



OVERCOMING RESISTANCE TO THE CLOUD

It sounds simple. Corporate data centers will align themselves with the public cloud and realize heretofore unattainable economies of scale from this new and accessible form of Internet computing. But it's not that simple. There are too many vested interests that are ready to place obstacles in the path of a smooth and common-standards-based migration to the cloud and between clouds.

If hybrid cloud computing is to become the data center of the future, as we concluded that it would in Chapter 5, it will be necessary for computer operations managers to be able to move workloads freely between their corporate data centers and a public cloud.

This movement between internal and external centers will need to occur in ways that minimize friction between the two. What we have today instead is friction and resistance to the notion of cloud computing at so many levels of the process that it's still hard to conceive of doing this on any sustained basis. Several technical barriers exist, but we can start with dislike for the term *cloud computing* on the part of the CEO and other top-level executives, incompatible file formats demanded by the different virtualization vendors, and proprietary moves by cloud service suppliers. But skillful users, an increasing number of standards, and a growing supply of open source code are keeping pressure on the artificial constraints, and some of them will soon fall away.

What's in a Name? CEO Opposition

Let's start at the top. As Bob Evans of *InformationWeek* reported, when HP CEO Mark Hurd, as no-nonsense a personality as they come, spoke to a group of CEOs in late 2009, he described the future possibilities of computing using the term *cloud* and was nearly jeered off the stage. "Here I am talking about the cloud and all kinds of cool things that can happen with the cloud, and I got a lot of boos, um, after that. It started with the whole term, 'cloud.'"

After that experience, Hurd stated that "cloud computing" was an inadequate phrase for the things he wanted to talk about. In a rare moment of harmony for two competitors, IBM's CEO Sam Palmisano agreed, saying that *cloud* was "an

unfortunate name” and suggested “highly virtualized infrastructure” instead. That phrase misses the mark when it comes to the self-service, end user empowerment, and multitenant nature of most cloud computing.

Leaving aside Palmisano’s proposed renaming, when you are looking for opposition to the term *cloud*, you don’t need to look far. It’s possible that Larry Ellison’s outspoken jibes have resonated with those who find the term confusing, misapplied, or misleading. Soon Oracle will be directly involved in cloud computing, as its second-tier executives well know, and perhaps Ellison will then clarify his remarks for the benefit of CEOs everywhere. In the meantime, opposition will continue to come from those who can’t take the time to wrestle with the implications of what it means.

Admittedly, “cloud” can be an awkward term to explain. It has evolved as a descriptive term that captures a new computing distribution pattern and business model, at a time when that pattern is still getting established. Most likely, “the cloud” will rapidly evolve into more specific forms of computing that reflect what particular clouds will do. These clouds will take on more specific names, reflecting a concrete form of computer service.

An example of a cloud with a more specific definition might be an IBM cloud, which will almost certainly include a combination of x86 instruction set servers, proprietary IBM servers, and perhaps IBM mainframe clusters. This will be a “heterogeneous cloud” that is capable of hosting a wide variety of workloads, or possibly a “legacy system cloud” that is capable of running old Unix and mainframe workloads as well as

new Linux and Windows systems. Amazon can't do that, so a large number of enterprises that are interested in cloud computing will have reason to look to IBM or elsewhere. Such a data center will sacrifice some simplicity of management and economies of scale in order to be able to host a variety of data center applications. So far, IBM hasn't drawn a road map of how its cloud facilities will be architected, although it offers some specific products that would lend themselves to private internal cloud operations.

For the moment, we're left with the generic term *cloud computing*, whether Mark Hurd's audiences like it or not. It captures the notion of a widely available, low-cost service that is available on the Internet, which is the ultimate network "cloud."

Data and Identity Security at Stake

In addition to etymological opposition, resistance on additional grounds may come from CEOs, chief security officers, chief information officers, and/or database administrators, all of whom will want to know how it's possible to send the company's most valuable asset, its data, outside the firewall.

The answers will emerge over the next two years as the largest vendors and innovative start-ups tackle the problem. In many cases, instead of solving the problem in its own labs, one of the established vendors will buy a start-up with a piece of technology that resolves some additional piece of the puzzle.

