



# 20 Things That May Be “Clouding” Your Choice About the Cloud But Shouldn’t

March 28, 2012

BY VISHWAS LELE



## VISHWAS LELE

Vishwas Lele is Chief Technology officer (.NET Technologies) at Applied Information Sciences, Inc., where he has worked for the last 14 years. In his current role, he is responsible for assisting organizations in envisioning, designing, and implementing enterprise solutions that are based on .NET technologies.

Vishwas also serves as the Microsoft Regional Director for the Washington D.C. area. As a Microsoft endorsed expert, he is regularly consulted for his insight and informed perspective on implementing .NET-based solutions.

A regular industry speaker and author, Lele has presented at a number of industry conferences as well as community user groups.



**TABLE OF CONTENTS**

- 1. Security .....4
- 2. Compliance .....4
- 3. Latency.....5
- 4. Availability .....5
- 5. Cost of Operation .....6
- 6. Loss of Control .....6
- 7. Tools .....7
- 8. Migration Costs.....8
- 9. Legacy Applications .....8
- 10. Tightly Coupled Application Portfolio.....9
- 11. Lock-in.....10
- 12. Private Cloud .....11
- 13. IT Staff.....11
- 14. Licensing .....12
- 15. Passing Fad/Immature State.....13
- 16. Outsourced .....14
- 17. What Goes Around Comes Around .....14
- 18. Budget .....15
- 19. Performance .....15
- 20. Not Built Here .....16
- Conclusion .....16



## 1. SECURITY

In the last four years that I have worked with cloud computing, I have heard from a wide cross-section of customers about the challenges they face in adopting the cloud. In this paper, I will discuss these concerns and provide my assessment of how these challenges are being incrementally alleviated. Not all challenges have been fully addressed, but a number of them have become less of a hurdle for cloud adoption.

*Note: My comments apply to IaaS and PaaS styles of cloud computing, which are increasingly being coalesced into one. In the interest of full disclosure, while I follow offerings from multiple cloud vendors closely, my most extensive hands-on experience is with the Windows Azure Platform.*

The obvious challenges—Security, Compliance, Availability and Latency—have been written about so much already that I will merely touch on them.

**Claim:** *“Multi-tenant/shared infrastructure will make our data and business logic vulnerable to attacks.”*

There are cloud platforms from different vendors that have been commercially available for multiple years with proven records of security success. While there have been well-published breaches, cloud computing vendors would be cavalier not to minimize risks to their multi-billion dollar investments. Many cloud vendors have invested in achieving industry standards, such as [ISO 27001](#), that are designed for datacenter operations while others have taken up [verification projects](#) to prove functional accuracy of their software building blocks on the platform level.

Today, cloud platforms are implemented with an array of controls to provide a secure environment for hosting applications. These controls cover **physical security**, including facilities access; redundant power; media disposal; **network security**, including VLAN isolation; packet filtering; **confidentiality** that includes identity and access control; **data isolation and integrity**; and **accountability**, including multiple levels of logging, monitoring and reporting.

In summary, customers will find, in most instances, the security properties exhibited by the cloud platforms equals or exceed those of on-premises-based infrastructure.

## 2. COMPLIANCE

**Claim:** *“The domain we operate in requires compliance with regulatory standards, such as FISMA, HIPAA, PCI, SAS 70, SOX, etc.”*

Of all the hurdles, the one related to compliance may be the hardest to overcome. Most of these standards were not designed with cloud computing in mind and include clauses that are not easy for cloud vendors to comply with. Frankly, it will take a collective industry and government effort to make the laws relevant and transparent for cloud computing.

Even when cloud vendors achieve these certifications, they are not transferrable to applications hosted in the cloud. For example, the Verizon cloud may be PCI DSS compliant, but the applications hosted in its environment are not automatically compliant. This situation is not different from having to certify on-premises-based

applications individually. For example, some applications have been independently HIPAA certified despite being built on platforms, like AWS, that do not have the HIPAA accreditation.

At the same time, it is encouraging to see cloud vendors achieve compliance standards, including [AWS adding FISMA Moderate authorization and accreditation](#) and [Office 365 becoming compliant with business associate HIPAA requirements](#).

In summary, workloads that have specific compliance requirements will be hard to migrate, but cloud providers are trying to meet various certifications and accreditation requirements.

