

# Four Ways to Achieve Quick Wins in the Cloud

More companies than ever want to use the cloud as part of their overall IT strategy. To do so, they often look to achieve some quick wins in the cloud to demonstrate its value. Achieving these quick wins also serves to give them some practical hands on experience in the cloud. Incorporating the cloud into your backup and disaster recovery (DR) processes may serve as the best way to get these wins.



Any company hoping to get some quick wins in the cloud should first define what a “win” looks like. For the purposes of this blog entry, a win consists of:

- Fast, easy deployments of cloud resources
- Minimal IT staff involvement
- Improved application processes or workflows
- The same or lower costs

Here are four ways for companies to achieve the quick wins in the cloud through their backup and DR processes:

## **#1 – Take a Non-disruptive Approach**

When possible, leverage your company’s existing backup infrastructure to store copies of data in the cloud. All enterprise backup products such as backup software and

deduplication backup appliances, save one or two, interface with public clouds. These products can store backup data in the cloud without disrupting your existing environment.

Using these products, companies can get exposure to the public cloud's core compute and storage services. These are the cloud services companies are most apt to use initially and represent the most mature of the public cloud offerings.

## **#2 – Deduplicate Backup Data Whenever Possible**

Public cloud providers charge monthly for every GB of data that companies store in their respective clouds. The more data that your company stores in the cloud, the higher these charges become.

Deduplicating data reduces the amount of data that your company stores in the cloud. In so doing, it also helps to control and reduce your company's monthly cloud storage costs.

## **#3 – Tier Your Backup Data**

Many public cloud storage providers offer multiple tiers of storage. The default storage tier they offer does not, however, represent their most cost-effective option. This is designed for data that needs high levels of availability and moderate levels of performance.

Backup data tends to only need these features for the first 24 – 72 hours after it is backed up. After that, companies can often move it to lower cost tiers of cloud storage. Note that these lower cost tiers of storage come with decreasing levels of availability and performance. While many backups (over 99%) fall into this category, check to see if any application recoveries occurred that required data over three days old before moving it to lower tiers of storage.

## #4 – Actively Manage Your Cloud Backup Environment

Applications and data residing in the cloud differ from your production environment in one important way. Every GB of data consumed and every hour that an application runs incur costs. This differs from on-premises environments where all existing hardware represents a sunk cost. As such, there is less incentive to actively manage existing hardware resources since any resources recouped only represent a “soft” savings.

This does not apply in the cloud. Proactively managing and conserving cloud resources translate into real savings. To realize these savings, companies need to look to products such as Quest [Foglight](#). It helps them track where their backup data resides in the cloud and identify the application processes they have running. This, in turn, helps them manage and control their cloud costs.

Companies rightfully want to adopt the cloud for the many benefits that it offers and, ideally, achieve a quick win in the process. Storing backup data in the cloud and moving DR processes to the cloud provides the quick win in the cloud that many companies initially seek. As they do so, they should also ensure they put the appropriate processes and software in place to manage and control their usage of cloud resources.

---

## Two Hot Technologies to Consider for Your 2019

# Budgets

Hard to believe but the first day of autumn is just two days away and with fall weather always comes cooler temperatures (*which I happen to enjoy!*) This means people are staying inside a little more and doing those fun, end of year activities that everyone enjoys – such as planning their 2019 budgets. As you do so, solutions from BackupAssist and StorMagic are two hot new technologies for companies to consider making room for in the New Year.

## [BackupAssist 365.](#)

[BackupAssist 365](#) backs up files and emails stored in the cloud. While backup of cloud-based data may seem rather ho-hum in today's artificial intelligent, block chain obsessed, digital transformation focused world, it solves a real world that nearly every size organization faces: how to **cost-effectively** and **simply** protect all those pesky files and emails that people store in cloud applications such as DropBox, Office 365, Google Drive, OneDrive, Gmail, Outlook and others.

To do so, [BackupAssist 365](#) adopted two innovative yet practical approaches to protect files and emails.

- First, it interfaces directly with these various cloud providers to backup this data. Using your login permissions (*which you provide when configuring the software,*) BackupAssist 365 accesses data directly in the cloud. This negates the need for your server, PC, or laptop to be turned on when these backups occur so backups can occur at any time.
- Second, it does **cloud-to-local** In other words, rather than running up more data transfer and network costs that come with backing up to another cloud, it backs the data backup to local storage on your site. While that

may seem a little odd in today's cloud-centric world, companies can get a great deal of storage capacity for nominal amounts of money. Since it only does an initial full backup and then differential backups thereafter, the ongoing data transfer costs are nominal and the amount of storage capacity that one should need onsite equally small.

LICENSING	
USERS	Per user per month (USD)
First 24 Users	\$1.00
25-49 Users	\$0.95
50+	\$0.90

Perhaps the best part about [BackupAssist 365](#) is its cost (*or lack thereof.*) [BackupAssist 365](#) licenses its software on a per user basis with each user email account counting as one user license. However, this one email account covers the backup of that user's data in any cloud service used by that user. Further, the cost is only \$1/month per user with a decreasing cost for greater number of users. In fact, the cost is so low on a per user basis, companies may not even need to budget for this service. They can just start using it and expense their credit cards to keep it below corporate radar screens.

### ***StorMagic SvSAN***

The [StorMagic SvSAN](#) touches on another two hot technology trends that I purposefully (*or not so purposefully*) left out above: hyperconverged infrastructure or HCI and edge computing. However, unlike many of the HCI and edge computing plays in the marketplace such as Cisco HyperFlex, Dell EMC VxRail, and Nutanix, StorMagic has not forgotten about cost

constraints that branch, remote, and small offices face.

As Cisco, Dell EMC, Nutanix and others chase the large enterprise data center opportunities, they often leave remote, branch, and small offices with two choices: pay up or find another solution. Many of these size offices are opting to find alternative solutions.

This is where StorMagic primarily plays. For a less well-known player, they play much bigger than they may first appear. Through partnerships with large providers such as [Cisco](#) and [Lenovo](#) among others, StorMagic comes to market with highly available, two-server systems that scale across dozens, hundreds, or even thousands of remote sites. To get a sense of StorMagic's scalability, walk into any of the 2,000+ Home Depots in the United States or Mexico and ask to look at the computer system that hosts their compute and storage. If the Home Depot lets you and you can find it, you will find a StorMagic system running somewhere in the store.

The other big challenge that each StorMagic system also addresses is security. Because their systems can be deployed almost anywhere in any environment, it does make them susceptible to theft. In fact, in talking to one of its representatives, he shared a story where someone drove a forklift through the side of a building and stole a computer system at one of its customer sites. Not that it mattered. To counter these types of threats, StorMagic encrypts all the data on its HCI solutions with its own software that is FIPS 140-2 compliant.

Best of all, to get these capabilities, companies do not have to break the bank to acquire one of these systems. The list price for the Standard Edition of the SvSAN software, which includes 2TB of usable storage, high availability, and remote management, is \$2,500.

As companies look ahead and plan their 2019 budgets, they need

to take care of their operational requirements but they may also want to dip their toes in the water to get the latest and greatest technologies. These two technologies give companies the opportunities to do both. Using BackupAssist 365, companies can quickly and easily address their pesky cloud file and email backup challenges while StorMagic gives them the opportunity to affordably and safely explore the HCI and edge computing waters.

---

## Proven Investment Principles Can Guide Your Cloud Strategy

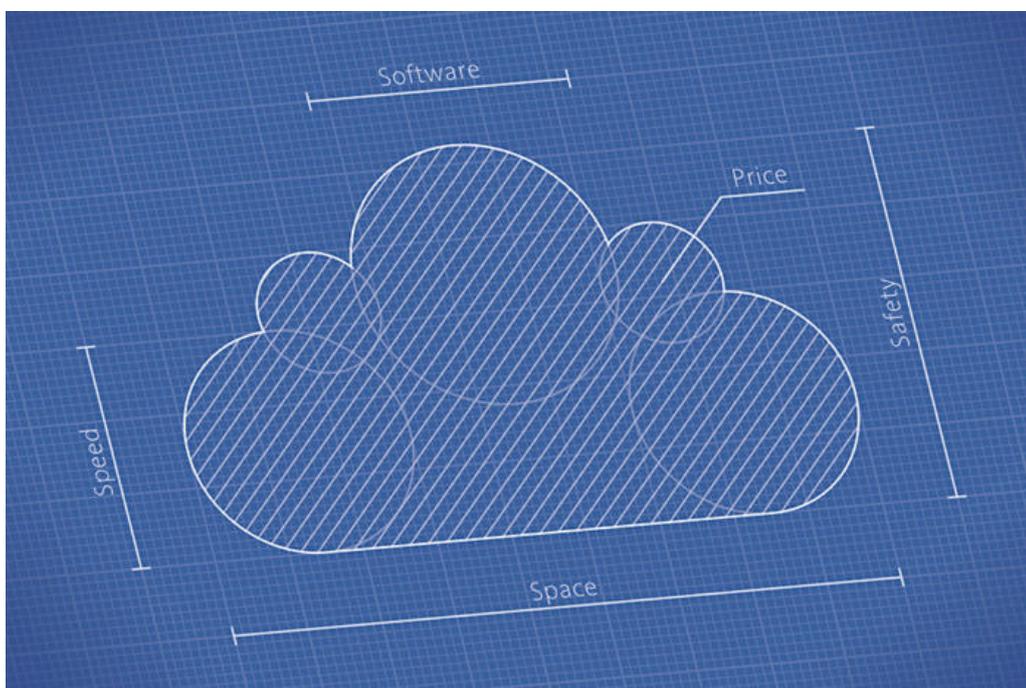
Living in [Omaha, Nebraska](#), one cannot help but be influenced by [Berkshire Hathaway](#) and its CEO, [Warren Buffett](#), one of the wealthiest men in the world, when it comes to making investment decisions. However, the process that Berkshire Hathaway uses to make investment decisions has multiple other applications that include guiding you in making decisions about your cloud strategy.

If there is a company and an individual that epitomize the concept of “*buy and hold*”, they are [Berkshire Hathaway](#) and [Warren Buffet](#). Their basic premise is that you thoroughly research a stock before making an investment decision. As part of that research, you investigates the financials of the company, its management team, its reputation, and the products and/or services it offers. Then you determine the type of growth that company will experience in the future. Once that decision is made, you then buy and hold it for a long time.

However, buy-and-hold is not the only principle that Warren Buffett follows. His first rule of investing is: ***Never lose***

**money.**

Companies should apply variations of both these principles when creating a cloud strategy. Once a company initiates and/or moves applications and/or data into the cloud, odds are that they will “*buy-and-hold*” them in the cloud for a long time assuming service levels and pricing continue to make sense. The more applications and data they store with a cloud provider, the more difficult it becomes for them to bring it back on-premise. Further, they can easily lose track of what data and applications their company has stored in the cloud.



The good and bad news is that public cloud providers such as [Amazon](#), [Google](#), and [Microsoft](#) have made and continue to make it easier than ever to get started with your cloud strategy as well as migrate existing applications and data to the cloud. This ease of implementing a cloud strategy can prompt organizations to bypass or shortcut the due diligence that they should take before placing applications and data in the cloud. Unfortunately, this approach leaves them without clearly defined plans to manage their cloud estate once it is in place.

To avoid this situation, here are some “*investment*” principles

to follow when creating a cloud strategy to improve your chances of success to get the return from the cloud that you expect.