Secure ways can be established to move sensitive data between the enterprise and a public cloud and handle it safely once it is there. But early cloud computing initiatives have not progressed to the point where they can do so and keep the data owner in compliance with all regulations, such as the PCI regulations that govern retail transaction data. As we've discussed, Amazon has announced that it will host "private" cloud computing within its public cloud infrastructure by imposing the use of a virtual private network—encrypted data moving over the public network—and other restrictions on how it deals with the "private" processing part of its business. This is not enough to meet businesses' objections to sending customer identity, health, or financial data outside the company, but it's a start. In the long run, if secure procedures are established and are proved to meet or exceed enterprise regulations, then the requirements may be changed to match the new conditions created by cloud computing. But revising regulations is a slow process. It will take established players—bankers, insurance professionals, equity traders—several years of illustrating the security of unregulated data exchange via cloud computing and lobbying for a review on regulated data to open the door to change.

After data management comes the ticklish issue of user identity as users migrate back and forth between applications in the enterprise and in the cloud. Already, Microsoft, Salesforce.com, and others say that they can provide a "federated identity"—a procedure by which one identity management system handles the requirements for user identity for several different

applications. The “federated” identity moves with the end user as she changes applications so that she doesn’t have to supply multiple user names and passwords. In effect, she logs in once and gets access to all the applications that she needs without having to do so again, regardless of whether the applications are in the enterprise or in the cloud. Microsoft says that it can do more than just enforce controls on end users as they cross the company boundary and move out into its Azure cloud. It can identify and authenticate users from other companies or from the public at large. It can use identity management from multiple directories, in addition to its own Active Directory, and use multiple types of identity confirmation. Its Identity Platform serves as a metadirectory for end user access control. Microsoft’s approach allows the application to demand a certain kind of unique identifier, a digital certificate, an Active Directory name and title, or a Windows token. The system retrieves that identifier, if it’s available, and submits it to the application, which accepts or rejects it. Identity under this system is “claims based,” or just a claimed identifier until the application accepts it. Some firms, such as start-up Simplified in Boulder, Colorado, say that they can also federate identity between enterprise and cloud users, relying on directory sources.

As this was being written, Fujitsu senior director Daniel Lawson said that in early 2010, his firm will launch cloud processing services at its Dallas, Texas, and Sunnyvale, California, data centers. The Dallas center will be secure enough to meet the PCI regulations. Fujitsu can do this by implementing secure FTP setups that ensure that the data that is sent arrives at

its destination intact and unchanged and is handled by secure processes afterward. That is, a process that might have been used by a financial institution to move data from one business unit to another has now been extended to the cloud.

Fujitsu goes a step further and says that it is planning ways to be able to handle Health Insurance Portability and Accountability Act (HIPAA) data, which includes patient health-care information. The privacy standards involved will make HIPAA a steep requirement to meet in the cloud environment, and such a development may still be years off. Fujitsu's Lawson acknowledges that not all health-care processing may be suitable for the cloud, but he believes that some of it can be executed there.

Unisys has also announced that it will provide services to support corporate cloud operations and is betting that its ability to deliver a more secure environment will give it a share of future cloud activity. Savvis, Verizon Business, and AT&T plan to offer VMware-based cloud services that go beyond the elementary controls contained in VMware's vCloud Express software. One area that they will emphasize is greater security of operations.

Avoiding Lock-In

Cloud advocates will then encounter their final barrier, vendor lock-in. Early cloud users will have to navigate the usual attempts by vendors to establish proprietary control in bids for

industry dominance. This vendor play for dominance has been a prominent feature of each previous phase of computing. Vendors have a right to seek a return on their investment. But I find it hard to believe that we really have to go through another protracted phase of attempted customer lock-in, the way the mainframe captured customers for IBM or Windows for Microsoft. After a certain period, these lock-ins have nothing to do with return on investment and everything to do with realizing long-term profits without having to compete on a level playing field. With luck, consumers won't put up with it this time around.