### 3. LATENCY

**Claim:** *“Cloud will add latency to my applications that my end users will not tolerate.”*

Unless all of the parts of an application can be co-located in the cloud, it is inevitable that latency will be introduced to deal with the distance and bandwidth uncertainty over public networks.

Fortunately, cloud vendors are offering solutions to this challenge.

- Technologies such as data sync can be used to synchronize cloud and on-premises databases, thereby removing the need to make a remote call.
- Cloud vendors can also offer durable queue functionality, which makes it easy to decouple incoming requests and responses. This functionality means the caller is not impacted by the latency hit associated with a roundtrip to and from the remote resource.
- Cloud vendors are also creating marketplaces for selling and buying datasets within the cloud that are directly available to cloud-based applications.
- Finally, cloud vendors are making it easy to move large amounts of data in and out of the cloud using portable storage devices. In many cases it is possible to establish a dedicated connection between on-premise and the cloud platform as well, thereby increasing the network’s bandwidth and consistency.

In summary, latency concerns can be minimized by using techniques such as data synchronization and durables queues and by taking advantage of datasets already available in the cloud.

### 4. AVAILABILITY

**Claim:** *“Cloud outages can take my business-critical application down for hours.”*

Recent, well published, cloud outages have raised questions about the availability and SLA guarantees associated with cloud-based applications. While the cloud platforms are designed to be fault tolerant by replicating data and compute instances, ensuring high availability requires careful [business continuity planning](#). For instance, deploying to a single datacenter will meet a two-hour [recovery time objective](#) (RTO), allowing the application to be moved manually to another location if the datacenter becomes temporarily unavailable.

However, if a zero or close to zero RTO is required, then the cloud application needs to be deployed simultaneously to more than one datacenter for a higher price. Cloud platforms facilitate such deployment topology by providing tools such as the [Azure Traffic Manager](#). The Azure Traffic Manager makes it possible to distribute incoming traffic to different datacenters.

In addition to RTO, it is also important to consider recovery point objective (RPO). If the acceptable RPO is a maximum of four hours, then customers need to ensure that adequate backups are being performed to meet this goal. For example, a Database-as-a-Service offering such as [SQL Azure](#) is designed to be fault tolerant by utilizing multiple replicas, but in order to ensure a four-hour RPO, data needs to be backed up to another datacenter using a service like the [DAC Import /Export Hosted as a Service](#).

In summary, concerns related to availability can be mitigated by careful business continuity planning.

## 5. COST OF OPERATION

**Claim:** *“How can a perpetual operational expenditure be less than one-time capital expenditure?”*

The key to determining if cloud computing can save you money is to take a holistic view of the costs involved. It is easy to take the unit cost of a cloud-based instance, multiply it by the hours and compare it to the monthly lease for a server. However, such a comparison would be lopsided. In addition to the direct hardware and software costs, it is important to consider indirect costs, such as applying patches and hot fixes, power, cooling, networking infrastructure, security, labor costs and so on. Finally, the ancillary costs—including purchasing, procurement, contracts/RFPs and disposal of hardware (like disk wiping), need to be considered as well.

Cloud vendors also have the following inherent cost advantages that on-premise setups will be hard-pressed to match:

- A multi-tenant/shared infrastructure that affords them huge economy of scale.
- Hardware homogeneity within their datacenter creates increases efficiency.
- A continuous [cycle of innovation](#).

Even if lower costs can be achieved for on-premises setups, consider this, you pay for cloud resources usage with the ability to turn off the resources at any time; however, with dedicated hardware you pay for all of the reserved capacity whether or not you are using the resources.

In summary, it is important to take a holistic view of costs, including direct and indirect costs. Cloud providers are able to offer a computing at a low cost because of economy of scale and hardware homogeneity.

## 6. LOSS OF CONTROL

**Claim:** *“I cannot have someone walk right up to an errant server in case of an emergency.”*

There is no question that there is going to be some loss of control when moving your application to the cloud.

- Cloud platforms typically split the application into smaller [parts](#) to improve availability. This structure makes it challenging to form a unified runtime view of the application.

- Cloud platforms split or partition data for disaster recovery and scale-out reasons, raising questions about data sovereignty.
- It is a challenge to track events, such machine shutdown and restarts in a consolidated manner.
- There is a level of difficulty in recreating a sequence of events that led to failure. For example, a two-instance Windows Azure service could have a hardware failure leading to one of the instances crashing. At the same time, the other instance could be brought down for scheduled maintenance, but Azure has mechanisms to protect against such as situations (in this case the planned maintenance will be delayed for a fixed amount of time). It takes a deeper understanding of various moving parts to diagnose an issue such as this.