1. ***Give preference to select proven, supported services from the cloud provider for critical applications and data.*** Most organizations when they move need to start with the basics such as compute, networking, security, and storage. These services are the bread and butter of IT and are the foundation for public cloud providers. These have been around for years, are stable, and are likely not going anywhere. Organizations can feel confident about using these cloud services for both existing and new applications and data and should expect them to be around for a long time to come.
2. ***Shy away from “speculative” technologies.*** Newly and recently introduced [Amazon](#) services such as [Lambda](#) (*serverless computing*), [Machine Learning](#), [Polly](#) (*text-to-voice*), and [Rekognition](#) (*visual analysis of images and videos*) among others sound (*and are*) exciting and fun to learn about and use. However, they are also the ones that cloud providers may abruptly change or even cancel. While organizations use them in production, companies just moving to the cloud may only want to use them with their test and dev applications or stay away altogether until they are confident they are stable and will be available indefinitely.
3. ***Engage with a trusted advisor.*** Some feedback that [DCIG](#) has heard is that companies want a more orchestrated roll-out of their computing services in the cloud than they have had on-premise. To answer that need, cloud providers are working to build out partner networks which have individuals certified in their technologies to include helping with the initial design and deployment of new apps and data in the cloud as well as the subsequent migration of existing applications and data to the cloud.

4. **Track and manage your investment.** A buy-and-hold philosophy does not mean you ignore your investment after you purchase it. You track cloud services like any other investment so take the time to understand and manage the billings. Due to the multiple options provided by each cloud service, you may need to periodically or even frequently change how you use a service or even move some applications and/or data back on-premise.



As organizations look to create a cloud strategy and make it part of how they manage their applications and data, they should take a conservative approach. Primarily adopt cloud technologies that are stable, that you understand, and which you can safely, securely, and confidently manage. Leave more “speculative” technologies for test and dev or until such a time that your organization has a comfort level with the cloud. While the cloud can certainly save you money, time, and hassle if you implement a cloud strategy correctly, its relative ease of adoption can also cost you much more if you pursue it in a haphazard manner.

---

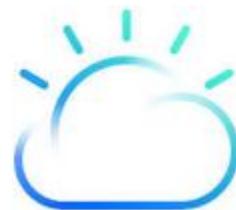
# Four Implications of Public Cloud Adoption and Three Risks to Address

Business are finally adopting public cloud because a large and rapidly growing catalog of services is now available from multiple cloud providers. These [two factors](#) have many implications for businesses. This article addresses four of these implications plus several cloud-specific risks.

## Implication #1: No enterprise IT dept will be able to keep pace with the level of services innovation available from cloud providers

The battle is over. Cloud wins. Deal with it.

Dealing with it does not necessarily mean that every business will move every workload to the cloud. It does mean that it is time for business IT departments to build awareness of the services available from public cloud providers. One way to do this is to tap into the flow of service updates from one or more of the major cloud providers.



For

Amazon Web Services, I like [What's New with AWS](#). Easy filtering by service category is combined with sections for featured announcements, featured video announcements, and

one-line listings of the most recent announcements from AWS. The one-line listings include links to service descriptions and to longer form articles on the AWS blog.

For Microsoft Azure, I like [Azure Updates](#). As its subtitle says, “One place. All updates.” The Azure Updates site provides easy filtering by product, update type and platform. I especially like the ability to filter by update type for [General Availability](#) and for [Preview](#). The site also includes links to the Azure roadmap, blog and other resources. This site is comprehensive without being overwhelming.

For Google Cloud Platform, its [blog](#) may be the best place to start. The view can be filtered by label, including by [announcements](#). This site is less functional than the AWS and Microsoft Azure resources cited above.

For IBM Cloud, the primary announcements resource is [What's new with IBM Cloud](#). Announcements are presented as one-line listings with links to full articles.

Visit these sites, subscribe to their RSS feeds, or follow them via social media platforms. Alternatively, subscribe to their weekly or monthly newsletters via email. Once a business has workloads running in one of the public clouds at a minimum an IT staff member should follow the updates site.

## **Implication #2: Pressure will mount on Enterprise IT to connect business data to public cloud services**

The benefits of bringing public cloud services to bear on the organization's data will create pressure on enterprise IT departments to connect business data to those services. There are many options for accomplishing this objective, including:

1. All-in with one public cloud

2. Hybrid: on-prem plus one public
3. Hybrid: on-prem plus multiple public
4. Multi-cloud (e.g. AWS + Azure)

The design of the organization and the priorities of the business should drive the approach taken to connect business data with cloud services.

### **Implication #3: Standard data protection requirements now extend to data and workloads in the public cloud**

No matter what approach is taken when embracing the public cloud, standard data protection requirements extend to data and workloads in the cloud. Address these requirements upfront. Explore alternative solutions and select one that meets the organization's data protection requirements.

### **Implication #4: Cloud Data Protection and DRaaS are on-ramps to public cloud adoption**

For most organizations the transition to the cloud will be a multi-phased process. Data protection solutions that can send backup data to the cloud are a logical early phase. Disaster recovery as a service (DRaaS) offerings represent another relatively low-risk path to the cloud that may be more robust and/or lower cost than existing disaster recovery setups. These solutions move business data into public cloud repositories. As such, cloud data protection and DRaaS may be considered on-ramps to public cloud adoption.

Once corporate data has been backed up or replicated to the cloud, tools are available to extract and transform the data into formats that make it available for use/analysis by that cloud provider's services. With proper attention, this can all be accomplished in ways that comply with security and data

governance requirements. Nevertheless, there are risks to be addressed.

## **Risk to Address #1: Loss of change control**

The benefit of rapid innovation has a downside. Any specific service may be upgraded or discarded by the provider without much notice. Features used by a business may be enhanced or decremented. This can force changes in other software that integrates with the service or in procedures used by staff and the associated documentation for those procedures.

For example, Office365 and Google G Suite features can change without much notice. This creates a “Where did that menu option go?” experience for end users. Some providers reduce this pain by providing an quick tutorial for new features within the application itself. Others provide online learning centers that make new feature tutorials easy to discover.

Accept this risk as an unavoidable downside to rapid innovation. Where possible, manage the timing of these releases to an organization’s users, giving them advance notice of the changes along with access to tutorials.

## **Risk to Address #2: Dropped by provider**

A risk that may not be obvious to many business leaders is that of being dropped by a cloud service provider. A business with unpopular opinions might have services revoked, sometimes with little notice. Consider how quickly the movement to [boycott the NRA](#) resulted in severed business-to-business relationships. Even an organization as large as the US Military faces this risk. As was highlighted in [recent news](#), Google will not renew its military AI project due in large part to pressure from Google employees.

Mitigate this risk through contracts and architecture. This is

perhaps one argument in favor of a hybrid on-prem plus cloud approach to the public cloud versus an all-in approach.

## **Risk to Address #3: Unpredictable costs**

It can be difficult to predict the costs of running workloads in the public cloud, and these costs can change rapidly. Address this risk by setting cost thresholds that trigger an alert. Consider subscribing to a service such as [Nutanix Beam](#) to gain granular visibility into and optimization of public cloud costs.

## **Its time to get real about the public cloud**

Many business are ready to embrace the public cloud. IT departments should make themselves aware of services that may create value for their business. They should also work through the implications of moving corporate data and workloads to the cloud, and make plans for managing the attendant risks.

---

# **Two Insights into Why Enterprises are Finally Embracing Public Cloud Computing**

In between my travels, doing research, and taking some time off in May, I also spent time getting up to speed on Amazon Web Services by studying for the AWS Certified Solutions Architect Associated exam in anticipation of DCIG doing more

public cloud-focused competitive research. While I know it is no secret that cloud adoption has taken off in recent years, what has puzzled me during this time is, *“Why is it **now** that have enterprises finally started to embrace public cloud computing?”*

From my first days as an IT user I believed that all organizations would eventually embrace cloud computing in some form. That belief was further reinforced as I came to understand virtualization and its various forms (compute, network, and storage.) But what has perplexed me to one degree or another ever since then is why enterprises have not more fully invested in these various types of virtualization and embraced the overall concept of cloud computing sooner.

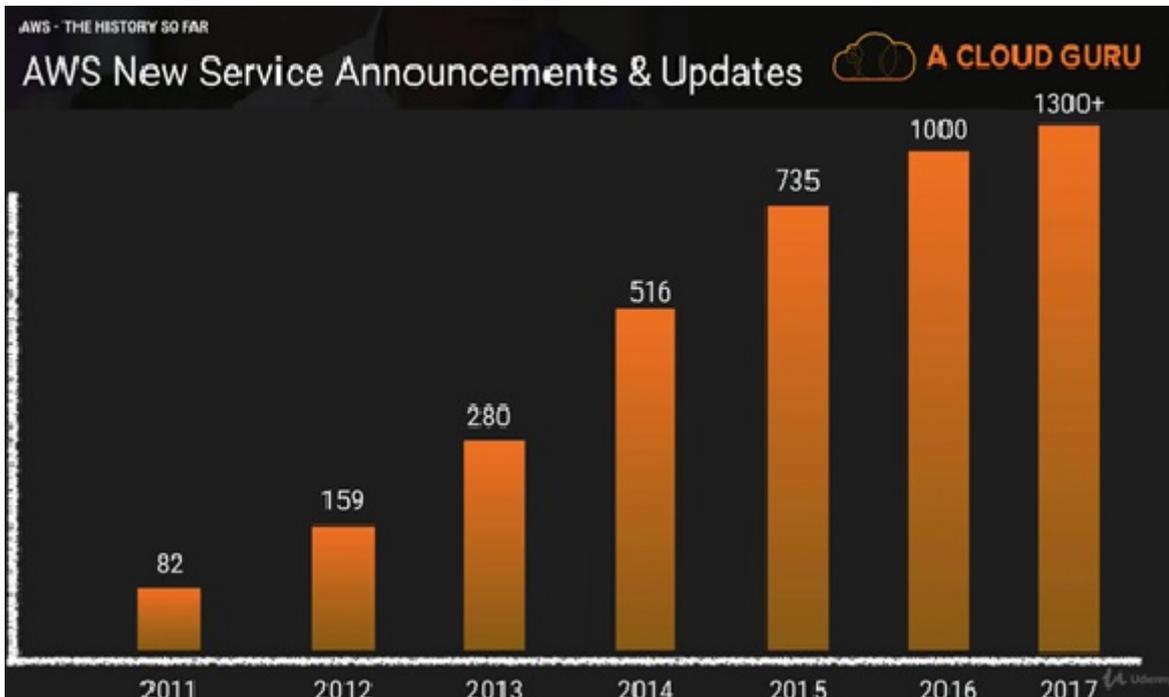
While there are various reasons for this, I sense the biggest reason is that most organizations view IT as a cost center. Granted, they see the value that IT has brought and continues to bring to their business. However, most organizations do not necessarily want to provide technology services. They would rather look to others to provide the IT technologies that they need and then consume them when they are sufficiently robust and mature for their needs.

Of course, establishing exactly when a technology satisfies these conditions varies for each industry. Some might rightfully argue that cloud computing has been around for a decade or more and that many organizations already use it.

But using public cloud computing for test, development, or even for some limited production deployments within an organization is one thing. Making public cloud computing the preferred or even the only choice for hosting new and existing applications is quite another. When this change in policy occurs within an enterprise, then one can say an enterprise has embraced public cloud computing. To date, only a relatively few enterprises have embraced the cloud computing at scale but I recently ran across two charts that help to

explain why this is changing.

The first chart I ran across was in one of the training videos I watched. This video included a graphic that showed the number of new service announcements and updates that [AWS](#) made each year from 2011-2017.



[Source:](#)

[A Cloud Guru](#)

It was when I saw the amount of innovation and changes that have occurred in the past three years at [AWS](#) that I got a better understanding as to why enterprises have started to embrace cloud computing at scale. Based on these numbers, [AWS](#) announced nearly five service announcements and/or updates **every business day** of 2017.

Many businesses would consider themselves fortunate to do five changes every month much less every day. But this level of innovation and change also explains why public cloud providers are pulling away from traditional data center in terms of the capabilities they can offer. It also explains why enterprises can have more confidence in public cloud providers and move more of their production applications there. This level of innovation also inherently communicates high degrees of

stability and maturity which is often what enterprises prioritize.