Until competition arises and populates the Internet with a daisy chain of cloud data centers around the globe, we are going to live through a period of attempts at dominance cloaked as proprietary initiatives. Proprietary initiatives in a free economy are a valuable thing; they're what's bringing us the first cloud data centers. But initiative is one thing and permanent, involuntary end user ensnarement is another. At the moment, there's practically no way for cloud customers to avoid some degree of lock-in.

For example, Amazon Web Services relied on open source code that was freely available in the public arena, such as the Linux operating system and the Xen hypervisor, to build its Elastic Compute Cloud (EC2), a move that made sense because freely downloadable open source code can be replicated over and over again as the cloud scales out, without incurring license charges. Although the code was based on Xen, Amazon Web Services tweaked the file format in which its EC2 cloud's virtual machines are built. It came up with a format,

the Amazon Machine Image, that was unique to EC2. The file format of a virtual machine allows it to be saved as a single file, combining the application, the operating system, and all its parts. That file, or virtual appliance, can then be stored, retrieved, and moved around like an iTunes or any other digital file. Amazon has published no details on what constitutes an AMI file or how it's different from other Xen hypervisor files. But it's different enough to prevent the standard, generally available Xen hypervisor from being able to run it.

If you like sending workloads to EC2, you accept the requirement that you use AMIs and find a way to build workloads in them. But if you decide that EC2 is no longer for you, those workloads are not easily extracted and moved someplace else, unless you are able to convert them on your own into some other format, such as VMware's Virtual Machine Disk format (VMDK), Microsoft's Virtual Hard Disk (VHD), or the neutral import/export Open Virtualization Format (OVF).

In addition, Amazon's AMI format is meant for use in the EC2 cloud only. It's not available for its customers to use in their internal data centers. In the long run, Amazon will surely provide tools that will make it easy to operate a hybrid cloud between EC2 and customer data centers and migrate workloads back and forth. But as of today, that's a stumbling block.

As lock-ins go, this is a modest one and, in various multi-step ways, reversible. But nevertheless, it exists as a barrier for the ill-prepared end user. To get workloads into EC2, Amazon supplies free tools to create AMIs. Tools to build AMIs are also available from independent suppliers, such as rPath. There are even vendors who will help you convert your existing virtual

machines into AMI workloads, and, for a fee, RightScale, FastScale, Elastra, and others will convert them or give you tools to convert them into formats that are capable of being run somewhere else. But this is not the frictionless back-and-forth migration that the cloud will need if it is to bring its full benefit to businesses. It's potential glitches and a need for services with fees attached.

There is also the previously mentioned neutral format, OVF. So far, Amazon has been noncommittal on this format. The Distributed Management Task Force standards body designed OVF to be a neutral format in which virtual machines may be moved around over the network. It is a mobility format, but the virtual machine can't actually be run in OVF. It's a freeze-dry pattern until the destination hypervisor is determined. Then OVF must be converted into that hypervisor's preferred proprietary format. A virtual machine cast in OVF can be moved under a VMware, Microsoft, or Citrix Systems hypervisor; each understands OVF and takes the files and builds them into the virtual machine of its choice. As it does so, it produces a virtual machine that is ready to run on its new host machine, unlike OVF. So the shared OVF format, which is used for importing files to a virtual machine host, represents a modest degree of cooperation among the competing vendors. As with AMIs, however, once you're in, it's hard to get out.

Why is this important? These barriers are being erected artificially. Providing a tool to convert AMIs back into OVFs would be relatively easy for Amazon, but it stays at arm's length from the prospect, just as technology pioneers before it have re-

mained aloof from neutral formats to preserve the proprietary advantage of being ahead of the crowd. But cloud computing didn't come about as the result of a breakthrough by any single vendor. There's a large public sphere contribution to the cloud in the standards of the Internet and Web services. In the long run, lack of ease of migration is going to slow the adoption of cloud computing until end users find so many ways around it that vendors back off from their proprietary formats. No one cloud is going to be good at every form of cloud computing, so users will naturally wish to move between clouds for different jobs. In the long run, those vendors that insist that the world conform to their (and only their) standard will find it increasingly difficult to find new customers.