However, the loss of control does not mean cloud platforms are a complete black box. One can associate detailed deployment policies with a cloud-based application. For example, even though Windows Azure storage geo-replicates content to another datacenter by default, it is possible to turn the geo-replication off. Similarly, it is possible to define the firewall rules, ports, tiers, etc. via the deployment model. In each of the aforementioned examples, the cloud platform is responsible for implementation details while adhering to the deployment model that was specified. Since the cloud platform is responsible for the implementation details, it can make necessary optimizations as needed and can ensure that different parts of the application are in close physical proximity in order to minimize latency without compromising the fault-tolerance requirements of the application.

Cloud platforms also check the software and hardware health of the application with [out-of-band monitoring](#) that uses an elaborate cascade of heartbeats spanning the guest OS, host OS and the controller. This monitoring allows for automatic restarting or replacing of machines in case of failure. In essence, by providing an application view verses a server-centric view of the application, cloud platforms are making up for the loss of control.

In summary, the loss of control of cloud-based applications needs to be weighed against the cloud's possible improvements in productivity.

## 7. TOOLS

**Claim: “Tools for managing cloud and on-premise resources are very different.”**

As stated earlier, cloud platforms are moving IT departments towards an application-centric view that fundamentally changes how we manage cloud-based resources. This is why the tools required to manage cloud-based applications are going to be different from more server-centric ones. Fortunately, platform vendors such as Microsoft and VMware are moving quickly to fill this void. For example, a recent announcement related to [System Center 2012](#) suggests it will provide a common management toolset for cloud and resources on-premise. Leading management tool vendors have also extended their offerings to include cloud-based resources. For example, IBM Tivoli Monitoring can be used to monitor the performance and availability of Amazon EC2 instances.

Finally, cloud computing is spawning an industry of new cloud-management-tool companies. For example, [Cumulux](#) and [Cerebrata](#) provide tools for monitoring, management and governance of Azure-based resources.

In summary, there is no doubt that cloud computing is changing the way IT departments will manage resources, but frankly, the recent consumerization of IT is already changing how we manage our resources.

## 8. MIGRATION COSTS

**Claim:** *"Migrating applications to the cloud entail non trivial costs."*

Before we talk about the costs involved in converting an on-premises application to the cloud, we must first discuss the business context for such a move.

If an organization is looking to free up on-premise resources (such as power, network and real-estate) by moving an infrequently-used application, they could **rehost** the application in the cloud. Such a conversion may not give them all the benefits of the cloud, but it will enable them to take advantage of shared infrastructure for power and their network as well as give them access to a pay-per-use model. The cost of rehosting an application is minimal, assuming the application is already enabled for virtualization.

If an organization is looking to go beyond rehosting and take advantage of cloud features, it will need to undertake at least some refactoring to the application. For example, to take advantage of horizontal scale-out, they must ensure the application is stateless. The changes also need to take advantage of cloud-based middle-tier services, such as caching, access control and Database as a Service.

In most cases, these changes can be made incrementally after rehosting the application. Furthermore, the cost of refactoring is minimized for platforms that share the same programming model between on premises and the cloud. For example, the Windows Azure platform programming model has minimal semantic differences between the cloud and Windows Server environment. Of course, the **API** is there to deal with cloud context, such as determining the number of cloud instances, status, etc., however, the key building .NET application building blocks are the same. Please review my detailed **tutorial** on converting an on-premises application to the cloud.

Finally, an organization may decide to completely replace the application with a SaaS platform such as **Dynamics CRM Online**—particularly for cases where the SaaS platform offers out-of-the-box capabilities that have a sizable overlap with the custom application. This organization could start with the existing features available in the online CRM and tweak them using light-weight programming techniques and configuration.

In summary, it is important to look at the cost of migrating an application to the cloud in the business context. Some applications can simply be rehosted while others will benefit from refactoring.

## 9. LEGACY APPLICATIONS

**Claim:** *"My applications rely on products that have fallen out of mainstream support."*

**Claim:** *"My applications rely on third-party products that are not cloud-friendly."*

Legacy applications can be moved to the cloud using the rehosting techniques described in the previous section. Since rehosting means that runtime, libraries and code assets are kept intact, it is possible for out-of-support and third-party components to be reused in the cloud. On the other hand, refactoring the application to take advantage of the cloud can be a herculean effort.

