the existing cluster. VMware doesn't just set aside an idle host, but rather distributes enough headroom across hosts in the cluster to accommodate the entire workload should one of the other servers keel over.

The VMware HA feature is licensed, so there's an associated cost for protecting each host. If you have more modest consolidation ratios with relatively smaller host systems, the host inventory of memory reserved with admission control goes down and saves on licensing costs. To look at it a different way, would you rather HA consumed four CPU licenses or just two?

Then there's the I/O load. Every virtualization environment will eventually come across a few high-performance workloads that will push the storage I/O capabilities. By provisioning less VMs per server, we reduce the risk of network contention and allow more flexibility in where workloads reside.

Here's the bottom line: As virtualization technology matures and servers continue to pack more memory and processing power, what once required high-end systems can now be done with commodity servers. Using two-socket servers loaded with memory offers a great, low-cost consolidation platform and provides flexibility for addressing VMs with both high I/O and high-availability needs. This is yet another case that proves bigger isn't necessarily better.

Rick Vanover (VCP, MCITP, MCSA) is an IT infrastructure manager at Alliance Data. He has years of IT experience and currently focuses on virtualization, Windows-based server administration, and system hardware. Write to us at comments@nwc.com.

Network Computing
For IT, By IT

READER SERVICES

InformationWeek.com The destination for breaking IT news, and instant analysis

Electronic Newsletters Subscribe to InformationWeek Daily and other newsletters at informationweek.com/newsletters/subscribe.jhtml

Events Get the latest on our live events and Net events at informationweek.com/events

Analytics
analytics.informationweek.com
for original research and strategic advice

How to Contact Us
informationweek.com/contactus.jhtml

Editorial Calendar
informationweek.com/edcal

Back Issues E-mail: customerservice@informationweek.com
Phone: 888-664-3332 (U.S.)
847-763-9588 (Outside U.S.)

Reprints Wright's Reprints, Brian Kolb, 1-877-652-5295

Web: wrightsreprints.com/reprints/?magid=2196

E-mail:
ubmreprints@wrightsreprints.com

Media Kits and Advertising Contacts
createyournextcustomer.com/contact-us

Letters to the Editor E-mail iweekletters@techweb.com. Include name, title, company, city, and daytime phone number.

Subscriptions
Web: informationweek.com/magazine
E-mail: customerservice@informationweek.com
Phone: 888-664-3332 (U.S.)
847-763-9588 (Outside U.S.)

Please direct all inquires to reporters in the relevant beat area.

TECHWEB

Tony L. Uphoff CEO

John Dennehy CFO

David Michael CIO

John Siefert VP and Publisher, InformationWeek Business Technology Network

Bob Evans Sr.VP and Global CIO Director

Joseph Braue Sr.VP, Light Reading Communications Network

Scott Vaughan VP, Marketing Services

John Ecke VP and Group Publisher, Financial Technology Network, InformationWeek Government, and InformationWeek Healthcare

Martha Schwartz VP, Group Sales, InformationWeek Business Technology Network

Brandon Friesen VP and Group Publisher, InformationWeek, bMighty, Network Computing, Dr. Dobb's

Beth Rivera VP, Human Resources

Jill Thiry Publisher, Microsoft Technology Network

David Berlind Chief Content Officer, TechWeb, and Editor in Chief, TechWeb.com

Fritz Nelson Executive Editor, Information Week Business Technology Network, and Executive Producer, TechWeb TV

UNITED BUSINESS MEDIA LLC

Pat Nohilly Sr.VP Strategic Development and Business Administration

Marie Myers Sr.VP, Manufacturing

INFORMATIONWEEK VIDEO

informationweek.com/tv

Fritz Nelson
Executive Producer
fnelson@techweb.com

INFORMATIONWEEK BUSINESS TECHNOLOGY NETWORK

DarkReading.com
Security
Tim Wilson, Site Editor
wilson@darkreading.com

IntelligentEnterprise.com
App Architecture
Doug Henschen, Editor in Chief
dhenschen@techweb.com

NetworkComputing.com
Networking, Communications, and Storage
Mike Fratto, Site Editor
mfratto@techweb.com

PlugIntoTheCloud.com
Cloud Computing
John Foley, Site Editor
jpfoley@techweb.com

bMighty.com
Technology for Small and Midsize Business
Frederick Paul, Publisher and Editor in Chief
fpaul@techweb.com

Dr. Dobb's
The World of Software Development
Jonathan Erickson, Editor in Chief
jerickson@techweb.com

United Business Media

BUSINESS CONTACTS

Senior VP and Publisher, InformationWeek Business Technology Network, John Siefert
(949) 223-3642, jsiefert@techweb.com

Publisher's Assistant, Esther Rodriguez
(949) 223-3656, erodriguez@techweb.com

VP of Marketing, Scott Vaughan
(949) 223-3662, svaughan@techweb.com

VP, Group Publisher, Brandon Friesen
(415) 947-6213, bfriesen@techweb.com

VP of Group Sales, InformationWeek Business Technology Network, Martha Schwartz
(212) 600-3015, mschwartz@techweb.com