Many people find Amazon's EC2 a useful place to do computing and know how to build AMIs. But even these users should stay watchful. New tools or start-up vendor services will spring into being to help you convert out of AMIs into OVF or one of the other familiar virtual machine formats. A request to your Amazon representative for a reverse converter, repeated enough times, might allow the message to sink in. Customers aren't quite in the driver's seat with cloud computing, but they're much closer to it than in the previous phases of computing.

And Amazon's per hour pricing has been competitive enough to set a de facto standard that other vendors have to try to meet. Microsoft positioned its Azure hourly charges only slightly higher than Amazon's, despite the fact that Microsoft can offer a more richly tooled environment with more cloud

services. By that standard, Amazon's success with AMIs has forced a major provider to a lower price point than it might have otherwise preferred.

Although the three leading x86 virtualization suppliers, VMware, Microsoft, and Citrix Systems, have agreed to support OVF, that doesn't mean that they've literally leveled the playing field among themselves. On the contrary, their support is rigged as a one-way street. Each is willing to convert a competitor's virtual machine into its own format, but it will not aid the customer in converting that virtual machine back into its original format or even back into OVF. Each supplier is thinking in terms of capturing a rival's customers, not making it easy for the customer to move workloads between clouds. In the previous phases of computing, even this modest level of cooperation would not have occurred, so OVF can be viewed as somewhat enlightened behavior. But as I say, one-way streets are just that and should not be confused with customer ease of transit.

Many people think that the possibilities of cloud computing will not be realized until there is a smooth, reliable path between the cloud and the enterprise data center and between different clouds. OVF and the current level of vendor cooperation aren't sufficient to guarantee that movement. So let the user beware. If you're a good customer of a cloud supplier, you should point out a specific purpose for which you want to use another vendor's cloud. If you get the cold shoulder, you might express some determination to find a way there—and not come back. The majority of your business is at stake. Sooner or later, the provider will get the message.

There are many reasons for businesses to adopt this demanding stance. Some cloud suppliers are specializing in setting up and tearing down software test environments, while capturing the test results (Skytap, SOASTA). Others may one day prove to be good at executing online transactions and storing those results securely. Others might provide a rich, hosted tool set for building software in the cloud (Salesforce.com, Microsoft, Engine Yard, IBM, Heroku) that will later be deployed to run in the same cloud or on Amazon's EC2. Such cloud "frameworks" can automate many underlying tasks, such as connection to the network or invoking specialized application programming interfaces, a way to speed software development.

For virtualization vendors and cloud suppliers to pretend that their customers need only one style of cloud computing (their style) is a false front. Business end users thrive on a diversity of choices, and vendors who stand in the way of diversity should be recognized as such and not rewarded. But the propensity to lock customers in remains strong.

Amazon is not alone in hanging on to the strength of a proprietary file format. The leading virtualization vendor VMware's VMDK is a proprietary format, with little information in the public sphere about it. VMware is a case where its technology strengths have kept customers from objecting too much.

Microsoft, in turn, wants to forestall VMware's dominance of the important and growing virtualization market. One of its few weapons for doing so is coordinating Hyper-V virtual machine operation in its Azure cloud with Hyper-V virtual machine operation in the enterprise data center. Doing so would

allow the creation of a hybrid Windows cloud and give Microsoft's approach to virtualization an advantage over VMware's.

VMware understands the link between enterprise virtualization and cloud computing, but it is not a cloud supplier itself. It is striving to generate a similar opportunity for its customers by seeding clouds that are compatible with its virtual machine format through vCloud Express. Announced in September, vCloud Express is a set of software and tools for a cloud vendor to use in setting up low-end cloud services, including self-provisioning, billing by the hour, and load balancing hundreds of VMware-based workloads. Terremark, Bluelock, RightScale, and Hosting.com are all similar cloud service providers or front ends to other service providers, who say that they are implementing vCloud Express.