The biggest benefit for rehosting is the pay-per-use model and the ability to easily change the underlying hardware environment (the ability to move to a larger VM as needed).

There are two approaches that one could take when moving a **classic ASP**-based application to Windows Azure:



## IaaS Approach

Assemble the virtual machine by installing and configuring all the necessary software and instantiate the virtual machine as a VM Role instance.

## PaaS Approach

Deploy the ASP application to a Web role within Azure. Since the ASP engine is not configured by default, a startup task is needed to make this change.

The IaaS-based approach is particularly useful if the installation is fragile or long-running. The benefit of the PaaS approach is that OS is now maintained by Azure.

In summary, the cloud platform can offer the flexibility to rehost binaries and runtime files, thereby allowing legacy applications to be moved to the cloud.

## 10. TIGHTLY COUPLED APPLICATION PORTFOLIO

**Claim:** *“Due to rounds of mergers and acquisitions, our portfolio of applications is a tangled mess. So unless we move to cloud en masse (which is unlikely because some the applications have strict regulatory requirements), we cannot take advantage of the cloud.”*

It turns out that this scenario is quite common, especially in large enterprises today. This is where a hybrid cloud comes into play. As the name suggests, a hybrid cloud is a blend of cloud and on-premises-based resources. The idea is to decompose the application portfolio into portions that take advantage of the best of features each environment has to offer. Let us look at a few examples.

Many organizations are leveraging the cloud for a specific capability such as burst compute and storage, access control and messaging. My organization is currently working with a professional services [company](#) that is using an on-premise High Performance Computing (HPC) -based infrastructure for performing actuarial calculations. During peak loads they rely on the [Windows Azure-based HPC scheduler](#) to scale up (and down) as needed.

Another example is the use of cloud storage to augment the on-premise storage, especially for workloads that are not sensitive to latency such as backup and disaster recovery. Access to cloud storage is typically via a hybrid storage appliance such as the one provided by [StorSimple](#). The StorSimple appliance provides a tiered access to the data-on-premise layer for more active data and the cloud-based tier for inactive data.

The final example is the use of cloud-based messaging service, such as the [Service Bus EDI](#) to exchange EDI messages with partners without setting up a dedicated integration infrastructure.

All of the above examples are about tapping into the cloud for specific functionality like storage or messaging. In addition to accessing a specific capability, hybrid clouds are also being used for core IT functions, including setting up development and test environments to take advantage of rapid provisioning and overall agility.

Let us move beyond the hybrid clouds that are designed for specific capabilities to the ones that are designed for offload-enterprise workloads to the cloud. Consider the hybrid cloud in use at [Coca-Cola Enterprises](#). It consists of an SAP on-premise instance that holds all the actual order history data and enterprise LOB apps, including pricing, promotions and delivery. The on-premise data is then enabled via a Web-tier that is hosted on

Azure. Hosting the Web-tier in the cloud makes it possible to handle the varying load generated by nearly a million customers worldwide. The Web-tier communicates with the on-premise ESB endpoints.

In summary, whether the need is to leverage a specific capability—like storage or to offload an enterprise workload—organizations are increasingly adopting hybrid cloud patterns.

## 11. LOCK-IN

**Claim:** *“With no standardization in sight, cloud computing is just another way for vendors to lock us in.”*

**Claim:** *“Didn’t Google App Engine (GAE) recently announce a price hike that stunned its user community?”*

Frankly, with IT systems getting more complex and vendors looking for differentiated offerings, the potential for lock-in exists whether it is a cloud or on-premise-based installation. In fact, the risk for lock-in may be higher for an on-premise vendor offering a vertically integrated stack. This is mainly because cloud vendors are forced to align with the de facto interface boundaries that have emerged – whether it is the virtualization layer or the RESTful storage API layer. However, this does not mean that lock-in concerns related to cloud providers are to be taken lightly.

The key lock-in argument against cloud vendors is that cloud-based applications cannot be moved back on-premises while remaining cost-effective (whether an existing application has been updated to run in the cloud or an application was developed to run in the cloud from the ground-up). As discussed in the section on migration costs, certain PaaS platforms offer the same programming environment (i.e., .NET and Java) across cloud and on-premises. Such platforms make it easier to move the business logic between the two environments. Of course, there are differences such as accessing the configuration and runtime context, but these differences can be abstracted inside a common interface if one plans ahead.

