

# VMware Central Region Customer Newsletter

May 2011

## In this issue

- **VMware acquires Slide Rocket**
- **Desktop Update**
- **VMware Cloud Foundry**
- **Clinical Desktop Nirvana**
- **ESXi and PXE Manager**
- **CEO Paul Maritz Interview**
- **Microsoft Clustering Support**
- **Cloud Outages?**

## Another Acquisition?

- **Scott Sauer**



On April 26 2011 VMware made another cloud based acquisition with the SaaS based software company named SlideRocket. SlideRocket has developed innovative presentation solutions that leverage modern concepts of cloud computing, collaboration, social media, and mobile computing platforms. More than 20,000 customers and 300,000 users are currently using SlideRocket to more effectively build, deliver, and share presentations. SlideRocket is designed and built to take advantage of a modern web and cloud based end-user computing model that blurs the lines between presentations, websites, and multi-media content. This approach represents a fundamental shift from PC-era presentation solutions.

## Free Training!

### Transition to ESXi Course is now Available On-line

<http://www.vmware.com/go/esxi-free-training>



#### Module 1: Transition to ESXi Overview

- Inspect why ESXi and not ESX any more
- Questions that customers might ask – when moving from ESX to ESXi

#### Module 2: Compare and Contrast VMware ESXi and VMware ESX

- Identify differences between ESX and ESXi

#### Module 3: Install and Configure VMware ESXi

- Analyze VMware ESXi Installable and ESXi Embedded editions
- Install ESXi installable
- Identify the tasks and procedures required to migrate to ESXi

#### Module 4: Script Interfaces and Command Line Tools

- Discuss primary scripting interfaces
- Install vSphere Management Assistant (vMA)
- Identify the scripting interfaces and programming tools for ESXi
- Identify differences for the commands that were used with ESX and what to use with ESXi
- Recognize any custom commands implemented and how to substitute them with vCLI or power vCLI commands

#### Module 5: Management Tasks

- Identify hardware monitoring techniques
- Discuss system management
- Identify the VMware back up tools
- Distinguish the log files
- Setup centralized logging
- Recognize file management strategies
- Analyze lockdown mode
- Enable Tech Support Mode (TSM)

# VMware Central Region Customer Newsletter

May 2011

## End User Computing

### - Josh Spencer

(Follow Josh on Twitter @VirtualSpence)

I joined the VMware, Inc. team just over sixth months ago. In that time, I've had the privilege of working with a variety of IT organizations as they evaluate VMware [View](#) and [ThinApp](#) technology. Having spent the previous four years implementing and utilizing these technologies in an IT shop, it has been fascinating to see the unique use cases, challenges that arise and the talented people there are to overcome them.

While I'd like to write a future article on some of the unique challenges and solutions in specific verticals such as healthcare or education, today I'll focus on a common issue that can affect anyone. The issue is how to take desktop virtualization technology from evaluation to production. There's a great [article by Jason Langone](#) on this topic. If you're now or have future plans to run a desktop virtualization pilot, I highly encourage you to have a look. It inspired me to write the next portion of this month's newsletter, which I hope you'll find to be complimentary.

There are four phases for evaluating and implementing desktop virtualization technology. While these phases are not set in stone, not always done sequentially, and vary in duration and complexity for different organizations, these seem to be the key pieces to successful implementations.

The first phase is **Proof of Concept** (POC). I do a lot of presentations where I talk about the power of VMware end user computing (EUC) technology, from the speed of provisioning [linked clone](#) desktop pools, to the efficiency of the [PCoIP protocol](#), and the simplicity of virtualizing an application and deploying it to an enterprise in a matter of minutes.

But who am I?! Those who know me understand that I'm a Desktop Engineer at heart who became so enthralled with VMware EUC technology that I couldn't pass up the opportunity to get out and talk to others about it. To the rest of you, I'm just a guy who works for VMware and wants you to use our products.

So in my mind, a POC is an opportunity for you to make sure that what I'm telling you is true; that you *can* actually build a pool of virtual machines that consume a fraction of the disk space normally required, that you *can* virtualize Internet Explorer 6 and have it running on your Windows 7 machines without any previous experience repackaging applications, that all work on your environment can *really* be done from a single management console.



So a POC is validation that the technology works in a known environment. A POC is *not* intended for use in production environments.