As a sign of how crucial success on this front is to VMware, it has made public its vCloud application programming interface (API), which specifies how any third party can connect to a vCloud Express supplier. It submitted a specification for vCloud Express to a standards body in the fall of 2009. That body was the Distributed Management Task Force (DMTF), the same standards group that produced OVF. VMware's submission makes its API a published specification that is headed toward becoming a public standard, a step that it hasn't taken with its virtual machine file format and other proprietary technologies. The vCloud API is the first such API from any cloud vendor to be submitted for standardization. (Fujitsu followed with its cloud API in December.)

Other cloud suppliers are seeking to capitalize on VMware's support for cloud computing. AT&T Synaptic Compute cloud,

Verizon Business, and Savvis all say that they will create more sophisticated cloud services, including in-depth security, that will host VMware virtual machines. AT&T actually launched its ability to host VMware virtual machines in June 2009. Mike Crandell, CEO of RightScale, says that his firm will create virtual machine templates that will allow a server, after it is configured by the customer, to be deployed to the cloud of the customer's choice. So far, two destinations are available: EC2 and Rackspace. In addition, RightScale will be able to configure workloads in the virtual machines of any of the three major vendors. The idea of being able to deploy servers to various clouds using different formats is likely to become a regular feature of front-end service providers.

On another front, Citrix Systems and Microsoft, who are close business partners, have both agreed to support Microsoft's VHD file format, combining the weight of the number two and number three vendors in x86 virtualization to counter VMware's better-established VMDK. Microsoft Azure will run the VHD file format. However, it's not compatible with VMware's ESX hypervisor or VMDK file format. It's the conversion problem again: VMware customers will have to find a way to convert if they are seeking a cloud based on VHD, and vice versa.

So far, few VMware customers have shown a tendency to migrate. VMware, the virtualization market leader with \$2 billion in revenues in 2009, keeps advancing the capabilities of the management environment that now surrounds its virtual machines in the enterprise. Even so, the virtualization market is expanding so rapidly that it's hard to say what it will

look like two or three years from now. Only 16 percent of data center applications or “workloads” have been virtualized, according to Gartner. Thus, much of the market remains up for grabs. Gartner predicts that 50 percent of data center workloads will be virtualized by 2012, so this picture is going to change.

All this competition to establish a dominant virtual file format is actually an indicator that cloud computing encourages open standards. In another bid to increase virtualization of servers with Microsoft’s Hyper-V, not only has Microsoft teamed up with Citrix to back VHD, but it has also promised that VHD will remain an open format, not subject to changes that leave the customer faced with the need to upgrade to a new product and subject to new license charges. It does so with a nonbinding but highly public statement: its Open Specification Promise.

The pressure of VMware’s current virtualization dominance has prompted Microsoft to adopt a stance of being more open than VMware on the virtual machine file format. The Open Specification Promise is different from actually putting a specification in the open under the authority of a standards body. Nonetheless, having some guarantee of openness, regardless of how it came about, is preferable to having a purely proprietary spec. Microsoft’s stance, and its growing influence with Citrix Systems in the virtualization market, may one day force VMware to follow suit with a greater openness on its VMDK.

What’s most important here is to realize that business users’ virtualization choices will end up guiding their cloud

decisions. When users are looking to move workloads between the data center and the cloud, compatible virtual machine formats will be an asset; incompatible ones, a drawback. The differences between VMDK, VHD, and AMI are small. They could be collapsed into one publicly referenced standard, allowing ease of migration between clouds. But that would open the dominant vendors to level playing field competition. I do not expect to see such a thing happen until cloud computing becomes widely established and the locus of competition moves to a new front. (For lock-in of a completely different sort, see Appendix C. The editor of *InformationWeek's* "Plug into the Cloud" blog, John Foley, has illustrated how the unwary can be locked into a cloud simply by the price of trying to move one's data out.)

One way to counter the vendor's proprietary interest, however, is for customers to form groups that list their own preferences and use them to serve notice to the vendors. The best form of pressure is a paying customer pointing out the advantage of ease of movement between clouds. If this mobility is granted sooner rather than later, the immense potential of cloud computing can be realized sooner as well, and I doubt that competent vendors would be injured by such a development. User groups often produce spokespersons who are skilled at producing such a message.