This leaves us with the hardest part of the migration – accounting for services, including elastic storage and compute. There are no easy solutions until a symmetric on-premises service becomes available, either via a “cloud-in-the-box” offering like [Windows Azure Appliance](#) or as an enhancement to the traditional server environment.

What about the architectural changes? Turns out cloud-dictated architectural patterns (statelessness, scale-out, rolling upgrades, etc.) forced by the cloud works quite well on-premises.

Let us visit the other key lock-in concern—the risk of unexpected price hikes. GAE price was not unexpected as the service went from preview state to “real” production. On the contrary, we have seen AWS and Windows Azure engaging in a price war of sorts, resulting in **lower** prices on several recent occasions. There are good reasons why this trend will continue too—as cloud providers become more efficient in managing their datacenters, they will find innovative ways to reduce energy usage and costs.

Finally, let us discuss the concern about data lock-in. What happens to the data if an organization terminates their contract? It turns out that customers can add a clause to their enterprise agreement (AWS is a good example) wherein the cloud provider will be contractually obligated to ship all the data back upon contract termination.

In summary, lock-in concerns can be partially mitigated by sound design techniques, well-thought-out enterprise agreements and careful planning.

## 12. PRIVATE CLOUD

**Claim:** *“Thanks... we already have a private cloud effort underway.”*

A private cloud is an approach for building *on-demand, scalable, self-service compute and storage resources* for a single customer. It should come as no surprise that this description sounds like the definition for public cloud computing—because the concept of a private cloud is an evolution of the public cloud computing concept. Private cloud computing is motivated by the desire of some organizations to bring the advantages of cloud computing in-house, such as agility, utility billing and self-service.

However, we must be careful. Simply labeling a virtualization infrastructure as a private cloud would not be accurate. Even though virtualization is an important ingredient for setting up a private cloud, unless it comes with a self-service, on-demand provisioning of resources, metered access, well-defined SLA- and service-based interface, it cannot be deemed private cloud computing.

Despite the fact that several vendors are now offering private cloud tools, setting up one is challenging. Beyond the tools, it requires a cultural shift within the IT department, significant upfront resources and time, and above all, a certain scale to be effective. As stated earlier, a private cloud is designed for a single customer. As a result, it is built on a single-tenant model, which flies in the face of the key tenet of cloud computing—elasticity of resources.

Organizations need to evaluate the available cloud computing options before they embark on a private cloud effort. A private cloud should only be set up if your workload includes security compliance or other requirement that cannot be met in a multi-tenant setup.

Note: Embarking on a private cloud effort does not preclude the use of a multi-tenant cloud. In fact, many organizations will find that a combination of a private and public cloud will help realize all the benefits of cloud computing. As it is possible to undertake a pilot project using the public cloud with minimal investment, it can become a good way to learn about the cloud even if your organization ultimately decides to go the private route. In fact, this is exactly the model followed by Zynga, the creator of FarmVille. Zynga used the Amazon EC2 to launch games in order to better understand the load patterns, but once the game reaches a more predictable level, **it moves it to the private cloud.**

In summary, having a private effort underway is not a reason to dismiss the public cloud—quite the contrary. Organizations should lead with the public cloud, even if they undertake a pilot/proof-of-concept project.

## 13. IT STAFF

**Claim:** *“How can I leverage my IT staff in the new scheme of things?”*

Cloud computing represents a dilemma for IT departments. As enticing as benefits like low cost-per-unit and rapid provisioning are for cloud computing, not all applications can be moved to the cloud—this may be due to a specific compliance requirements, allowable latency, or a packaged product that is simply not ready for the cloud. But IT should not simply ignore or work against the adoption of the cloud. Instead IT professionals need to find ways to tap into the benefits of the cloud. They can start by figuring out which applications are well-suited

for the cloud and can start with targeted-use cases, like storage or access control. By incrementally moving these applications to the cloud, they can not only reduce the load on their ostensibly overloaded datacenters, but also give themselves more time to focus on the remaining on-premise applications.