Many IT shops will conduct a POC themselves. For those who are constrained by resources, VMware has a fantastic [partner community](#) that can help you out.

The second phase is **Assessment**. This phase is critical and probably overlooked most often. A good assessment will give you invaluable data about the current state of desktops in your environment. By understanding what your users are doing today: applications used, IOPS consumed, peak hours, etc. you'll be in a *much* better position to plan your future infrastructure. I can't begin to tell you how many times I've heard surprised IT admins talking about things occurring in their environments that they were unaware of. There are tools like [Lakeside SysTrack](#) and [Liquidware Labs Stratusphere](#) that will provide a view in to your environment, helping you to make good decisions for planning and future expansion.

# VMware Central Region Customer Newsletter

May 2011

## Desktop Update

### (...Continued)

The third phase is **Pilot**. In a pilot we actually build a scaled down version of what will ultimately become a production implementation. Most organizations have a variety of use cases in which they'd like to leverage desktop virtualization. My advice is "pick one". The danger of picking too many use cases up front is that both end users and IT can become frustrated and give up. If users have two minor issues that IT can work through relatively quickly, things are successful. If ten different groups find two minor issues and IT takes months to resolve them...well, you get the idea.

The other suggestion I've got is that you *not* pick your hardest use case. If you've got that one group who constantly has problems, is not particularly happy with IT, and isn't going to be open to helping you "work out the bugs" of new technology, they're probably not the best candidates to start with. Instead, choose a group where you have good rapport, you'll be able to get solid and frequent feedback, and can see the value in what you're trying to do

Armed with data from the assessment and confidence in the technology from the POC and pilot, you're now ready for the final phase, **production implementation**. One of the common pitfalls is that a pilot works so well, it simply morphs in to the production environment. In most cases there are lessons learned and architectural designs that the IT admins would like to change before moving to production. But if your pilot *becomes* production, you may not get that chance. Ultimately, IT becomes forced in to managing a less-than-optimal environment.

Best of luck on your desktop virtualization initiatives! Remember, if you have questions, the VMware team is here to support you. Please don't hesitate to reach out to your local SE.

In other news, there is a great [FAQ for the PCoIP Secure Gateway](#) that's been posted. If you're considering an implementation it's well worth the read.

Been waiting on the View client for your Android device? Take a look at this [InfoWorld article](#)!

## VMware Cloud Foundry

= Hypervisor 2.0... The Freight Train is coming. -Aaron Sweemer



Ever since the acquisition of Springsource nearly two years ago, VMware has been generating a lot of excitement in the application development space. That excitement was kicked into high gear a few weeks ago when VMware announced the industry's first implementation of open PaaS, CLOUD FOUNDRY.



But I have a feeling much of that excitement is not felt or even understood by the average reader of this Newsletter. The reason largely has to do with the fact that most of us have an IT infrastructure/operations background. We are really good at troubleshooting low-level infrastructure stuff, we can rattle off the differences between RAID5 and RAID10, and we can debate iSCSI vs NFS until we are blue in the face. However, while we may be able to go crazy, Einstein deep into infrastructure technologies, there are very few us who would have a single clue about things like MVC software architecture, Object/Relational Mapping, or Dependency Injection.

Sure, some of us (and probably not many of us) may have the ability to create useful automation scripts in PowerShell or PERL, but that's a far cry from being able to create a full-blown application for end user consumption. And I'm here to tell you the application development world is, now more than ever, something we all need to embrace. Because worlds are colliding and CLOUD FOUNDRY is a glimpse of things to come.

# VMware Central Region Customer Newsletter

May 2011

## Cloud Foundry

### (...Continued)

#### What is CLOUD FOUNDRY?

Well you already know that Cloud Foundry is a PaaS, which means that at a very high level, you can think of Cloud Foundry as something on-par with Microsoft's Azure, or Google's AppEngine, or Salesforce's Force.com, or Engine Yard. Not familiar with those services? Or not 100% clear on what a PaaS is? OK, then for now, let's think of Cloud Foundry as a Hypervisor for cloud based applications. To be clear, I am NOT saying Cloud Foundry is a Hypervisor (because it is not); but let's just start there.

So today, what do we do when we want to deploy an application in our virtual datacenters? First, we start with a VM or a collection of VMs, and we either deploy them from a template, or we start from scratch and install an Operating System. Then, after some routine IT processes (patching, updating, configuration management, etc.) we either install and configure the application, or we hand it off to an application team to do the rest. The key point I want to make here is you start with an Operating System and build up from there. Meaning, the primary point of abstraction, the place upon which we begin to start build, is the Hypervisor.