In 2007, AMD's director of software development, Margaret Lewis, in a master stroke of stagecraft, if not statecraft, put representatives of VMware, XenSource, and Microsoft on stools on a raised platform at the end of a San Francisco virtualization conference, then filmed the results. Each was asked

whether interoperability between their virtualization products was a good idea.

VMware's Patrick Lin, senior director of product management at the time; Microsoft's Bob Tenzar, director of product management for Windows Server; and John Bara, vice president of marketing at what was at the time XenSource (now part of Citrix Systems), all agreed that it would be better if the virtual machine formats could work together and said that they were working behind the scenes to make it happen. In a report on the occasion, I termed this evening declaration on the benefits of interoperability "virtual kumbaya." By night, we sing around the campfire; by day, we go our separate ways. Nevertheless, the big three are on the record as saying that they are working on interoperability.

Two years later, I was reminded of the backward state of the industry on this point when I attended the Cloud Computing Forum in San Francisco in February 2009 and asked a panel of cloud experts when we would achieve a shared virtual machine runtime format as well as the migration format OVF. The answers were diplomatic.

"I don't think we're holding back any genuine progress by not documenting the AMI format," said Amazon's Jeffrey Barr.

Joseph Tobolski, Accenture's director of cloud computing, who was on the panel, later backed up Barr. "Jeff's point is perfectly valid. You've got to wait until the time is right to reconcile those different formats," he said in an interview.

This panel illustrated the industry's understanding that vendors have a right to use proprietary formats until the mar-

marketplace undergoes a shakeout and everybody can tell who the winners are. If there's any reconciliation to be done, let it follow the marketplace decision.

At this event, Lewis defended virtualization vendors' practices as better than in the past. Citrix has aligned its format with Microsoft's VHD, Microsoft and Red Hat have agreed to support each other's operating systems in virtual machines, and the DMTF has published OVF, with everyone's assent. "We see our software partners working more cooperatively than they have in years. Agreements are being reached and alliances are being made," she said.

I concluded a blog entry on these responses by noting how easy it is for strong technology vendors to agree that it's reasonable to pursue their own interests, despite the fact that a simple remedy to a customer problem was at hand.

"Knowledgeable parties inside ongoing software concerns may have a disdain for those users, those small minded individuals, who just can't understand why things need to be done the way they are. But I for one say bring on those revolting end users. After this gang, I'd like to hear from them." I still think an end user revolt is one of the few ways to get powerful vendors to listen.

Rather than let this issue lie dormant, cloud users should acquaint themselves with several open source code options that are exerting pressure on the proprietary nature of cloud computing. In some cases, open source code will provide a means of knocking down closed doors and building a private cloud that interoperates with a proprietary one, regardless of

whether the cloud vendor has exposed its format. Open source code may prove to be one of the ways to gain mobility between clouds.

The Eucalyptus Project, which we introduced earlier, is offering cloud APIs that can mimic what the Amazon EC2 APIs do in simple functionality, including loading a workload, calling Simple Storage Services (S3), or employing the temporary Elastic Block Store. Using these Eucalyptus APIs means that a private cloud can interoperate with Amazon's EC2. Amazon must understand that it is in its interest to tolerate this open source code as a way to extend the future reach of EC2. It has made no move to block or otherwise object to the Eucalyptus implementers.

Ubuntu, the Linux-based open source operating system from Canonical, now includes the Eucalyptus open source code as part of its package. Canonical and Eucalyptus Systems, the firm formed from the Eucalyptus Project, offer consulting services on how to build a private cloud that is compatible with Amazon's.

Eucalyptus Systems is extending what the project's original open source code can do with additional proprietary products. The Eucalyptus APIs originally supported use of open source code hypervisors only [known as Kernel-based Virtual Machine (KVM) and Xen]. The product, Eucalyptus Enterprise Edition, adds support for VMware's ESX Server hypervisor. Enterprise Edition thus could become a widely used building block of the private cloud. In the past, a wall existed between VMware's virtual machines, which are built in a VMDK file format, and EC2's Amazon Machine Image (AMI) format.