The other chart brought to my attention is found on Microsoft's website and provides a [side-by-side comparison](#) of [Microsoft Azure](#) to [AWS](#). This chart provides a high-level overview of the offerings from both of these providers and how their respective offerings compare and contrast.

Most notable about this chart is that it means organizations have another competitive cloud computing offering that is available from a large, stable provider. In this way, as an enterprise embraces the idea of cloud computing in general and chooses a specific provider of these services, they can do so with the knowledge that they have a viable secondary option should that initial provider become too expensive, change offerings, or withdraw an offering that they currently or plan to use.

Traditional enterprise data centers are not going away. However, as evidenced by the multiple of enhancements that [AWS](#), [Microsoft Azure](#), and others have made in the past few years, their cloud offerings surpass the levels of auditing, flexibility, innovation, maturity, and security found in many corporate data centers. These features coupled with organizations having multiple cloud providers from which to choose provide insight into why enterprises are lowering their resistance to adopting public cloud computing and embracing it more wholeheartedly.

---

# Hackers Say Goodbye to

# Ransomware and Hello to Bitcoin Mining

Ransomware gets a lot of press – and for good reason – because when hackers break through your firewalls, encrypt your data, and make you pay up or else lose your data, it rightfully gets people's attention. But hackers probably have less desire than most to be in the public eye and sensationalized ransomware headlines bring them unwanted attention. That's why some hackers have said goodbye to the uncertainty of a payout associated with getting a ransom for your data and instead look to access your servers to do some bitcoin mining using your CPUs.

A week or so ago a friend of mine who runs an Amazon Web Services ([AWS](#)) consultancy and reseller business shared a story with me about one of his clients who hosts a large SaaS platform in [AWS](#).

His client had mentioned to him in the middle of the week that the applications on one of his test servers was running slow. While my friend was intrigued, he did not at the time give it much thought. This client was not using his managed services offering which meant that he was not necessarily responsible for troubleshooting their performance issues.

Then the next day his client called him back and said that now all his servers hosting this application – test, dev, client acceptance, and production – were running slow. This piqued his interest, so he offered resources to help troubleshoot the issue. The client then allowed his staff to log into these servers to investigate the issue

Upon logging into these server, they discovered that all instances running at 100% also ran a [Drupal](#) web application. This did not seem right, especially considering that it was

early on a Saturday morning when the applications should mostly be idle.

After doing a little more digging around on each server, they discovered a mysterious multi-threaded process running on each server that was consuming all their CPU resources. Further, the process also had opened up a networking port to a server located in Europe. Even more curious, the executable that launched the process had been deleted after the process started. It was as if someone was trying to cover their tracks.

At this point, suspecting the servers had all been hacked, they checked to see if there were any recent security alerts. Sure enough. On March 28, 2018, [Drupal](#) issued a [security advisory](#) that if you were not running Drupal 7.58 or Drupal 8.5.1, your servers were vulnerable to hackers who could remotely execute code on your server.



However, what got my friend's attention is that these hackers did not want his client's data. Rather, they wanted his client's processing power to do bitcoin mining which is exactly what these

servers had been doing for a few days now on behalf of these hackers. To help their client, they killed the bitcoin mining process on each of these servers before calling his client to advise them to patch [Drupal](#) ASAP.

The story does not end there. In this case, his client did not patch [Drupal](#) quickly enough. Sometime after they killed the bitcoin mining processes, another hacker leveraged that same [Drupal](#) security flaw and performed the same hack. By the time his client came to work on Monday, there were bitcoin mining

processes running on those servers that again consumed all their CPU cycles.

What they found especially interesting was how the executable file that the new hackers had installed worked. In reviewing their code, the first thing it did was to kill any pre-existing bitcoin mining processes started by other hackers. This freed all the CPU resources to handle bitcoin mining processes started by the new hackers. The hackers were literally fighting each other over access to the compromised system's resources.

Two takeaways from this story:

1. ***Everyone is rightfully worried about ransomware but bitcoin mining may not hit corporate radar screens.*** I doubt that hackers want the FBI, CIA, Interpol, MI6, Mossad, or any other criminal justice agency hunting them down any more than you or I do. While hacking servers and "stealing" CPU cycles is still a crime, it probably is much further down on the priority list of most companies as well as these agencies.

A  
bitco  
in  
minin  
g  
hack  
may  
go  
unnot  
iced  
for  
long  
perio  
ds of  
time



and may not be reported by companies or prosecuted by these criminal justice agencies even when reported because it is easy to perceive this type of hack as a victimless crime. Yet every day the hacker's bitcoin mining processes go unnoticed and remain active, the more bitcoin the hackers earn. Further, one should assume hackers will only become more sophisticated going forward. Expect hackers to figure out how to install bitcoin mining processes that run *without* consuming all CPU cycles so these processes remain running and unnoticed for longer periods of time.

**2. *Hosting your data and processes in the cloud does not protect your data and your processes against these types of attacks.*** [AWS](#) has all the utilities available to monitor and detect these rogue processes. That said, organizations still need someone to implement these tools and then monitor and manage them.

Companies may be relieved to hear that some hackers have stopped targeting their data and are instead targeting their processors to use them for bitcoin mining. However, there are no victimless crimes. Your pocket book will still get hit in cases like this as Amazon will bill you for using these resources.

In cases like this, if companies start to see their [AWS](#) bills going through the roof, it may not be the result of their businesses. It may be their servers have been hacked and they are paying to finance some hacker's bitcoin mining operation. To avoid this scenario, companies should ensure they have the right internal people and processes in place to keep their applications up-to-date, to protect infrastructure from attacks, and to monitor their infrastructures whether hosted on-premise or in the cloud.

---

# Cool New Features and Offerings from AWS

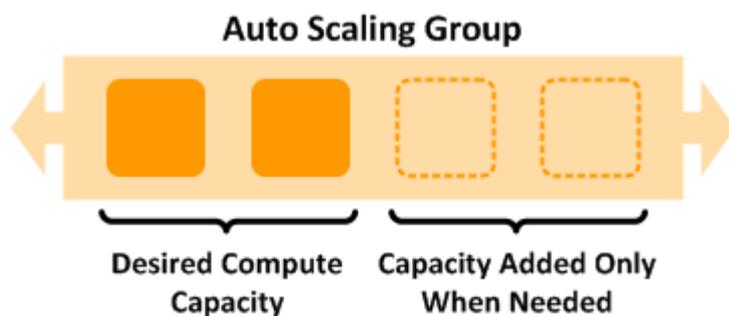
Amazon has made significant progress in the last few years to dispel the notion that Amazon Web Services' (AWS) primary purpose is as a repository for archives and backups. During this time, it has demonstrated time and time again it is well suited to host even the most demanding of production applications. However, what companies may still fail to realize is just how far beyond being a leading provider of cloud storage services that AWS has become. Here are some recent cool new features and offerings available from AWS that indicate how far it has come in terms of positioning itself to host enterprise applications of any type as well as satisfy specific enterprise demands.

- **Take a tour of Amazon's data centers – virtually.** As organizations look to host their mission critical applications, sensitive data, and regulated data with third party providers such as Amazon, the individuals who make these types of decisions to outsource this data have a natural inclination to want to physically inspect the data centers where this data is kept.

While opening up one's data center to visitors may sound good on the surface, parading every Tom, Dick, and Harry through a "secure site" potentially makes a secure site insecure. To meet this demand, Amazon now gives individuals the opportunity to take virtual tours of its data centers. Follow this [link](#) to take this tour.



- **Get the infrastructure features you need when you need them at the price you want.** One of the most challenging and frustrating aspects of managing any application within a data center is adapting to the application's changing infrastructure requirements. In traditional data centers, applications are assigned specified amounts of CPU, memory, and storage when they are initially created. However, the needs and behavior of the application begin to change almost as soon as it is deployed and to try to manually adapt the infrastructure to these constantly changing requirements was, at best, a fool's game.



### [Amazon Auto Scaling](#)

changes this paradigm. Users of this service can set target utilization levels for multiple resources to maintain optimal

application performance and availability even as application workloads fluctuate. The beauty of this service is that it rewards users for using it since it only charges them for the resources they use. In this way, users get better performance, optimize the capacity available to them and only use the right resources at the right time to control costs.

- **Amazon has its own Linux release.** Watch out Red Hat, SUSE, and Ubuntu – there is a new version of Linux in town. While [DCIG](#) has not yet taken the opportunity to evaluate and see how [Amazon Linux 2](#) compares to these existing, competing versions of Linux, perhaps what makes Amazon's release of Linux most notable is that it runs both on-premise and in the Amazon cloud. Further, it makes one wonder just how far Amazon will develop this version of Linux and will it eventually compete

head-to-head with the likes of VMware vSphere and Microsoft Hyper-V?

- **Corporate world: Meet Alexa.** Many of us are already familiar with the commercials that promote a consumer version of Alexa that enables us to order groceries, get answers to questions, and automate certain tasks about the home. But now Alexa has grown up and is entering the corporate world. Using [Alexa for Business](#), companies can begin to perform mundane, business-oriented tasks such as managing calendars, setting up meetings, reserving conference rooms, and dialing into meetings.



---

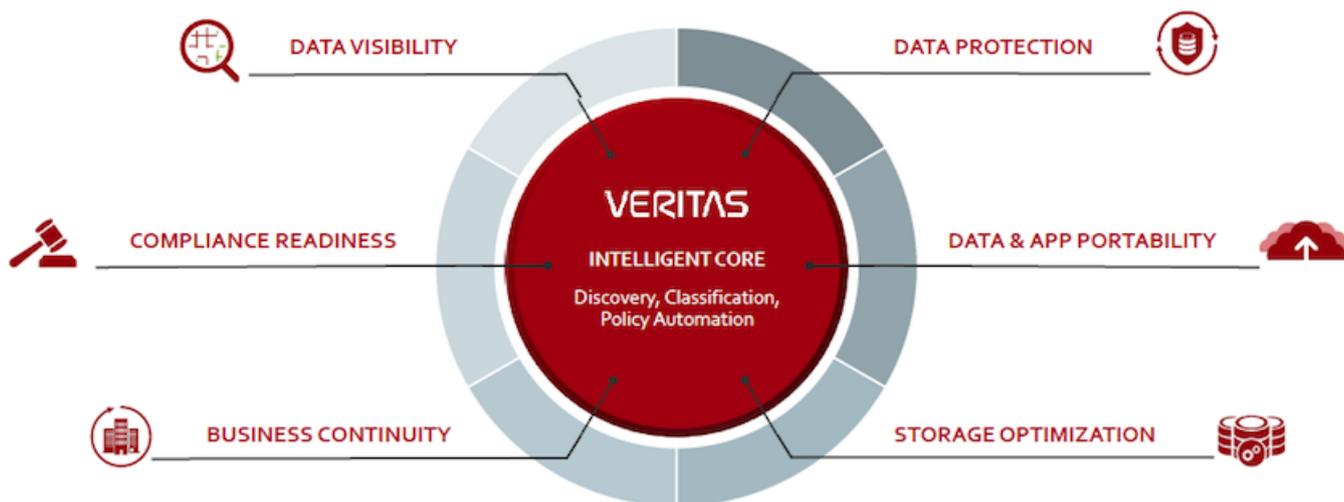
## **Veritas Delivering on its 360 Data Management Strategy While Performing a 180**

Vendors first started bandying about the phrase “*cloud data management*” a year or so ago. While that phrase caught my attention, specifics as what one should expect when acquiring a “*cloud data management*” solution remained nebulous at best. Fast forward to this week’s Veritas Vision 2017 and I finally encountered a vendor that was providing meaningful details as

to what cloud data management encompasses while simultaneously performing a 180 behind the scenes.

Ever since I heard the term cloud data management a year or so ago, I loved it. If there was ever a marketing phrase that captured the essence of how every end-user secretly wants to manage all its data while the vendor or vendors promising to deliver it commits to absolutely nothing, this phrase nailed it. A vendor could shape and mold that definition however it wanted and know that end-users would listen to the pitch even if deep down the users knew it was marketing spin at its best.