How does this translate to Cloud Foundry? Well, Cloud Foundry allows us to start building applications directly on Cloud Foundry. There is no need to install an Operating System, nor is there a need to patch it, apply configurations, and install application components. That's all taken care of behind the scenes. So Cloud Foundry becomes the main point of abstraction, the place upon which we directly build our new cloudy applications. Another way to look at it would be, the Hypervisor switches our focus from managing hardware to managing VMs. Similarly, Cloud Foundry switches our focus from managing VMs to managing applications. In the former case, the hardware doesn't go away and in the latter case, the VMs won't go away either. But the way we interact with, manage and even think about hardware has fundamentally changed ... and so it will be with VMs and Cloud Foundry.



#### How is CLOUD FOUNDRY different from other PaaS offerings?

Now that we understand a bit about what Cloud Foundry is, I'm sure you're wondering what makes Cloud Foundry any different than the other PaaS offerings out there. The biggest differentiator can be summed up with one word: Choice.

Prior to Cloud Foundry, PaaS meant limited choices and ultimately PaaS meant vendor lock-in. Writing an application for Microsoft's Azure, as an example, means you will only be able to run your application on Azure. I suppose that's not a big deal if you're 100% committed to Microsoft's Azure solution and you're OK with an off premise only solution (Azure doesn't have an on premise option). This is true for any PaaS offering out there today. Whichever PaaS you go with, you are either limited in terms of the developer frameworks and application services the PaaS makes available to you, or you are limited in your deployment options (i.e. public vs. behind-my-firewall), or both. For the customers I talk to, this is a very big deal.

But the good news is Cloud Foundry brings a big change to all of this. Cloud Foundry has been designed eliminate vendor lock-in by offering:

- choice of developer frameworks
- choice of application services, and
- choice of deployment (internal vs. external cloud).

Of course you might be thinking, "that sounds great, but ultimately we'll have to run Cloud Foundry on top of vSphere, so we'll still be locked in to VMware." Well, you would be wrong. Yes, Cloud Foundry does run on vSphere, but it can run on non-VMware Infrastructure clouds as well.

Choice. It's super attractive, and it makes Cloud Foundry unique.

# VMware Central Region Customer Newsletter

May 2011



## Achieving Clinical

### Desktop Nirvana

- Scott Jobe



What do you get when you add an adaptive, intelligent, and real-time UDP-based protocol to a network without thoroughly understanding the end-user environment.....fire drills.....every time. Many of the discussions I have with my healthcare customers revolve around VMware View, our award winning Virtual Desktop offering. Of all environments begging for a virtual desktop solution, healthcare is very high on the list. The promise of true session roaming for clinicians while being able to work from an array of end-user devices (which may be in different locations) is clinical nirvana. In no way am I saying this is a plug and play operation; instead I am suggesting that it takes a calculated approach to achieve success.

There are many items to consider before jumping into a Clinical Desktop as a Service. These items will cross the boundaries of Server, Storage, and Desktop IT, clinical, and line-of-business staff. The new virtual desktop reality unlocks the ability to provide a desktop as-a-service that will fundamentally change how IT approaches backup, disaster recovery, and service levels. All this is happening while the clinicians expect a greater amount of freedom and control.

Adding a View Manager/Connection Server and View Composer to the existing environment can give customers a false sense of reality. As easy as it is to install the components, the real work lies within the analysis of the end-user environment and the impact of implementing a Virtual Desktop solution. Many of our customers have "stood up" a VMware View environment and tried to grow it organically, much like they did with the server environment over the past number of years. I can tell you that this approach generally leads to disaster. We frequently spend time talking to the IT teams about items like boot storms, security, the PCoIP protocol, and refresh/recompose cycles in a 24x7 operation.

Equally important is the conversation with clinical liaisons about end-user experience, session roaming, and general workflow. Unfortunately, the clinical conversations are the ones least likely to happen. This omission can derail any Virtual Desktop solution. The PCoIP protocol is unbelievably powerful but like I mentioned before, it is a real-time, adaptive protocol that can be best compared to VoIP. Many early adopters of VoIP failed in their implementation because they didn't architect a solution that could handle real-time requirements effectively. There are many parallels with PCoIP and VoIP. You cannot underestimate the impact on the network.