The ultimate goal of cloud computing is to make IT systems more efficient to operate. One measure of datacenter efficiency is the average number of servers for which an IT worker is responsible. An IT worker in a traditional datacenter is typically responsible for maintaining 30 to 40 servers. On the other hand, an IT worker in a cloud-based datacenter, such as Windows Azure, may be responsible for as many as **5,000 servers**. Such a comparison may not seem fair to an IT worker responsible for maintaining a heterogeneous environment that includes not just different versions of an OS, but also different OSs altogether. But businesses have come to expect this kind of low-cost computing offered by clouds today. So whether an organization goes with cloud computing or to its own private variant, they are expecting a lean and mean IT department in the future. This is understandably causing consternation amongst IT workers.

So what can IT workers do to prepare for a future that includes cloud computing? They can adapt to the new role. IaaS-based cloud platforms do not need IT workers to provision servers and networks. PaaS platforms go even further; they do not need IT workers to maintain the OS by applying patches or hotfixes. IT workers will need to shift their focus to defining runtime architecture of cloud applications. For example, while developers build code for an application running on Windows Azure, IT workers should be able to create the **application topology** that defines the firewall rules, number of role instances, endpoints, etc. IT workers should acquaint themselves with cloud-management tools, defining and setting up governance policies, and establishing and setting up infrastructure needed to meet ROT and RPO objectives laid out by the organization.

Finally, a discussion of IT workers' roles in the cloud-era would not be complete without mentioning a new movement in IT management known as **DevOps** that is itself influenced by cloud computing. Most cloud platforms expose APIs for provisioning and managing infrastructure resources. DevOps sit somewhere between development and IT and relies on the aforementioned APIs to automate the management of cloud resources. To transition into a role of a DevOps, IT workers need to learn scripting languages and get familiar with agile development techniques such as TDD.

In summary, the role of an IT worker is undergoing a transformational change. IT workers who prepare for this change will be best to benefit the most.

## 14. LICENSING

**Claim: “I don’t understand how the cloud licensing works. What happens to my on-premises software assurance agreements?”**

Fortunately, with many hurdles associated with licensing software in the cloud now removed, the aforementioned question has a simple answer.

Let us start with a PaaS platform such as Windows Azure. Since Azure provides and maintains the guest OS environment on behalf of the application, the OS license is included in the per-unit compute charge. This is also true for a Virtual Machine role where customers assemble and upload a hyper-V image on their own. The PaaS platform also offers software components as a service. For example, while SQL Server on-premises is licensed by server configuration and a number of clients, its cloud equivalent, SQL Azure, is offered on a monthly subscription basis.

Let us look at how licensing works for a pure-play IaaS vendor. The IaaS vendors bundle the guest OS as part of per-unit charge. The whole point of IaaS is to have the flexibility to install whatever software you need on top of the guest OS. Of course, any software you install on top of the OS has to be appropriately licensed. Fortunately, most ISVs allow server licenses to be deployed in the cloud by applying them to a pre-configured Amazon Machine Image. One such arrangement is the [License Mobility Through Software Assurance](#) offered by Microsoft. Similar arrangements are available from Oracle, Sybase, IBM and others.

For organizations that are unwilling to procure licenses upfront through traditional or volume channels, vendors such as [Oracle offer on-demand](#) DB instances include licenses.

While we are on the subject of licensing, it is worth mentioning there are special perks for startups in the cloud. Microsoft just announced the [Biz Spark Plus](#) program that offers up to \$60,000 over two years.

In summary, licensing for software running in the cloud has seen tremendous progress. Licensing options, like the ability to “rent” applications by the hour are options that are simply unavailable on-premises.

## 15. PASSING FAD/IMMATURE STATE

**Claim:** *“Cloud computing is just the latest fad that will pass as soon as we hit the proverbial “Trough of Disillusionment.”*

**Claim:** *“Cloud platforms are immature in their current state. We will wait on the sidelines until the dust settles down.”*

Cloud computing, like many other technologies before it, has been over-hyped. Frankly, much of hype is due to vendors slapping the “cloud” label on just about everything they sell, whether it is traditional hosting or dedicated email service. This is why it is important to always keep the key tenets of cloud computing in mind: *on-demand, scalable, multi-tenant, self-service compute and storage resources*. Beyond these key tenets there is a never-ending supply of related concepts, such as Integration as a Service, private cloud, community cloud or even “cloud in a box.” But focusing on the key tenets will help you cut through the hype.