Of course, Veritas promptly blew up these pre-conceived notions of mine this week at Vision 2017. While at the event, Veritas provided specifics about its cloud data management strategy that rang true if for no other reason that they had a high degree of veracity to them. Sure, Veritas may refer to its current strategy as “360 Data Management.” But to my ears it sure sounded like someone had finally articulated, in a meaningful way, what cloud data management means and the way in which they could deliver on it.



Source: Veritas

The above graphic is the one that Veritas repeatedly rolls out when it discusses its 360 Data Management strategy. While

notable in that it is one of the few vendors that can articulate the particulars of its data management strategy, it more importantly has three important components to it that currently makes its strategy more viable than many of its competitors. Consider:

1. ***Its existing product portfolio maps very neatly into its 360 Data Management strategy.*** One might argue (*probably rightfully so*) that Veritas derived its 360 Data Management strategy from its existing product portfolio that it has built-up over the years. However, many of these same critics have also contended that Veritas has been nothing but a company with an amalgamation of point products with no comprehensive vision. Well, guess what, the world changed over the past 12-24 months and it bent decidedly bent in the direction of software. Give Veritas some credit. It astutely recognized this shift, saw that its portfolio aligned damn well with how enterprises want to manage their data going forward, and had the hutzpah to craft a vision that it could deliver based upon the products it had in-house.
2. ***It is not resting on its laurels.*** Last year when Veritas first announced its 360 Data Management strategy, I admit, I inwardly groaned a bit. In its first release, all it did was essentially mine the data in its own NetBackup catalogs. Hello, McFly! Veritas is only now thinking of this? To its credit, this past week it expanded the list of products to which to which its Information Map connectors can access to over 20. These include Microsoft Exchange, Microsoft SharePoint, and Google Cloud among others. Again, I must applaud Veritas for its efforts on this front. While this news may not be momentous or earth-shattering, it visibly reflects a commitment to delivering on and expanding the viability of its 360 Data Management strategy beyond just NetBackup catalogs.
3. ***The cloud plays very well in this strategy.*** Veritas

knows that plays in the enterprise space and it also knows that enterprises want to go to the cloud. While nowhere in its vision image above does it overtly say “cloud”, guess what? It doesn’t have to. It screams, “Cloud!” This is why many of its announcements at Veritas Vision around its [CloudMobility](#), [Information Map](#), [NetBackup Catalyst](#), and other products talk about efficiently moving data to and from the cloud and then monitoring and managing it whether it resides on-premises, in the cloud, or both.

One other change it has made internally (*and this is where the 180 initially comes in,*) is how it communicates this vision. When Veritas was part of Symantec, it stopped sharing its roadmap with current and prospective customers. In this area, Veritas has made a 180, customers who ask and sign a non-disclosure agreement (NDA) with Veritas can gain access to this road map.

Veritas may communicate that the only 180 turn it has made in the last 18 months or so since it was spun out of Symantec is its new freedom to communicate its road map to current and/or prospective customers. While that may be true, the real 180 it has made entails it successfully putting together a cohesive vision that articulates the value of products in its portfolio in a context that enterprises are desperate to hear. Equally impressive, Veritas’ software-first focus better positions it than its competitors to enable enterprises to realize this ideal.

---

# A Full Embrace of the Cloud Has Occurred ... Now the Challenge is to Successfully Get and Stay There

Ever since I got my first job in IT in the mid-1990's, everyone has used a cloud in some form. Whether they referred to it as outsourcing, virtualization, central IT, or in some other way, the cloud existed and grew but it did little to stem the adoption of distributed computing. Yet at some point over the past few years, the parallel growth of these two technologies stopped and the cloud forged ahead. This shift indicates that companies have now fully embraced the cloud but remain unclear about how best and how soon to transition their IT infrastructure to the cloud and then manage it once it is there.

One of my first jobs in IT was as a system administrator at a police department in Kansas. During my time there, I was intimately involved in a project that involved setting up a cloud that enabled it along with other police departments throughout the state to communicate with state agencies. Setting this cloud up would enable our department along with others to run background checks as well as submit daily crime reports. While we did not at that refer to this statewide network as a cloud, it did provide a means to send and receive data and centralize store it.

However, the data that the police department sent, received, and stored with various state agencies represented only a fraction of the total data that the department generated and used daily. There were also photos, files, Excel spreadsheets, accident and incident reports, and many other types of data that officers and civilians in the police department needed

and used to perform their daily duties. Since the state agencies did not need this data it was up to the police department to manage and house it.

This example is a microcosm of what happened everywhere. Private and public organizations would choose to store some data locally and only store certain data with cloud providers which, in the police department's case, were the systems provided by the various state agencies.

The big change that has occurred this decade and particularly over the past two years is that the need to host any applications or data on-premise has essentially vanished. This change has freed organizations of all sizes to fully embrace the cloud by hosting most if not all internal application processing and data storage with cloud providers.

Technology largely exists at the application, compute, network, operating system, security and storage layers that make it more cost-effective and efficient to host all applications and data with cloud providers rather than trying to continue to host it on premise. Further, the plethora of powerful endpoint mobile devices that are available as phones, desktops, tablets, and/or laptops along with ever larger network pipes make it easier than ever to access and manipulate centrally stored data anywhere at any time.

Organizations must accept ... and probably largely have ... that the technologies exist in the cloud to support even their most demanding applications. Further, these technologies are often more mature, cost-effective, and efficient than what they possess in-house.

The challenges before them are to now identify and execute upon the following:

1. Identify the right cloud provider or providers for them
2. Securely and successfully migrate their existing applications and data to the cloud

### 3. Manage their applications and data once hosted in the cloud

These objectives represent a fundamental shift in how organizations think and make decisions about their applications and data, and the IT infrastructure that supports them. This “cloud-first” view means that organizations must assume all new applications and data will end up in whole or in part in the cloud either initially or over time. As such, the new questions they must ask and answer are:

- How soon should their applications and data end up in the cloud?
- How much of their data should they put in the cloud versus retaining a copy onsite?
- If they choose not to put an application or data in the cloud, why not?

Organizations have officially embraced the cloud and what it offers as evidenced by the “cloud-first” policies that many have implemented that require them to deploy all new applications and data with cloud providers. However, migrating existing applications and data to the public, private, or hybrid clouds and then successfully managing all migrated applications and data in the cloud, as well as determining when to bring cloud-based applications and data out of the clouds, becomes more complicated.

Helping organizations understand these challenges and make the right choices will become a point of emphasis in DCIG’s blogs, research, and publications going forward to help organizations successfully migrate their data to the cloud, and then have a good experience once they get there.

---

# Facebook's Disaggregated Racks Strategy Provides an Early Glimpse into Next Gen Cloud Computing Data Center Infrastructures

Few organizations regardless of their size can claim to have 1.35 billion users, have to manage the upload and ongoing management of 930 million photos a day or be responsible for the transmission of 12 billion messages daily. Yet these are the challenges that Facebook's data center IT staff routinely encounter. To respond to them, Facebook is turning to a disaggregated racks strategy to create a next gen cloud computing data center infrastructure that delivers the agility, scalability and cost-effective attributes it needs to meet its short and long term compute and storage needs.

At this past week's [Storage Visions](#) in Las Vegas, NV, held at the [Riviera Casino and Hotel](#), Facebook's Capacity Management Engineer, [Jeff Qin](#), delivered a keynote that provided some valuable insight into how uber-large enterprise data center infrastructures may need to evolve to meet their unique compute and storage requirements. As these data centers daily may ingest hundreds TBs of data that must be managed, manipulated and often analyzed in near real-time conditions, even the most advanced server, networking and storage architectures that exist today break down.

Qin explained that in Facebook's early days it also started out using these technologies that most enterprises use today. However the high volumes of data that it ingests coupled with end-user expectations that the data be processed quickly and securely and then managed and retained for years (and possibly

forever) exposed the shortcomings of these approaches. Facebook quickly recognized that buying more servers, networking and storage and then scaling them out and/or up resulted in costs and overhead that became onerous. Further, Facebook recognized that the available CPU, memory and storage capacity resources contained in each server and storage node were not being used efficiently.

To implement an architecture that most closely aligns with its needs, Facebook is currently in the process of implementing a [Disaggregated Rack](#) strategy. At a high level, this approach entails the deployment of CPU, memory and storage in separate and distinct pools. Facebook then creates virtual servers that are tuned to each specific application's requirements by pulling and allocating resources from these pools to each virtual server. The objective when creating each of these custom application servers is to utilize 90% of the allocated resources to use them as optimally as possible.

Facebook expects that by taking this approach that, over time, it can save in the neighborhood of \$1 billion. While Qin did not provide the exact road map as to how [Facebook](#) would achieve these savings, he provided enough hints in his other comments in his keynote that one could draw some conclusions as to how they would be achieved.

For example, Facebook already only acquires what it refers to as "vanity free" servers and storage. By this, one may deduce that it does not acquire servers from the likes of Dell or HP or storage from the likes of EMC, HDS or NetApp (though Qin did mention Facebook did initially buy from these types of companies.) Rather, it now largely buys and configures its own servers and configures them by itself to meet its specific processing and storage needs.

Also it appears that Facebook may be or is already buying the component parts that make up server and storage such as the underlying CPU, memory, HDDs and network cabling to create its

next gen cloud computing data center. Qin did say that what he was sharing at Storage Visions represented what equated to a 2 year strategy for Facebook so exactly how far down the path that it is toward implementing it is unclear.

Having presented that vision for Facebook, the presentations at Storage Visions for the remainder of that day and the next were largely spent showing why this is the future at many large enterprise data centers but why it will take some time to come to fruition. For instance, there were some presentations on next generation interconnect protocols such as [PCI Express](#), [Infiniband](#), [iWarp](#) and [RoCE](#) (RDMA over Converged Ethernet).

This high performance, low latency protocols are needed in order to deliver the high levels of performance between these various pools of resources that enterprises will need. As resources get disaggregated, their ability can achieve the same levels of performance that can within servers or storage arrays diminishes since there is more distance and communication required between them. While performance benchmarks of 700 nanoseconds are already being achieved using some of these protocols, these are in dedicated, point-to-point environments and not in switched fabric networks.

Further, there was very little discussion as to what type of cloud operating system would overlay all of these components so as to make the creation and ongoing management of these application-specific virtual servers across these pools of resources possible. Even assuming such an OS did exist, tools that manage its performance and underlying components would still need to be developed and tested before such an OS could realistically be deployed in most production environments.

Facebook's Qin provided a compelling early look into what the next generation of cloud computing may look like in enterprise data centers. However the rest of the sessions at Storage Visions also provided a glimpse into just how difficult the

task will be for Facebook to deliver on this ideal as many of the technologies needed are still in their infancy stages if they exist at all.

---

# Five Key Changes for CIOs to Adopt to Successfully Evolve their Roles Within Today's Enterprise

The role of the Chief Information Officer (CIO) is evolving as enterprises worldwide attempt to navigate their way through the fundamental changes required to keep pace with the explosion of Cloud Computing, Social Media, Big Data and Mobile Computing. Information Governance, Compliance, eDiscovery, Data Security and Business Intelligence are now more important than ever. If the CIO can't keep pace, the fate of the entire enterprise may be at stake.

According to the [FreeDictionary](#) by Farlex, the definition of [evolve](#) is to develop or achieve gradually; to work (something) out or devise; and, to undergo gradual change. In the world of the CIO, the requirement to evolve is accelerating at an increasing rate and therefore CIOs have to *evolve at a much quicker pace* than is expressed in this standard definition of evolution.