VP, Site Publisher, Chris Harding
(831) 426-2344, charding@techweb.com

Sales Assistant, Group Sales, Stacey Harden
(212) 600-3327, sharden@techweb.com

SALES CONTACTS—WEST

Western U.S. (Pacific and Mountain states) and Western Canada (British Columbia, Alberta)

Western Regional Director, Matt Stovall
(415) 947-6245, mstovall@techweb.com

District Sales Manager, Rachel Calderon
(516) 562-5338, rcalderon@techweb.com

Account Executive, Jenny Casey
(415) 947-6105, jcasey@techweb.com

Sales Assistant, Vesna Beso
(415) 947-6104, vbeso@techweb.com

Strategic Accounts

VP, Site Publisher, Chris Harding
(831) 426-2344, charding@techweb.com

Account Manager, Kevin Flint
(415) 947-6349, kflint@techweb.com

Account Executive, Krystle Liew
(415) 947-6121, kliew@techweb.com

Sales Assistant, Justin Kischinevzky
(415) 947-6214, jkischinevzky@techweb.com

SALES CONTACTS—EAST

Midwest, South, Northeast U.S. and Eastern Canada (Saskatchewan, Ontario, Quebec, New Brunswick)

Eastern Regional Director, Linda Hooper
(212) 600-3153, lhooper@techweb.com

District Manager, Jenny Hanna
(516) 562-5116, jhanna@techweb.com

Account Manager, Alex Dardac
(212) 600-3019, adardac@techweb.com

Account Executive, Michael Greenhut
(516) 562-5044, mgreenhut@techweb.com

Junior Account Executive, Darren Chu
(212) 600-3029, dchu@techweb.com

Junior Account Executive, Ray Capitelli
(212) 600-3045, rcapitelli@techweb.com

Strategic Accounts

Eastern Regional Director, Mary Hyland
(516) 352-2021, mhyland@techweb.com

Account Manager, Tara Bradeen
(212) 600-3387, tbradeen@techweb.com

Account Manager, Cori Gordon
(516) 562-5181, cgordon@techweb.com

Account Executive, Jennifer Gambino
(516) 562-5651, jgambino@techweb.com

Junior Account Executive, Julie Ann Kuczynski
(212) 600-3157, jkuczynski@techweb.com

SALES CONTACTS—NATIONAL

Global CIO, InformationWeek 500 Conference

National Account Manager, Amy Neidlinger
(212) 600-3163, aneidlinger@techweb.com

Sales Assistant, Stacey Harden
(212) 600-3327, sharden@techweb.com

SALES CONTACTS—EVENTS

Director, Event Operations, Jennifer Russo
(516) 562-5094, jrusso@techweb.com

MARKETING

Associate Publisher, Marketing, Winnie Ng-Schuchman
(516) 562-5982, wng@techweb.com

Director of Marketing, Sherbrooke Balser
(949) 223-3605, sbalser@techweb.com

Senior Marketing Manager, Ellen Asuncion
(949) 223-3622, easuncion@techweb.com

Field Marketing/Sales Ops Manager, Monique Luttrell
(415) 947-6304, mluttrell@techweb.com

AUDIENCE DEVELOPMENT

Director, Karen McAleer
(516) 562-7833, kmcaleer@techweb.com

ADVERTISING AND PRODUCTION

Publishing Services Manager, Lynn Choisez
(516) 562-5581 Fax: (516) 562-7307

Publishing Manager, Ruth Duggan
(516) 562-5111

MAILING LISTS

MeritDirect LLC (914) 368-1083

REPRINTS AND RIGHTS

For article reprints, e-prints, and permissions, please contact: Wright's Reprints, 1-877-652-5295, ubmreprints@wrightsreprints.com

BUSINESS OFFICE

General Manager, Marian Dujmovits

EDITORIAL OFFICE

(Fax) 516-562-5200
United Business Media LLC
600 Community Drive
Manhasset, N.Y. 11030 (516) 562-5000

TECHWEB

Tony L. Uphoff CEO

John Dennehy CFO

David Michael CIO

John Siefert Sr. VP and Publisher, InformationWeek Business Technology Network

Bob Evans Sr. VP and Global CIO Director

Joseph Braue Sr. VP, Light Reading Communications Network

Scott Vaughan VP, Marketing Services

John Ecke VP and Group Publisher, Financial Technology Network, InformationWeek Government, InformationWeek Healthcare

Martha Schwartz VP, Group Sales, InformationWeek Business Technology Network

Brandon Friesen VP and Group Publisher, InformationWeek, bMighty, Network Computing, Dr. Dobb's

Beth Rivera VP, Human Resources

Jill Thiry Publisher, Microsoft Technology Network

David Berlind Chief Content Officer, TechWeb, and Editor in Chief, TechWeb.com

Fritz Nelson Executive Editor, InformationWeek Business Technology Network, and Executive Producer, TechWeb TV

UNITED BUSINESS MEDIA LLC

Pat Nohilly Sr. VP, Strategic Development and Business Admin.

Marie Myers Sr. VP, Manufacturing

United Business Media