Please consider these stories from 2011 as evidence that cloud computing has gone beyond being just a fad or a shiny new object:

- AWS reached \$1 Billion in revenue in 2011.
- 30,000+ customers are now hosting applications on the Windows Azure platform. Azure represents the most strategic companywide initiative across Microsoft. Almost every major product (SQL Server, SharePoint, BizTalk, CRM, etc.) is being refactored for the cloud.
- Salesforce is on track for a \$3 Billion in revenue in 2012.
- IBM and SAP were compelled to add capabilities via large acquisitions of cloud vendors such as Demand Tech and Success Factors, respectively. Even Oracle, despite their initial misgivings about the cloud, decided to acquire Right Now Technologies.

In summary, cloud computing is evolving rapidly. Simply waiting on the sidelines while the market shakes itself out may not be prudent. In fact, it may even backfire as competitors come out with differentiated offerings that take advantage of the cloud.

## 16. OUTSOURCED

**Claim:** *“We have outsourced all our IT, please check with our service provider.”*

Most large organizations today are committed to long-term IT contracts that were written well before the advent of cloud computing, but that does not mean organizations have to be left out from tapping into the benefits offered by cloud platforms. As this [GigaOM](#) story suggests, organizations can:

- Review their current contracts and amend them if it makes sense.
- Identify workloads that are well-suited for the cloud and initiate pilot programs with or without the cooperation of the service provider.
- Be sure to include the cloud in the selection criterion for the next rebid.

As you would imagine, large service providers like IBM and HP are entering cloud computing as well. Needless to say, these companies will try to incorporate cloud computing into their existing outsourcing contracts, which is why customers need to educate themselves about upcoming disruptions. While the large service providers specialize in maintaining cost effective heterogeneous systems (Windows, Linux, Oracle, mainframes, etc.), they may not match the low cost-per-unit offered by pure-play cloud platforms like AWS and Windows Azure. In fact, according to a [PricewaterhouseCoopers survey](#), 55 percent of respondents would prefer to go with a pure-play provider in the next three years versus a traditional outsourcing services provider.

In summary, even if you have a long-term IT contract, there is little doubt cloud computing will play an important part in the next rebid.

## 17. WHAT GOES AROUND COMES AROUND

**Claim:** *“We are a mainframe shop— cloud computing is simply "mainframe rediscovered.”*

Indeed cloud-based datacenters are like giant mainframes where data and code are co-located (except cloud datacenters are built on racks of commodity servers), but the analogy goes only so far. Cloud platforms enable highly scalable and elastic services at a low cost. Compare that to the expensive capacity upgrade cycles mainframe customers have been accustomed to over the years. Not only do these upgrades represent a licensing challenge (whether to buy capacity in steps or go for a big-ticket upgrade), but predicting when an upgrade is needed can be a challenge in itself. Combine these factors with rising licensing costs, aging technology stacks and a shortage of skilled workers and you quickly begin to see why [Gartner](#) is predicting that 40 percent of the IBM mainframe application workloads will be modernized by 2015.

Of course, mainframe modernization efforts are not new. Customers have wanted to move their workloads off of mainframes for a while. I have worked on the development of a mainframe modernization product for six years. The key concern for these customers has been to select a runtime environment that:



- Can be as robust as mainframes are known to be.
- Will be able to support large workloads common on mainframes today (upwards of 1000 to 2000 MIPS).

While Windows- and Linux-based solutions, with the right target architecture, have proven to be sufficiently resilient, cloud computing with elastic scalability and fault tolerance makes an ideal environment to host moving mainframe workloads. Companies like [Alchemy Solutions](#) have announced [NeoKicks for Azure](#), making it possible to rehost CICS-based workloads on Windows Azure. They do so by recompiling the existing COBOL code to .NET and hosting the resulting binaries within an Azure Web- or worker-role.

In summary, if you are a mainframe shop looking to modernize some of your workloads, cloud computing may be an option you want to explore.

## 18. BUDGET

**Claim:** *“We have zero budgets for any new initiatives.”*

Given the recessionary state of our economy, private businesses and public organizations alike are saddled with budget cuts across the board. So it should come as no surprise if the suggestion to move to the cloud is met with a tepid response. Interestingly though, the Opex (Operational Expense)-based cost model makes cloud computing one of the few viable approaches for modernizing the IT infrastructure in such conditions. An Opex-based model removes the need for capital expenditure upfront. Clearly, there is a cost for migrating the application to the cloud or customizing an SaaS-based service, but cloud vendors and system integrators, in part to promote the use of cloud, are willing to subsidize the implementation costs—of course, they do hope to recover their investments through subscription fees.