According to the [annual study of the digital universe](#) by [IDC](#) and sponsored by [EMC Corporation](#), the amount of digital information created and replicated in 2010 surpassed 1.8 zettabytes (1.8 trillion gigabytes) – growing by a factor of 9 in just five years. The ICD study goes on to say

that while 75% of the information in the digital universe is generated by individuals through email, social media and texting, the enterprise will have some compliance and/or legal liability or business requirement to analyze 80% of that digital information at some point in its digital life. This requirement, within most enterprises, is the responsibility of the CIO.

However, given the historical perception of the CIO within many enterprises as being the *technology geek* with the legacy baggage of impeding the management and analysis of the data required by the business stakeholders to successfully complete their jobs, many CIOs are faced with an internal uphill battle to prove their value.

The best strategy for CIOs that want to *successfully evolve their role* is to redesign and restructure their organization(s) from *centralized technology and information gatekeepers* to an *enterprise-wide information service bureau*. [Eric Lundquist](#), VP & Editorial Analyst for InformationWeek Business Technology Network, stated in an article on August 14, 2012, titled, "[CIO: The Four Headed Monster?](#)" that this strategy is counter to past history in business (and politics) *where you increased your influence by increasing your power*.

Mr. Lundquist went on to say that a bigger budget, more direct reports, and participation in board room level committees was the mark of success. However, the consumerization of business technology and the growing phenomena of the mobile workforce with [Bring Your Own Device](#) (BYOD) that are making business-capable applications available to all employees via the Web is profoundly changing the enterprise information technology playing field. Mr. Lundquist contends that CIOs need to adapt to that change or risk being seen as a costly department without a clear mandate.

Five key changes that CIOs can adopt to successfully evolve

their roles and the role of the IT department within the enterprise are:

**1. Evolve from Providing Technology to Enabling Information and Collaboration.** Historically, enterprise CIOs have built and managed the complex hardware, software and networking infrastructure required to support the back office needs of the enterprise such as Enterprise Resource Planning (ERP) and Customer Relationship Management (CRM) systems. Access was structured, cumbersome and provided very little opportunity for users to collaborate.

Today's CIO can no longer hide in their office behind the firewall and dictate what IT solutions will be available. The successful CIO is going to evolve the role of the IT organization and embrace new technologies that *enable users to access and analyze business critical information and collaborate* without having to adhere to the rigid and centralized IT structures of the past.

A simple example would be to embrace a Software-as-a-Service (SaaS) based CRM system such as [salesforce.com](https://www.salesforce.com) that enables users to collaborate both internally and externally with clients and prospects outside the firewall and without the assistance (or gate keeping) of the IT department.

**2. Move to the Cloud.** With all of the vendor propaganda and cases studies available on the Internet that prove the cost savings and increases in productivity, you would think that every enterprise in the world had moved to cloud computing. However, according to an [IDC study](#), by 2015, only 24% of all new business software purchases will be of service-enabled software with SaaS delivery being 13.1% of worldwide software spending. IDC further predicts that only 14.4% of applications spending will be SaaS-based in the same time period.

These statistics don't mean that the cost savings and increases in productivity afforded by moving to the cloud

aren't real; it just means that many CIOs are reticent to move to the cloud and as a result are impeding their enterprises from realizing the benefits. Therefore, the successful CIO is going to move to the cloud in a responsible manner as quickly as possible.

As an example, CIOs might consider enabling their users to store and share data via the cloud. As I stated in "[Enterprise Alternative to Dropbox](#)," there are enterprise class storage and user collaboration solutions that support public, private and hybrid cloud deployment options for a variety of user devices such as legacy desktops, laptops, iPhones, iPads, Androids, and BlackBerrys.

**3. Embrace BYOD.** The very thought of *allowing* employees to Bring Their Own Devices to work and access and share corporate data is counter intuitive to the security responsibilities of the CIO. However, *the BYOD train has left the station* and therefore the successful CIO is going to embrace BYOD and evolve their enterprise security systems to meet the growing wave of BYOD.

The Federal CIO Council working group has released [a new guidance document](#) intended for federal agencies that are implementing bring-your-own-device (BYOD) programs. The document contains case studies from Delaware, and a handful of federal agencies. The council also presents a list of key considerations to keep in mind – such as a cost/benefit analysis for BYOD, security and policy obstacles, and roles and responsibilities.

**4. Adapt Alternative Development and Delivery Methodologies.** Historically, enterprise users learned through experience and disappointment that new software solutions development took the IT department years to design and develop and cost millions of dollars. In today's quickly moving global information world, the CIO and the IT department no longer have the luxury of such long and expensive deployment cycles.

Therefore, the successful CIO will adopt next generation [agile development](#) and deployment methodologies and utilize pre-build solutions that enable user requirements to be met in a matter of days or months as opposed to years.

As an example, the successful CIO could utilize one of the new enterprise application market places. As I indicated in "[Enterprise Application Market Places in the Cloud Provide Effective Alternative to Legacy Off-the-Shelf Software](#)," I would recommend any global 2000 IT department that is considering the development of next generation applications to meet the evolving demands of cloud computing, mobile computing devices, Social Media, Big Data, enterprise collaboration to investigate the new generation of application development platforms that support Enterprise Application Marketplaces in the Cloud. It could turn your IT Department into heroes.

**5. Be Proactive, Not Reactive.** Whether deserved or not, the enterprise CIO and their IT departments don't have a good reputation among enterprise users for meeting their requirements. Successful CIOs are going to meet this challenge head on and become *proactive* in identifying, developing and supporting the next generation solutions to meet the needs of their enterprise users. A good place to start is to get out of your office and actually talk to your users about their needs—you may be surprised by what you learn.

As enterprises worldwide attempt to navigate their way through the fundamental changes required to keep pace with the explosion of Cloud Computing, Social Media, Big Data and Mobile Computing, expectations for a successful CIO in 2012 have also changed. If the CIO can't keep pace with this new generation of requirements, the fate of the entire enterprise may be at stake. CIOs that embrace these five key changes will keep the enterprises they serve—as well as their own careers—on track to a successful future.

---

# Nirvanix Takes Another Step Toward Becoming the De Facto Standard in Enterprise Cloud Storage

*"Nirvanix was about a year ahead of everyone else in terms of what it could offer for enterprise cloud storage services."* Making this claim is Fred Rodi, the CEO of DRFortress, who over the last year had to look ahead to determine which storage provider could best position DRFortress and its customers for the future of cloud storage. So when it came time for DRFortress to make the choice, Nirvanix was the hands down winner.

Over the last few years Rodi has carefully watched the development of the cloud storage market for a variety of reasons. Chief among them, [DRFortress](#) now provides cloud services in the form of cloud computing, continuous data access and co-location among other services for its clients.

So Rodi anticipated that cloud storage was going to be a logical extension to DRFortress' existing services. The key was to deliver it in a way that DRFortress resellers could sell it in such a way that their customers could easily implement it.

This was easier said than done. Other cloud storage offerings from traditional storage providers required that DRFortress:

- Drop their storage boxes in at the DRFortress data center and another location
- Use the same vendor's storage at each site

- Rent space at *someone else's data center* to put this new storage there
- Configure the software necessary to virtualize all of the storage
- Dedicate his staff to managing the storage
- Use multiple portals to manage the storage
- Configure the replication software
- Open new WAN circuits anytime new replication options were configured
- Put a server in the cloud to have access to on-demand storage
- Buy more equipment and rent more space if additional data center locations were needed

Going to emerging cloud storage providers really was not an option either. He considered reselling Amazon cloud storage services but he almost immediately ruled that out. Being located in Hawaii, DRFortress is a long way away from the US mainland where the nearest Amazon data center is located.

Rodi already had nagging, unresolved questions regarding Amazon's availability, performance and security issues on the storage side. However the unpredictable latency and bandwidth costs associated with putting his customer's data into the Amazon storage cloud that was located on the US mainland and then pulling it back out in a timely manner made it an unacceptable option for him.

He also looked at leveraging other cloud providers like Rackspace, Terremark and Savvis that claim to have a cloud storage offering. However upon closer inspection he found that they did not offer it in a manner in which his customers wanted to consume it and they were not willing to deploy a cloud storage node in Hawaii and manage it as a service.

For instance, if a customer went to one of them and only needed storage (say 300 TBs) but had no corresponding compute requirement, they did not offer a solution that enabled them

to decouple their cloud computing solution from their existing cloud storage solution.

Further, there were two other intangible requirements that Rodi knew DRFortress had to deliver on.

- ***The first was to meet an expectation of “Excellence” from its customers.*** DRFortress’ existing customer base includes health care providers, the military, and telecommunication providers just to name a few. It is not like DRFortress is a startup that could call any solution “cloud storage” as they will find out it was not what the vendor originally claimed it was. Its customers look to DRFortress to be their trusted partner so for DRFortress to introduce the wrong cloud storage solution could adversely affect the rest of its business and the trust its partners had placed in it.
- ***The second was not to compete with his customers.*** DRFortress wanted to provide cloud storage services to its customers – not replace his customer’s existing storage solution with one from DRFortress. Yet almost every other cloud storage solution he looked at that met their enterprise requirements essentially put DRFortress in this position.

It is for all of these reasons and more that led Rodi to conclude that [Nirvanix](#) was the only viable enterprise cloud storage offering for DRFortress. While Nirvanix addressed all of his aforementioned concerns and exceeded his expectations in multiple different ways, there were three particular features Rodi cited as the primary reasons DRFortress elected to deploy Nirvanix.

- ***Geographically-diverse NameSpace.*** All of the global Nirvanix data centers appear as one common file system to all end users. Once they subscribe to Nirvanix, they

can configure their data to reside in as few or as many data centers as their application or business needs dictate. Rodi says, *“The ease in which Nirvanix enables us to set this up and then copies data to other locations is seamless for both us and our customers. This was huge for us.”*

- ***Eight Global Data Centers (with DRFortress becoming the Ninth.)*** Reselling Nirvanix cloud storage gave DRFortress and every one of its customers the freedom to put data on any one of eight (8) different data centers. Further, the specifications of the DRFortress data center qualified it to become the ninth data center in the global network of Nirvanix data centers. Rodi adds, *“Nirvanix will manage all of these data centers. It maintains the gear. It already has the WAN links to the other data centers. It is all just included in the service.”*
- ***Object based storage.*** The ability to decouple cloud computing and cloud storage offerings was critical for DRFortress to meet the demands of its clients. Using Nirvanix, neither DRFortress nor its clients needed to have a server in the cloud in order to provide storage on demand. Rodi quips, *“Using Nirvanix our customers now can have storage when they need, where they need it, anytime they need it.”*

Many storage providers have been vocal about the need for creating hybrid cloud over the past year. But to date they have not explained clearly yet how they plan to federate public cloud with their “private cloud”-labeled storage systems.

As Nirvanix owns and operates its own public cloud, it can easily deploy a node at a customer site and manage it as part of their overall network environment so existing corporate data storage solutions simply become part of what Nirvanix

already manages on a day-to-day basis. This is a major differentiator that enables Nirvanix to keep posting wins like this.

Yesterday's [announcement](#) that DRFortress, the largest service provider in Hawaii, has become a Nirvanix cloud storage reseller should come as no surprise to anyone familiar with Nirvanix or who regularly follows DCIG. DRFortress over the last year did an in-depth study of every cloud storage service offering available on the market that covered the gamut from well-known storage providers to emerging start-ups.

Yet in the end it came to the same conclusion that [Cerner](#), [IBM](#) and [USC](#) recently reached. If you are going to offer cloud storage as a service offering privately, publicly or both, you have to use Nirvanix.

But what organizations should not overlook is that these ***decisions*** by some of the world's largest service providers ***to choose Nirvanix are occurring with such regularity and on such a grand scale***

that it signals that Nirvanix is becoming more than just the best cloud storage solution on the market. It is ***on the cusp of becoming the de facto standard*** in how enterprise cloud storage is implemented and delivered.