So how do we get there? The most successful implementations started with a Service Level and a thorough assessment of the environment including operational and technical readiness. Cross-functional teams that represent the stakeholders are incredibly important as well. It is not wise to view this as a technical project so a cross-functional team will help avoid this pitfall. Make sure to include the ISV's in the project as well. You'll find that several ISV's already support their applications in VMware View. Make sure to capture that knowledge by including the ISV's in your conversations. Tools such as Lakeside Software provide fantastic analytics that provide a clear picture of the "before" and "after" state of an implementation. VMware has a terrific partner network that contains talented engineers trained to follow best practices for implementations. Do not confuse a best-practice implementation of the software bits with a successful implementation of the Virtual Desktop solution. What I mean is that the workflow elements associated with the Virtual Desktop play a large part in the success and separating the technical team from the clinical teams will, in most cases, lead to failure. Lastly, focus on the clinicians and the service being provided. Using a calculated approach is the best way to avoid the fire drills that can arise.

To understand how VMware View can help your Healthcare environment, please visit: <http://www.vmware.com/solutions/industry/healthcare/index.html>

# VMware Central Region Customer Newsletter

May 2011

## ESXi and PXE Manager

- Tom MacKay

So here I am again, promoting ESXi enablement! This time, I would like to call your attention to some great work by Max Daneri from our Cloud Services group. Max has created a PXE manager for vCenter, available on our VMware Labs "Flings" site.

<http://labs.vmware.com/flings/pxe-manager>

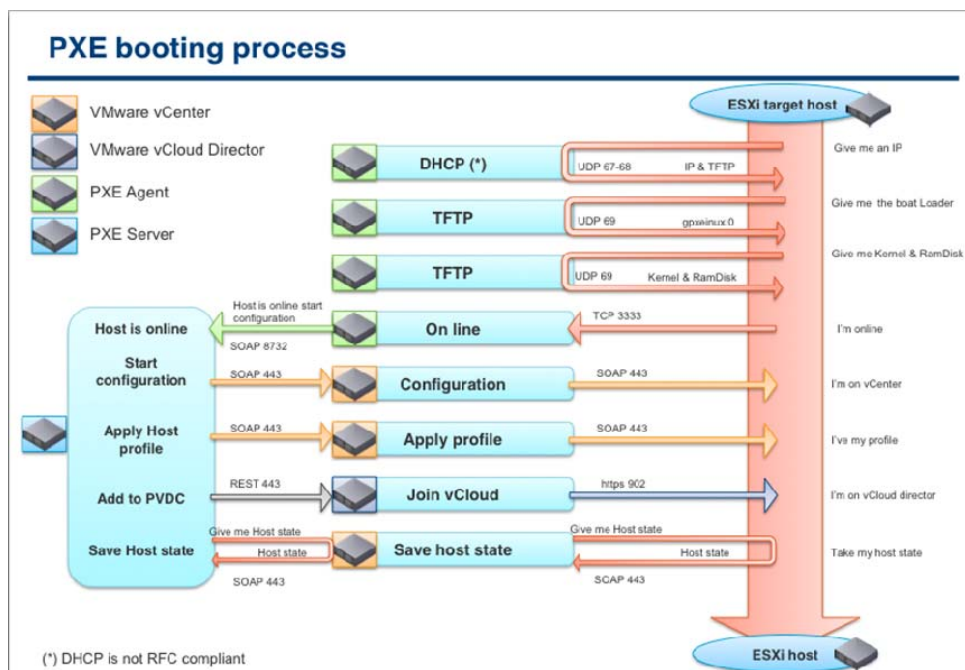
Not familiar with the Flings? Flings are tools that are created by our Engineers as "Cool Tools" that are not officially sanctioned or supported by VMware, but are made available through the VMware Labs site. Sort of spare time work for these guys...I particularly like ESXPlot, the vCMA (to use the Ipad VC Client), and the Thinapped vSphere Client. So take them for a spin, but remember, you cannot call support on these, though the engineers monitor the forum and answer questions about their respective flings...