The two formats do not build virtual machines in the same way and are incompatible. The Enterprise Edition software, however, invokes a converter that changes the VMware's VMDK virtual machine into an AMI recognized by EC2. A workload in the VMware private cloud can now migrate across the boundary to function in the Amazon cloud. This opens up a path for coordination between public and private clouds.

At this point, Eucalyptus has stopped short of trying to create look-alike APIs for some of Amazon's more advanced services, such as the SimpleDB database service, Amazon Elastic MapReduce, or Amazon Relational Database Service. Nevertheless, Eucalyptus has broken down several barriers to constructing the private cloud. Because Eucalyptus is open source code, its core APIs are in the public arena.

A related effort is Simple API for Cloud Application Services, an open source project led by Zend Technologies. It seeks to provide an API for types of services that are found in the public cloud, and then let different clouds support that API, if they choose to do so. Zend's aim is to allow an application running in an enterprise to invoke, say, a Simple API for storage and receive the storage service that is available from the cloud it's dealing with—if that cloud supports Simple API. Simple API may become a way to level the playing field and give new cloud service providers a shot at attracting business from emerging private clouds. Simple API already works across the Nirvanix Storage Delivery Network, a public cloud storage provider, and Amazon's S3. That means an application built to run in one cloud could be moved to another and make use of the same services without being changed.

It's still very early in the game, but these open source initiatives show how private clouds may soon be built and find the means to synchronize their operations with public clouds. In some cases, front-end management services, such as Skytap and RightScale, already accept and manage an enterprise's virtual workloads for the cloud, even if they are generated by different hypervisors. They may extend that ability and start directly navigating the man-made barriers between private cloud operations and the public cloud.

Forces Line Up behind Cloud Standards

The Distributed Management Task Force has launched an Open Cloud Standards Incubator in which it will host early work on specifications, APIs, and other candidates to become standards of cloud computing. In November 2009, the DMTF published a 21-page white paper, "Interoperable Clouds," which makes the point that we've been emphasizing throughout this chapter: "It is important for users to use standard interfaces to provide flexibility for future extensions and to avoid becoming locked into a vendor." This white paper can be found at http://www.dmtf.org/about/cloud-incubator/DSP_IS0101_1.0.0.pdf.

The Cloud Security Alliance seeks to promote shared standards and best practices in cloud computing security. It is partnering with the DMTF to cooperate on cloud systems management standards.

Another group, the Open Grid Forum, debates proposed standards for managing large clusters known as grids and teams up with the DMTF, the Cloud Security Alliance, the Storage Networking Industry Association, and the Open Cloud Consortium to discuss standards for cloud computing. Many cloud vendors and a few cloud users belong to these groups. “Fostering trust in cloud computing services is a key criteria for enabling its growth,” said Jim Reavis, cofounder of the Cloud Security Alliance. This is true, but unless these groups enlist the support of the market leaders, they will end up talking to one another as cloud customers march off to one vendor or another’s proprietary drum. Too often, the open standards bodies consist of the vendors who didn’t lead in a technology innovation but want a piece of the action. Open standards give them entrée to the market and allow them to invest in products that interact with those of the market leader, if they can get that leader to follow the standard.

Thus, Simple API, a potentially valuable approach to cross-cloud computing, is supported not only by little Zend Technologies, but also by IBM and Microsoft. The party that is missing among these backers is Amazon Web Services, which is by far the dominant supplier of public cloud infrastructure. The cloud customer needs to remain wary, shopping around, accepting some proprietary control when necessary to engage to the degree he wants to in cloud computing. But customers should never lose their willingness to fight lock-in.

Cloud suppliers themselves rely on the Internet, built on open standards such as Berkeley Internet Name Domain

(BIND) and Sendmail, and they frequently depend on open source code in their own infrastructure, which makes them half-open even when they'd rather not be. They understand the customer's interest in more open standards and ease of movement across vendors very well. But they won't move in that direction voluntarily. It's up to you, the cloud user, to object when they put barriers in your way. It's up to customers to pry open that door, already slightly ajar, that vendors lean against so persistently.