---

## USC Upshifts to the Cloud

Right now many organizations are debating about who to select as their preferred cloud storage provider. But for organizations like USC that already manage petabytes of unstructured data, the decision is not about which provider to choose. Rather it is about deciding on the right technology that can transform it into both a private cloud storage user

and a public cloud storage provider.

This mindset helps to explain why USC announced earlier this week that it will deploy over 8 petabytes of unstructured data on a Nirvanix Private Cloud Storage solution so USC may act as a private cloud provider for its internal users even as it lays the foundation to offer public cloud storage services—as part of the [USC Digital Repository](#) banner—to outside companies looking for deep content archival and preservation. To accomplish this USC worked closely with Integrated Media Technologies ([IMT](#)) and [Nirvanix](#) to architect a private cloud storage solution that will enable it to manage internal data growth and external business opportunities.

To say that enterprise organizations that have hundreds of terabytes or even petabytes of data under management are reticent to move all of their data to a public cloud provider is an understatement. Aside from the requirements to first determine the viability of the provider and the stability of its solution, these enterprises have more practical concerns such as:

- Where and on what storage platforms will the provider physically store the data?
- Do they understand the nature of the data being stored with them?
- How well trained are the provider's IT staff?
- What are the provider's liabilities?
- What is my organization going to do with my staff and data center floor space if this is outsourced?
- Will the solution have resources on demand?
- Will the solution flexibly expand or shrink the storage pool as needed?
- Can the solution deliver expected service levels on a consistent basis?
- Do they only get charged for the storage they actually consume?

These and many other questions give enterprise organizations pause as they look at how to take advantage of public cloud storage services. It is as they do so that they are coming to two conclusions.

- Put in place a private cloud for their own needs
- Transform themselves into a public cloud provider to solve the needs of others in their particular vertical industry where they have the right mix of relationships and expertise

***The key now is for them to pick the “right” solution that delivers on these two requirements.*** So while the verdict may not yet officially be in as to what the “right” solution is, more and more of the largest enterprises on the planet are turning to the Nirvanix Private Cloud Storage solution to deliver on these cloud storage requirements.

Further evidence of that was one full display earlier this week when Nirvanix [announced](#) that USC ***will put over 8 PBs of unstructured data into a Nirvanix Private Cloud Storage solution.*** While Nirvanix has announced just in the last few weeks that IBM Global Services and Cerner are also putting in place a Nirvanix Private Cloud Storage solution that they can offer to their clients as their own public cloud services, the USC announcement is even more significant for the following reasons.

***First, it strongly suggests that current cloud storage solutions are not meeting enterprise cloud storage needs for data that will be backed up, archived and used for content collaboration.*** USC has very specific needs that are global in nature, especially its own internal USC Digital Repository that is serviced by USC. The USC Digital Repository may gather data from multiple contributors around the globe so it needs a cost-effective means to confidently aggregate, store and then share this data with the appropriate individuals within USC.

There are other storage offerings that are positioned as “private cloud storage” but some of these are basically existing storage systems that have been rebranded with a cloud label. What Nirvanix is doing is taking the same software, services and infrastructure that make up its global grid of nodes known collectively as the Cloud Storage Network and placing those exact nodes in customers’ data centers, in essence building them a miniaturized version of its public cloud.

A Nirvanix Private Cloud requires a minimum of two data centers for two reasons. First, one data center acts as the primary data center and the other as the secondary location. Second, because a single data center does not constitute a “storage cloud” because in Nirvanix’s view there is no failover and a cloud needs to have redundancy.

Using the Nirvanix Private Cloud, USC does not have to share resources with anybody that it does not want to. Instead USC has its own global namespace, its own multi-tenant file structure that it controls, and its own object store that can handle billions of objects. So while USC likes the elastic flexibility of a public cloud, USC also wants the security and peace of mind of using its own data centers that Nirvanix affords.

In this respect, DCIG is not aware of any other public storage cloud provider that is taking the same architecture that it offers as its public storage cloud network and also deploying it in customer sites and then using it to manage that storage infrastructure as a service with usage-based pricing.

***Second, it confirms the trend that internally IT is getting serious about converting their data centers from cost centers to profit centers.*** USC can now resell access to its Private Cloud and treat it essentially as a Public Cloud under the USC Digital Repository brand. By deploying the Nirvanix Private Cloud, USC has created a cloud-based archival service based on

a virtual construct that can scale to meet its needs.

So just as Cerner intends to sell cloud storage to hospitals and doctors, USC can similarly sell cloud storage archival services to the entire entertainment industry in Los Angeles and other media-centric cities. With NBC Universal already storing in excess of [2 petabytes](#) of digitized movies and TV shows in the Nirvanix Cloud Storage Network and adding 100 – 150 TBs monthly, USC is in a similar position as it has multiple post production houses and studios to call on that have similar digital content storage requirements.

***Third, USC is going full throttle to the cloud—not experimenting, not testing the waters, but loading up by moving over 8 unstructured petabytes of data to the cloud immediately.***

USC's decision to immediately put 8 PBs of data under managements indicates that it is immediately shifting its cloud strategy into high gear as it transforms its IT strategy and direction. In so doing, it is also making a clear message to other universities that the time is now to deploy, leverage and monetize the cloud.

Many companies and universities in the IT space are still struggling to articulate a cloud storage strategy and then put a solution in place. But in the case of USC, the debate is over. USC grasped its storage challenges, identified a cloud storage solution that aligned with its needs and now is aggressively moving to deploy it. In so doing, it sets a new standard for enterprise cloud storage deployments by which others will be measured going forward.

---

# IBM says “Don’t Cloud Wash!” Then Proceeds to Cloud Wash; SNIA Commits to New Testing Standard

Last week the DCIG team attended the [Fall 2011 Storage Networking World \(SNW\)](#) show in Orlando, FL. While there were a lot of cool storage companies, only two meetings left any kind of impression on me: one with IBM and another with SNIA.

My impressions are driven based on my focus in the electronic storage industry. As an analyst I spend more time looking at Information Management, Governance, Risk and Compliance from the perspective of product adoption and business intelligence, e.g. predictive analytics. So, eating lunch and dinner with our readers is always a great experience.

While eating lunch with Edil Vicenty, Director of Enterprise Architecture at the Central Florida YMCA, he commented “*The educational sessions are packed with business justification for the technologies discussed.*”

He identified a Monday morning session, SNIA Tutorial: [Data Center Evolution and Network Convergence.](#), by Joseph White, Distinguished Engineer at Juniper as a good one. While lunch with Edil was great, IBM left me understanding, “*There is no such thing as a free lunch.*”

Dan Galvan, VP of Storage Systems Marketing and Strategy, IBM Systems and Technology Group, had little to say about the cloud announcement with Nirvanix. When I asked if he was going to comment on it, he answered flatly “*No.*” ( DCIG,

however, did take the opportunity to comment on it in a blog entry last Friday: [Cerner and IBM Send Industry Message that Nirvanix is How Enterprise Cloud Storage Will Be Done](#))

Galvan did share that a lot of vendors had been cloud washing, but was reluctant to state names. Galvan's message that vendors are cloud washing was well received but in retrospect, Dan was preparing us for his message. After he talked about cloud washing, he then provided us with a new marketing story for IBM's [Storwize V7000](#) and [SONAS](#) products. Now, according to IBM, these are "cloud storage solutions."

In a nod to cloud washing, IBM's Galvan first tried to tell a cloud story, but then shared product updates. The three for the Storwize V7000 included:

- Clustering to scale performance
- File storage based on SONAS technology
- IBM Active Cloud Engine

Of the three product updates, the most interesting is IBM's [Active Cloud Engine](#). Active Cloud Engine is the policy system that IBM supports across its SONAS and Storwize product lines. This policy system currently supports moving files based on last access though IBM left me wondering when Active Cloud, Storwize, SONAS and Nirvanix were going to work together.

***Tragically, the Storwize V7000 will be using the SONAS file protocol engine.*** This engine is a distributed file system architecture based on GPFS that requires a distributed lock manager (DLM).

Distributed lock managers only work well in a LAN or MAN networking environment and using a DLM in a WAN environment means the WAN must be low latency and highly available, e.g. MAN. Thus, the Storwize V7000 and SONAS are not well suited for distributed cloud storage environments due to the limitations of DLMs.

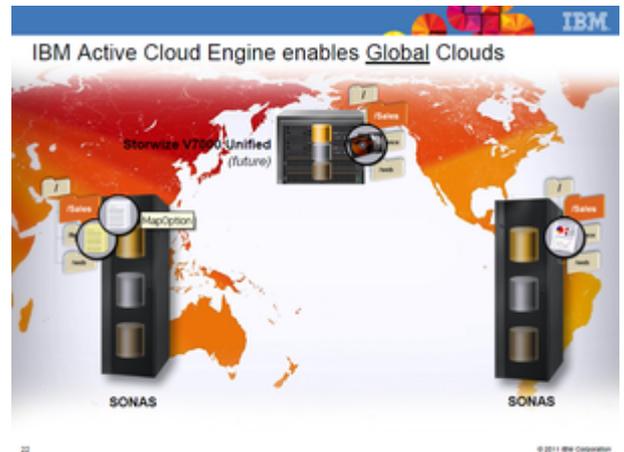
Where local or metro oriented data centers are concerned the Storwize V7000 combined with SONAS can be a good solution. Further, I don't disagree that Storwize combined with SONAS is a replacement for expensive file storage from the top tier NAS companies, like EMC and NetApp.

When asked if announcing Storwize integrating with SONAS was going to be competitive to IBM reselling the N-Series, Galvan responded, *"No, IBM's messaging discipline will ensure it won't be."*

This messaging is wishy washy. In a local or metro data center IBM won't comment on the co-opetition alignment with NetApp. Further, at the cloud level, IBM skimmed over SONAS remedial support for distributed storage. In addition, there was little talk about its Active Cloud Engine, its new agreement with Nirvanix and whether or not IBM will release RESTful APIs, e.g. Content Management Interoperability Services (CMIS). In this respect, IBM, Galvan and his team need to stitch together a better story.

"Moving the Clouds" was the sunny disposition and clarity delivered by Wayne Adams, Senior Technologist in the Office of the CTO at EMC, on the Emerald Program. The Emerald Program is a SNIA-funded-cross-vendor program sharing information on storage system "power usage and efficiency."

One of the crucial points Adams shared was regarding the SNIA Emerald taxonomy. Getting an industry to agree on taxonomy related to anything about performance is near impossible. The challenges posed for "energy efficiency and use" are no different than those posed by product and performance testing with [DCIG previously covering the difficulty in product testing in a prior blog entry](#).





The taxonomy is a major milestone because companies who compete head to head for billions of dollars don't like to agree on much of anything. Two key terms from the taxonomy are Idle and Active metrics. To identify those terms, SNIA Emerald Program team evaluated models used by the EPA, Energy Star, etc. For example, a consumer equivalent is the EPA's City versus Highway MPG rating for automobiles.

As is the case with Highway MPG, an Active State system has certain metrics that must be met. For example, some systems perform routine housekeeping. To properly understand energy use in an Idle State, the period of measurement must exceed the time required for the system to enter housekeeping.

After hearing about this, I asked Adams, "*As analysts we have a baseline for measurement and valuation?*" Not so fast, responded Adams. The Emerald Program is only **ONE** of several data points to consider when purchasing storage infrastructure.

At DCIG we understand these challenges and recommend standing up a [DCIG Buyer's Guide](#) next to the key points from the Emerald Program. Further, as the [Emerald Program terms and conditions allow](#), vendors and end users may see some representation of the Emerald Program results in our Buyer's Guides.