Back to PXE Manager...PXE Manager supports automatic PXE booting for deployment and automatic customization of ESXi hosts. It also provides state management and patch management, apart from other tools like VUM. Features from the Flings

page:

- Automated provisioning of new ESXi hosts stateless and stateful (no ESX)
- ESXi host state (firmware) backup, restore, and archiving with retention
- ESXi builds repository management (stateless and statefull)
- ESXi Patch management
- Multi vCenter support
- Multi network support with agents (Linux CentOS virtual appliance will be available later)
- Wake on Lan
- Hosts memtest
- vCenter plugin
- Deploy directly to VMware Cloud Director
- Deploy to Cisco UCS blades

One of the questions asked to Max was what the difference between the Auto Deploy and he points out:



# VMware Central Region Customer Newsletter

May 2011

## ESXi and PXE Manager (...Continued)

pxe manager manage all lifecycle of your ESXi not only the deployment, this is the first difference. What you can do “more” with pxe manager: 1) stateless ESXi aren’t supported from VMware but a lot of time are very usefull for testing 2) manage more network segment 3) DHCP works in stealth mode (“no problem” with audit) 4) backup and restore of ESXi firmware (scheduled backup) 5) patching and updating of ESXi (you don’t need a pxe manager for any vCenter ) 6) run a memory test on hosts 7) very easy to configure and use 8) you can add a host to VMware Cloud Director instead of a vCenter folder 9) automatic deployment on Cisco UCS (very very cool feature, I’m preparing a video on this) A) vSphere plugin

with pxe manager you can save a state of host and restore it when you need to.

You WILL require DHCP in your DC, so that might be an issue for many administrators, but we see it becoming more prevalent throughout our customer base. Also, the requirements include .NET Framework, Visual C++ Agent, and Novell Mono 2.6.7.

PXE Manager leverages SOAP API calls to “build itself” after initially obtaining IP services and base loader and kernel capability. It registers itself with VC, and uses the REST API to join vCloud Director if required. Though we always mention that Flings may NEVER see the light of day as an official product, I am betting that this will be around and be improved upon in the near future! So if you have a few spare lab cycles, check out the Flings!

## VMware CEO Paul Maritz Podcast Interview

Do you know how many different versions of vSphere are in development at any given time?

Do you know what kind of mobile phone device CEO Paul Maritz carries?

How does Paul manage the explosive growth in a quickly expanding software company like VMware?

You can find the answer to all of these questions by tuning into the [Infosmack Podcast](#). This interview was conducted last August but highlights the VMware strategy, where we are going and Paul’s thoughts around this Journey to the Cloud.

**Infosmack**  
Podcasts

<http://infosmackpodcasts.com/infosmack-63-vmwares-paul-maritz/>



# VMware Central Region Customer Newsletter

May 2011



## VMware Clarifies Support for Microsoft Clustering - Jeff Szastak

VMware published KB Article 1037959 (<http://kb.vmware.com/kb/1037959>) on April 18, 2011 in an effort to clarify VMware's position on running Microsoft Clustering technologies on vSphere. Below is a snapshot of the support matrix published by VMware in the KB (always refer to KB 1037959 for the most current information).

For those familiar with VMware's previous position on Microsoft Clustering, you will notice a couple changes. First, VMware has made a distinction in Microsoft Clustering technologies by segmenting them into Shared disk and Non-shared Disk.

- Shared Disk – solution in which the the data resides on the same disks and the VMs share the disks (think MSCS)
- Non-shared Disk – solution in which the data resides on different disks and uses a replication technology to keep the data in sync (think Exchange 2007 CCR / 2010 DAG).

Next, VMware has extended support for Microsoft Clustering to include In-Guest iSCSI for MSCS.

For those interested in leveraging Microsoft SQL Mirroring, the KB states that VMware does not consider Microsoft SQL Mirroring a clustering solution and will fully support Microsoft SQL Mirroring on vSphere.

Under the Disk Configurations section of the KB, the KB discusses how if using VMFS, the virtual disks used as shared storage for clustered virtual machines must reside on VMFS datastores and must be created using the eagerzeroedthick option.

The KB provides detail on how to create the eagerzeroedthick disks for both ESX and ESXi via command line or GUI. Additional information regarding eagerzeroedthick can be found in KB article 1011170 (<http://kb.vmware.com/kb/1011170>). Something to note in KB 1011170, at the bottom of the article it states using the vmkfstools -k command you can convert a preallocated (eagerzeroed) virtual disk to eagerzeroedthick and maintain any existing data. Note, the VM must be powered off for this action. In closing, the VMware support statement exists to explicitly define what VMware will and will not support. It is very important for you to remember these support statements do not make any determination (either directly or indirectly) about what the software ISV (Independent Software Vendor) will and will not support. So be sure to review the official support statements from your ISV and carefully choose the configuration that makes sense for your organization and will be supported by each vendor.