Moving specific workloads to the cloud can also be attractive to organizations with limited IT capabilities. Moving workloads like ERPs or CRMs (a well-established model now) to a cloud provider will allow organizations to focus their scarce resources on mission-critical applications.

In summary, given the limited budget for new capital expenditures, businesses should consider a subscription-based operational expense as a way to modernize their infrastructure.

## 19. PERFORMANCE

**Claim:** *“Cloud datacenters rely on commodity hardware that can degrade performance.”*

Cloud providers rely on low-cost commodity hardware to build humongous datacenters. While this is great for applications that can take advantage of scale-out-based architectures, it can lead to interesting performance limitations.

Consider a situation where the usual techniques for performance optimization, such as examining query plans, has been applied to no avail – the I/O bottlenecks continue to persist. In such situations, the usual recourse is to add a *faster* I/O controller. Unfortunately, this is not an option available in SQL Azure since it is based on a multi-tenant shared infrastructure.

However, there are techniques to get around some of these limitations. Notice that the word “faster” has been italicized in the previous paragraph; this is because moving to a faster I/O controller is not the only option for improving I/O—you can go with *additional* controllers as well. In fact, SQL Azure recently added the SQL Federations feature that makes it easier to partition your database (in essence adding *additional* controllers).

It is also interesting to note IaaS providers such as [CloudSigma](#) are now offering the ability to create a personalized combination of CPU, storage and RAM that includes SSDs. This is in contrast to AWS and Windows Azure where one needs to choose from a standard set of bundled configurations. It is probably also worth noting that DynamoDB, introduced by AWS, recently began running exclusively on SSDs.

In summary, whether it is making the partitioning of data simpler or making it possible to select a custom-hardware configuration, cloud providers are going to be looking for ways to deliver improved performance.

## 20. NOT BUILT HERE

**Claim:** *“Government regulations are going to require us to use a national cloud provider.”*

Concerns about data leaving the national boundaries, warrantless data surveillance and, frankly, protectionism are sparking demands for national regulations around cloud computing. Recent stories including Microsoft’s UK managing director [admitting](#) the reach of the Patriot Act outside the U.S. has exacerbated these concerns. According to a recent Bloomberg [article](#), the German government is working on stricter data-protection rules and French companies are creating a “Built-in-France” version of the cloud infrastructure.

Fortunately, industry realizes that an important facet of cloud computing is its global scale. U.S.-based tech companies including Google, IBM and Microsoft are urging the U.S. government to establish [global treaties](#) that will enable a free flow of information. Certifications, such as ISO 27001, are also being recognized by countries like India. The Windows Azure Platform recently [achieved](#) the ISO 27001 certification (Microsoft’s Global Foundation Services division has a separate [ISO 27001 certification](#) for the datacenters in which Azure is hosted.)

In summary, cloud customers and providers alike are realizing that creating nationalization of the cloud is not in anyone’s interest. Instead, we need transparent and effective international laws that will foster innovation and allow the technology to realize its full potential.

## CONCLUSION

Moving to the cloud is not going to be without challenges. However, cloud computing is turning out to be such a disruptive and transformational technology that few businesses can afford to ignore it.





## WHO WE ARE

Founded in 1982, AIS provides software and systems engineering services to companies and the U.S. government. AIS builds software solutions for forward-thinking IT departments, including human workflow applications, portals and websites, enterprise content solutions, digital archives and more – in and out of the cloud.

AIS is a top-tier Microsoft partner and a recognized SharePoint and .NET expert, having built many solutions that provide organizations with the ability to collaborate and share information across the enterprise. AIS has early adoption expertise with Azure, Microsoft's cloud services platform. AIS' CTO, Vishwas Lele, is a Microsoft Regional Director, and other AIS staff are members of several exclusive Microsoft partner advisory councils, Virtual Technology Specialist (VTS) programs and other Microsoft readiness programs. AIS is also a member of Microsoft's NSI partner program.

AIS has a Top Secret Facilities Clearance and over 75 percent of employees hold Secret, Top Secret and Top Secret/SCI clearances. AIS is headquartered in Reston, Virginia with additional offices in Columbia, Maryland; Dayton, Ohio; Austin, Houston, and San Antonio, Texas; and Hyderabad, India.



[WWW.APPLIEDIS.COM](http://WWW.APPLIEDIS.COM)

800-AIS-4553