### **Credible, Cloudy and Consistent**

It is clear to me that SNW and SNIA event management are taking a closer look at the presenters and their material. This is a good sign that **future shows will deliver** credible technology details and business roadmap for adoption. IBM **must clear up their story** around the Storwize V7000, Active Cloud Engine, Nirvanix and SONAS. Lastly, the Emerald Program created consistency on energy use testing and terms. We can

**ALL benefit** from its adoption. Where DCIG Buyer's Guides are concerned, DCIG is listening.

---

## **Cerner and IBM Send Industry Message that Nirvanix is How Enterprise Cloud Storage Will Be Done**

Anyone who still doubts that Nirvanix is [poised](#) to deliver the same type of solution for cloud storage that VMware already delivers for cloud computing got a serious wake-up call this past week. Announcements that both Cerner and IBM entered into strategic relationships with Nirvanix are more than just validations of Nirvanix's cloud storage technology. They signal that Nirvanix is poised to become how enterprises of all sizes will eventually implement cloud storage.

There are seminal points in the adoption of any game-changing technology. This happened in late 2003 when [EMC acquired VMware](#) and then proceeded to introduce it into its enterprise accounts. In so doing EMC almost single-handedly transformed VMware vSphere from a compelling technology into one that enterprises felt they could begin to safely run their business and mission critical applications on.

Something similar occurred this past week in the maturation of cloud storage space. While no acquisitions have yet occurred, what may be a precursor to such an acquisition occurring is IBM's [announcement](#) that it has formally entered into a strategic five year OEM relationship with [Nirvanix](#). This should give the entire industry – vendors, VARs and enterprise

organizations – pause for a number of reasons.

First, embedded in the first paragraph of that announcement is IBM's intent **to integrate cloud storage technology from Nirvanix as part of an expanded [IBM SmartCloud Enterprise storage services portfolio](#)**.

The term **integration** suggests that this is much more than just IBM offering Nirvanix to pass the time until a better cloud storage comes along. It **implies** that this is **THE** cloud storage solution that IBM wants and that now is the time for IBM to bet its global cloud storage strategy on Nirvanix's cloud storage solution.

Making this announcement more interesting is that Dan Galvan, IBM's VP Storage Systems Marketing and Strategy and who is responsible for marketing IBM [SONAS](#) (an IBM scale-out storage solution that competes with Nirvanix at the private cloud level) **refused** to comment on this announcement when asked about it at the Fall Storage Networking World (SNW) 2011 conference. **This refusal suggests that at higher levels within IBM that IBM is placing a premium on being a cloud storage provider even over selling its own cloud storage solutions.**

Second, these announcements on consecutive days from [Cerner](#) and [IBM](#) about their decision to offer Nirvanix cloud storage instead of either their own solution or solutions from other vendors indicate that:

1. A change is occurring in how the clients of Cerner and IBM are asking them to deliver storage
2. Nirvanix offers cloud storage that meets the demands of how the clients of Cerner and IBM want it delivered: by the GB

Their clients only want to pay for the amount of storage they use, not pay for a bunch of storage capacity up front in anticipation of eventually using it. Even IBM's Galvan at the Fall SNW 2011 went so far as to say that IBM's clients are

“sick of cloud washing.” They do not want today’s existing storage solutions repackaged under the guise of “cloud storage.” They instead want real cloud storage solutions that are “storage capacity in the cloud” not “storage capacity in a box.”

So is it any real surprise that Cerner and IBM have to look beyond current storage solutions to find a cloud storage solution that better aligns with how their customers are defining “cloud storage?” Cerner and IBM are both smart, profit-driven companies and they know their customers are no longer buffaloed by cloud storage terminology. They want a cloud storage solution that:

- Automatically distributes data across multiple geographic sites and regions so it is protected
- Dynamically scales up (or down)
- Is configurable as either a private or public cloud (or both)
- Metered billing
- Multi-tenancy to securely store data from multiple clients on the same physical platform
- No up-front capital investment to get started
- Offers “Pay by the drink” (only pay for what you use)
- Provides enterprise level support
- Puts data on the right tier of storage according to its usage
- Scales to handle **petabytes** of storage capacity

So when Cerner and IBM look at what their clients want and what storage providers have to offer, Nirvanix is the solution that best aligns with this diverse set of customer requirements.

But possibly the biggest implication coming out of Nirvanix’s [announcements](#) with Cerner and IBM is that cloud storage as a whole overcomes a larger perception problem: *that it is not enterprise ready*. By both Cerner and IBM selecting Nirvanix as

the cloud storage solution that they will deliver to their enterprise customer base, cloud storage as a whole gets the validation it needed from two enterprise level service providers with Nirvanix being the sole beneficiary of this validation.

As DCIG knows from the regular conversations it has with end-users, they do want great technology. But they also want a great company with an established reputation to deliver it so if things go south (*and things inevitably go south at some point*) that there is someone reputable to call upon anytime day or night to fix it and make things right. So while Nirvanix's cloud storage technology is sound, Nirvanix does not have the same reputation as either Cerner does in the health care industry or IBM does globally.

So by Cerner and IBM both offering Nirvanix's cloud storage technology **and** putting their name behind it, their customers have more confidence to move ahead with cloud storage in general and Nirvanix specifically without lingering questions about how it will be supported. They **know** it will be supported in the same manner as other products that they currently use from these two companies.

In much the same way the future of enterprise cloud computing got a boost when EMC acquired VMware, the future of what enterprise cloud storage will look like got a big boost this week with both Cerner and IBM announcing that they will deliver Nirvanix as their enterprise storage cloud offering. In so doing, enterprise organizations get more than a message that cloud storage is now a viable option for them. They just got a strong message from Cerner and IBM that Nirvanix is the means by which they can confidently implement it.

---

# What Your Cloud Will Look Like: Final Thoughts from VMworld 2011

I realize VMworld 2011 ended over a week ago and everyone is by now probably looking ahead to the next big thing. But before we leave VMworld 2011 behind in the annals of history, I wanted to take one final look at how VMware went about promoting cloud ownership. Because rather than telling users they should own “VMware’s cloud” or “NetApp’s Cloud” or “EMC’s Cloud” or even some cloud service provider’s cloud, it touted “Own Your Cloud.”



This is a different spin than what I expected going into VMworld. Having been an end-user for most of my normal life prior to delving into the analyst realm, the only thing I found most vendors interested in doing is making my environment look like an extension of their environment and me feeling good about it as I do so.

VMware’s CEO Maritz almost admitted as much during his keynote that this is what was going to happen if applications become too addicted to VMware. He stated, *“The danger that lurks is as vendors tune their applications for VMware, they become locked in and VMware thereby creates a new mainframe environment.”*

As I have mentioned before, lock-in is not necessarily a bad thing. In fact, it would probably in some respects make the

whole world of computing much easier since everyone would have a common platform from which to develop and run applications – aka Microsoft.

So in this regards VMware is no different than Microsoft. VMware is clearly interested in having as many of the world's applications as possible use vSphere as they are virtualized and move into the cloud. This was evident by Maritz citing an IDC statistic that over 50% of the world's applications are already virtualized. Further, VMware is introducing more initiatives to expand that percentage.

But VMware also needs to play it smart. It knows it is in a market leadership position and rapidly becoming, if not already is, the dominant player in server virtualization. So the danger that VMware runs is unintentionally creating a counter culture that hates VMware just because VMware is the dominant player in server virtualization.

This is arguably what happened with IBM and its mainframe platform which led to the advent of open systems computing (UNIX, Windows, etc.) It then happened again to Microsoft and its Windows platform which arguably contributed first to the rise of Linux and then eventually to VMware's emergence.

So the problem that VMware faces is this. How does it become the next dominant platform on which businesses lock in on while not creating this counter culture that eventually gives rise to a company that is dedicated to replacing VMware?

In this regards, VMware having Maritz at its helm is serving it in good stead as he is playing it both smart and shrewd. Rather than getting arrogant over VMware's recent and ongoing success, it appears he has learned from the lessons he gleaned while working at Microsoft. Rather than allowing a counter culture to emerge which VMware has no part of or insight into, VMware is embracing and facilitating it.

Perhaps the best evidence of this is its [Cloud Foundry](#)

initiative. While it is still in beta, Cloud Foundry supports multiple frameworks, multiple cloud providers and multiple application services all on a cloud scale platform. But maybe most importantly from VMware's perspective, it is a service project initiated by VMware so VMware is aware of innovations occurring outside of VMware and can take advantage of them so it ultimately benefits VMware in the long run.

VMware may be touting "Own Your Cloud" but at the end of the day the cloud it really wants end-users, businesses and organizations to own is VMware's cloud. The trick for VMware is trying to make "Your Cloud" look remarkably like VMware's cloud without having them resent the choice either now or in the future which gives rise to a counter culture.

This helps to explain why VMware is putting initiatives like Cloud Foundry in place. It gives the illusion that everyone has the freedom to choose whatever cloud they want and then take it in any direction they wish. So when they choose VMware's cloud without harboring any resentment about it either now or in the future, they get the cloud they want and want to own while VMware gets what it really wants – end-users who feel like VMware's cloud is their own.

---

## **Defining a Cloud as "Good" or "Bad" May Come Down to Whether or Not It Works**

This past Thursday I became aware of David Linthicum's [Cloud Computing blog](#) over at InfoWorld for the first time as a result of an email that was promoting a [blog entry](#) he wrote earlier this week. In that particular blog entry he warns why

a shortage of cloud architects will soon lead to “bad clouds.” That’s interesting because I did not realize that the industry had really settled on what defines a “good” or a “bad” cloud.

As an individual who first majored and received an undergraduate degree in Theology before getting a second undergraduate degree in Computer Information Systems, the concepts of “good” and “bad” have always intrigued me.

You would think that having a degree in Theology would result in me believing that the lines between “good” and “bad” are pretty much black and white. Yet what I concluded from my studies is that while the Bible clearly defines some behaviors as “bad” or “evil” and others as “good,” there are a far greater number of behaviors (I would say an almost infinite number) that it deems as “acceptable.”

These lessons as to what constitutes “good,” “bad” and “acceptable” in the spiritual realm have had an interesting carryover into the computer realm. Once I got into the world of computer science, I found most of my colleagues define computer architectures in terms of “good” and “bad” (though they may use terms like “smart” and “stupid” to describe them.) For example, I find many UNIX folks consider Microsoft Windows a curse that mankind is forever doomed to suffer under while mainframe folks look at today’s distributed computer solutions as “tinker toys.”

Unfortunately it just is not that easy to look at a particular computer design and then label it “good” or “bad” (or “smart” or “stupid” as the case may be.) The same holds true when creating a cloud. After all, how can you create a cloud and holistically classify it as “good” or “bad” when each company’s definition as to what a cloud needs to do or provide may be different?

By way of example, here are a few attributes as to what I would define as a “good” cloud providing:

- Appropriate service levels for each application
- Data mobility
- Data protection
- Ease of management
- Flexibility to independently scale capacity – memory, networking, processing or storage
- Multi-tenancy
- Security
- Uninterrupted service

So assuming you agree with me there is a lot of room for interpretation as to how each of those features is delivered. Consider data protection. Data protection can be delivered in a multitude of ways – snapshots, traditional backup and replication just to name a few.

So are two of these forms of data protection “bad” and the other one “good”? Or do all of these forms of data protection fit within the spectrum of “acceptable”? I would argue the latter but it really depends on what the business wants to accomplish.

Even then, the business may not be (and likely is not) aware of all of the data protection options available in the market. As such, they are likely not going to have the “best” data protection solution as analysts like myself may define it and consider “good” implemented in their environment. More than likely they are going to have a solution that is “acceptable” to them.

Organizations should also not think that they can turn to vendors or solutions providers in anticipation of getting the “best” cloud available either. Vendors are typically limited to offering whatever solutions they have in their portfolio while solutions providers are motivated by a variety of incentives not the least of which is which of their providers is giving them the most incentive to sell their technology.

While I certainly do not think few if any solutions providers would purposely deliver a cloud that does not work (a “bad” cloud) if they can deliver an “acceptable” solution while making some extra money in the process, who is to fault them?