	Microsoft Clustering on VMware	vSphere support	VMware HA support	vMotion DRS support	MSCS Node Limits	Storage Protocols support				Shared Disk	
						FC	In-Guest iSCSI	Native iSCSI	FCoE	RDM	VMFS
Shared Disk	MSCS with Shared Disk	Yes	Yes <sup>1</sup>	No	2	Yes	Yes	No	No	Yes <sup>2</sup>	Yes <sup>2</sup>
	Exchange Single Copy Cluster	Yes	Yes <sup>1</sup>	No	2	Yes	Yes	No	No	Yes <sup>2</sup>	Yes <sup>2</sup>
	SQL Clustering	Yes	Yes <sup>1</sup>	No	2	Yes	Yes	No	No	Yes <sup>2</sup>	Yes <sup>2</sup>
Non-shared Disk	Network Load Balance	Yes	Yes <sup>1</sup>	Yes	Same as OS/app	Yes	Yes	Yes	Yes	N/A	N/A
	Exchange CCR	Yes	Yes <sup>1</sup>	Yes	Same as OS/app	Yes	Yes	Yes	Yes	N/A	N/A
	Exchange DAG	Yes	Yes <sup>1</sup>	Yes	Same as OS/app	Yes	Yes	Yes	Yes	N/A	N/A

# VMware Central Region Customer Newsletter

May 2011

## Cloud Outages?

-Chris Everett



We would be remiss if we didn't talk about Amazon's outage (since everyone else has) this past month. I found several blogs that didn't want to put the blame on Amazon for the Outage and the blame on customers for not designing their application to handle cloud outages. Just like an enterprise data center, clouds WILL have outages. Public Clouds are run by humans on technology made by humans, so outages are unfortunately inevitable. I guess the difference of where to put failover architecture is where I differ from several of the blogs that have been written. Many of the blogs talk about how to write an application so the application can handle failures in the cloud. Even though this strategy is effective (and required with the Amazon cloud), it puts the effort to architect availability and places typical infrastructure requirements on the developer. This, in turn, also takes away focus away from business functionality for the application. We at VMware believe that infrastructure functions should remain at the infrastructure layer as much, as possible. This frees up the developers to focus on developing for business functionality and they don't need to think about infrastructure issues. In the case of Amazon, their infrastructure doesn't have some "basic" high availability features that are inherent with a VMware cloud. Amazon has a data replication strategy that in "theory" should provide data redundancy, however this failed because of a networking configuration change. The networking change routed replication to a secondary network that couldn't handle the replication load. This then caused a domino effect whereas the replication APIs and corresponding queues got flooded. Amazon then had to then turn off the replication and this then caused some data loss for several customers.

So, what does this mean for your "public cloud" decision? We believe that there are many advantages to choosing a VMware based cloud. This provides a customer with "options". A customer has the option to run workloads at many different cloud providers as well as run workloads, internally. If one specific cloud provider is having problems, a customer has the ability to have workloads run somewhere else. Customers at Amazon, don't have that option. In addition, if you have applications that are already written and aren't "cloud aware", they certainly aren't going to be written specifically for cloud availability. A VMware cloud has some availability "built in" at the infrastructure layer, with features that we are all familiar with. Including HA, DRS, vmotion and in the future even SRM. An application doesn't have to be written to take advantage of these features or even be aware that these features exist in the infrastructure. VMware based clouds can run applications that have no knowledge that they are running in an IaaS cloud.

We certainly believe that new applications that understand that they are running in a cloud do have advantages, and that developers will code some infrastructure cloud features into their applications. For example, cloud application can "auto scale" by deploying copies of portions of themselves, automatically. Cloud aware application can also provide some specific actions to work around infrastructure outages. However, but why have to code basic availability into every cloud based application? The Amazon outage has caused everyone to think about "cloud" and what it really means. A "public cloud" certainly doesn't mean that it is infallible. But then again, are we surprised if we work in the IT industry? Unfortunately, there will be more outages with cloud providers, including VMware based cloud providers. Will you be "locked into" your public cloud provider when that happens?