Defining any particular cloud design as “good” or “bad” is a risky proposition. If anything, the only clouds that are easy to define as “bad” are those that provide little or no business value or simply do not work. Conversely classifying a cloud “good” may be as simple as a cloud that works and meet the needs of both the business and the requirements of the IT staff that have to support it.

---

## **Being Anti-Cloud and Cautious Cloud are Two Very Different Mindsets**

Recently an individual brought to my attention that I had created a perception that some people thought I was “anti-cloud” and that I believe the cloud is “bad” or “evil.” I am not exactly sure how that perception got created or how that conclusion was reached. But whether or not that perception is accurate, he does raise two valid points. What is my opinion of “the cloud” and how do I think organizations should proceed with it?

To a certain degree I can see how people, based upon some of my prior blog entries, might conclude that I am “anti-cloud.” For instance, back in late April I wrote a [blog entry](#) entitled ***“Cloud! Cloud! Cloud! A Not so Authoritative Look at What Cloud Terms Mean.”*** In it I took a tongue-in-cheek look at how vague various cloud terms are and how they are frequently

being bent to serve almost any purpose.

I similarly wrote a [blog entry](#) a few weeks ago asking the question, "**Who owns the cloud?**" In that blog entry I examined a situation where a VAR was telling me about a company that was building out a private cloud. However this cloud is apparently being built without any rhyme or reason as the right hand in the company buying the storage does not know what the left hand is doing. This lack of strategic direction and corporate guidance as to what that company's cloud should eventually look like is resulting in a tremendous amount of overbuying and waste.

So if someone read only those two blog entries, I could see how one might conclude that I am anti-cloud.

In short, that is not correct. However I will certainly position myself as someone who sees himself as "Cautious Cloud," at least when it comes to businesses running any of their applications or putting any of their data in the cloud.

For instance, businesses should ideally verify:

- How financially stable their cloud provider is
- What their physical infrastructure looks like
- How their data center is architected
- How prepared they are to recover from an outage or a disaster
- How well trained their people are
- The processes and procedures that they internally follow to manage the cloud

What I have found is that many if not most cloud providers have facilities, processes and procedures that exceed the capabilities and expertise of what many businesses have in place. In these circumstances, businesses may be pleasantly surprised to find that by moving some or all of their applications and infrastructure to the cloud may be the best decision they ever made. They could conceivably lower IT

costs, improve application up time, discover new options for disaster recovery and ultimately make IT an enabler of business initiatives.

So when a cloud provider offers all of those services, I am a 100% believer in the cloud and I say, "**Go for it!**"

But there are just as many or more circumstances where an organization still has not internally defined what the cloud means to it, who is internally going to own the cloud or set its strategic direction. In those situations, I would encourage companies to proceed very slowly with any cloud initiative until they resolve these questions of authority and direction.

Notice I did not say they should not proceed at all. Sometimes the only way to learn what you want or need in order to move forward is to experiment and try something new. But it appears some organizations are jumping feet first into private or public clouds without any clear expectations as to how it will ultimately improve their situation.

So for those of you out there who thought or perceived that I was anti-cloud, that was not the message I intended to convey or impression I wished to leave you with. But as I voice my support for the cloud, I **DO** want to leave people with the impression that the cloud should be implemented carefully and cautiously.

Because simply implementing a cloud (public, private or hybrid) or hosting applications in it does not necessarily make all problems go away. Absent any clear direction or strategic objectives, it may only make them worse.

---

# VMware Starts to Build Its Case for Trusting the Cloud

VMware shared a pretty astounding statistic this past Tuesday when it rolled out vSphere 5. It stated that 50% of application workloads will be virtualized by the end of 2011 with that ratio continuing to grow at a rate of 10% per year for the next few years. That's pretty remarkable considering ten years ago when I proposed starting to virtualize my prior company's infrastructure that I was scoffed at by many of my peers.

There are any number of reasons why ten years ago people scoffed at my idea to start virtualizing my company's environment. But the reasons that I cited as to why my company needed to embark upon a path toward virtualization still hold true today. Further, what I saw at that time, and which is becoming more evident with every passing day, is that applications that are **NOT** virtualized will be the exception not the rule.

Yet it is ironic that when I went to VMware's website last night as I prepared for this blog entry, the slogan I saw on its website was:

## It's Time to Trust Your Cloud

You know what that implicitly tells me? People today do not yet trust the cloud even as they did not trust in its early form ten years ago. (*Anyone remember the phrase "storage utility"?*)

Frankly, ten years ago I did not trust the storage utility in its early form. At that time I could not answer some of the

thornier questions that were thrown at me by my counterparts as to how virtualization (server or storage) would behave in a mission critical environment.

For instance, I had to field questions such as:

- If a server has multiple paths (say 8) to a storage subsystem, how does that work when the storage system is virtualized?
- If virtualized storage is presented to separate AIX, Linux, UNIX, and Windows servers on the same physical port, does that always work? Always??
- Are all of the port flags on the storage array set correctly and is working in someone else's environment who I can call and verify that with?

I could not answer yes to those questions at the time because I did not know. Further most of the vendors providing these storage solutions did not know either. So the status quo prevailed.

This is what I strongly suspect VMware is running up against as it has to deal with similar questions as it moves up the application stack. Granted, VMware proudly proclaimed in its webinar this past Tuesday that its OS is being used to host more business critical applications such as Microsoft Exchange, Microsoft SharePoint, SQL Server, and Oracle. But as that occurs, I can just see skeptical, battle-hardened IT managers who have been burned one too many times looking their VMware reps square in the eye and asking:

- If I virtualize applications running on Linux and Windows OSes on a VMware server and each of these applications needs multi-pathing does that always work?
- Are you sure VMware makes all of the interoperability issues go away? Are you absolutely sure?
- If I have a virtualization application that has its data spread across 2 or 3 arrays from different vendors, does

that work?

Those sort of questions have to make those at VMware squirm just a bit. While VMware may want to say that with virtualization all of those issues are virtualized away, it is never that easy or simple when you get into enterprise environments when a single application outage may result in millions of dollars of lost revenue or garners headlines on the evening news.

VMware is admittedly probably already doing this in many environments. But this slogan tell me it recognizes it has an uphill battle to convince corporate IT that it is ready to assume the rest of these workloads. VMware is probably coming to find out if it does not know already that once you get into some of their environments, it is rarely Plug-N-Play but more like Plug-n-Pray.

All of the testing in the world cannot prepare you for the idiosyncrasies that each of these environments is bound to have and that VMware is going to encounter first hand and for the first time. To VMware's credit it admittedly already has a lot of experience to draw upon as it prepares to fight these battles.

But every enterprise has its skeleton applications in the closet that, as VMware runs across and virtualizes them, will likely cause even those internal to VMware to shake their heads and marvel how these enterprises have operated all of this time in the state they are in.

VMware painted an exciting picture of the future this week and I do firmly believe that VMware will evolve to become more than just a provider of virtualization software but a key enabler of the cloud. But VMware's biggest journey may yet lie in front of it.

It now needs to convince people that it is time to trust the cloud. Whether or not VMware succeeds in persuading

enterprises of that concept remains to be seen but if VMware does (and it is arguably in the best position to do so) then we will likely in the next few years see what my former colleagues thought was ludicrous less than a decade ago.

---

## **Storage Companies Cautiously Celebrate Their Independence**

Independence Day on July 4th in the United States is only a few days away but as it approaches storage companies are cautiously celebrating their independence. As they do they are either looking to survive or aggressively looking to be acquired to avoid becoming a footnote in the annals of history with Pillar Data Systems becoming the latest storage company to join the ranks of the acquired that now pledges its allegiance to a new master.

One only needs to look over the last few years to get some sense as to how many storage companies have been acquired:

- 3PAR (HP)
- Compellent (Dell)
- Data Domain (EMC)
- Diligent Technologies (IBM)
- EqualLogic (Dell)
- ExaNet Systems (Dell)
- Isilon (EMC)
- LeftHand Networks (HP)
- LSI Engenio (NetApp)
- Pillar (Oracle)
- Ocarina (Dell)
- Sun StorageTek (Oracle)

This is not to say that the storage technologies that these companies developed are gone. If anything, the technologies originally associated with them are doing just fine with the possible exception of Sun StorageTek. Further, there are plenty of independent storage companies still in existence in addition to EMC, HDS and NetApp which include:

- Aberdeen
- Celeros
- IceWEB
- Nexsan
- Nimbus Data Systems
- Nirvanix
- Quantum
- Xiotech

But to say that the majority of these companies are going to become the “Next Big Thing” is probably a stretch. While there may be one or two on that list that may achieve a level of success that parallels EMC, HDS or NetApp, the best the rest of them can hope for is to simply survive.

The outlook for the ones that “survive” is certainly not doom and gloom. There will always be companies out there looking for viable storage companies that offer alternative solutions that are both economical and reliable.

But this is only a formula for staying in business, not one for explosive growth or acquisition. One only has to look at Oracle’s [acquisition](#) of Pillar that it made for 100% earn-out with no upfront payment which, at the end of the day, Pillar was probably lucky to get.

The major technology players (Dell, EMC, HDS, HP, IBM, and NetApp) already possess SAN storage solutions that are comparable to Pillar. So while the Pillar Axiom consistently scored very well in DCIG Buyer’s Guides, the absence of available suitors and the low probability of Pillar

experiencing explosive growth at this late stage of the storage game for the type of technology that Pillar offered suggested that this was as good an offer as Pillar was going to get.

The upside for Pillar is that if Oracle properly manages the intellectual property that it gets with Pillar, it could be a boon for both Oracle and Pillar. It certainly helped the other companies (3PAR, Compellent, Data Domain, etc) listed above when their technology was acquired and then resold by their new masters..

If it does succeed, ***it will require a change of heart on Oracle's part*** as to date Oracle has had a poor track record of executing on its storage acquisitions. So it remains to be seen if this acquisition of Pillar will turn around Oracle's storage fortunes or end up being yet another failed venture in storage by Oracle.

The real question is where do all of these storage acquisitions leave the big independent storage providers: EMC, NetApp and maybe HDS? While they are publicly claiming the future looks brighter than ever as all of their competitors have been scooped up (either by themselves or "server companies" like Dell or HP) their confidence may be misguided.

The quality and breadth of their storage offerings are certainly on par and, in a number of cases, arguably above what Dell, HP, IBM and Oracle have to offer. But the gap between these traditional "application" and "server" companies is no longer as great as it once was. While I discussed this in greater detail in a [blog entry](#) I wrote a few weeks ago, it may force these companies to abandon their "all storage" philosophy and come to market with a more holistic offering.

Already we are beginning to see this happen at some level with alliances like "[VCE](#)" (VMware – Cisco – EMC) and NetApp re-selling its technology through Fujitsu and IBM. But can these

alliances stay in place long term or will companies like Cisco finally be forced to pull the trigger and buy EMC or IBM and/or Fujitsu have to buy NetApp not because they want to. Rather they will find themselves in the same positions that Dell and HP found themselves in. They have to acquire them in order to capture all of the margin they are currently giving up to EMC and NetApp respectively because they have to resell their products.

Currently both EMC and NetApp want huge premiums for their respective companies but they may have to back off as the months and years pass. As companies virtualize their environments, the underlying hardware (server, network and storage) may command less of a premium.

Instead companies will place a greater premium on integrated end-to-end virtualized infrastructures that enables them to move their applications wherever they want when they want. As this transition unfolds, this would seem to put companies like Dell, HP and to a lesser degree IBM in the driver's seat assuming they execute on their storage strategies.

So as July 4th approaches and the United States celebrates its independence, celebrations for independent companies in the storage industry are a bit more subdued. Virtualization is taking over in customer accounts of every size which puts a greater emphasis on the features that enterprise storage systems have to offer than ever before. But with more major vendors providing "end-to-end" solutions optimized for virtualized environments, the argument for selecting an independent storage provider has to be more articulate and justifiable than ever before.