

OWNER'S **IT** MANUAL



# Cloud Computing: Optimizing application performance

IT pros share insights on their cloud deployments,  
and how to create a great user experience



# Deliver a better cloud experience

Experts explain how to successfully navigate the cloud waters, and weigh in on the big issues: Evaluating cloud providers, cost, latency, provisioning, performance monitoring, and creating the ultimate user experience

BY DAVID STROM

**BUILDING THE BEST CLOUD** apps isn't a matter of just copying some software to a cloud-based server - it's understanding that "your apps need to be engineered for the cloud from the beginning," says Muhammad Chaudhry, the Director IT Operations at The Motley Fool investment site. "This is because the on-premises infrastructure is configured differently and has different resource requirements from what is needed in the cloud." The Motley Fool site was designed from the beginning to operate in the cloud, long before the cloud was even used as a descriptive term, and they worked carefully to create the most appropriate apps for their business.

Not every company has the same opportunity from the start. Most grapple with various legacy systems, and do a fair bit of retrofitting along the way. But

wherever you are with your cloud projects, there are lessons to be learned from other's hard-earned experiences. This Owner's Manual reviews some of the major issues concerning how to make cloud projects successful, both from the standpoint of everyday users as well as the IT departments that watch over them.

As you get deeper into cloud computing, there's a lot to know and do. You must understand network latencies of your cloud apps, and how to reduce them. You'll also need to evaluate cloud providers, as well as improve on cloud provisioning, monitoring and cloud performance -- all while anticipating and managing costs. Last, you'd be wise to involve your users in the decision-making process when building any new cloud apps. Are you ready?

## What does the cloud cost?

One of the big benefits about the cloud is a reduction in the cost of IT infrastructure. According to

## By the numbers

**61%**

IT leaders who said their companies have at least one application or a portion of their computing infrastructure in the cloud

**\$1.5M**

Average investment in cloud-based services during the next 12 months -- a 10% increase over 2012

**24%**

Cloud spending that will originate outside of the IT organization within 3 years vs. 20% today

**LOWER TCO**

Most effective internal selling point for a private cloud

SOURCE: 2013 Cloud Computing: Key Trends and Future Effects, 1,358 respondents, conducted across IDG Enterprise brands: CIO, Computerworld, CSO, InfoWorld, ITworld, NetworkWorld



# OWNER'S IT MANUAL



Computerworld's 2013 State of the Enterprise survey, 49 percent of 345 respondents agreed that the cloud would reduce their operating costs.

But actually figuring out your cloud computing bill is a real challenge. More than half of 1,682 respondents to a 2012 Cloud Computing survey by IDG Enterprise said that, in general, their contracts with cloud vendors weren't structured to easily evaluate costs. This is because the cloud service providers do a great job – like a 1960s-era car salesman – of hiding the actual pricing numbers, and presenting you with long lists of a la carte pricing that can confuse and confound. Some don't even give you access to a full price list until you sign up for their service, or hide this information behind a Web calculator page that has dozens of variables, many of which you won't really know until you begin using their services. Some vendors sell servers in pre-packaged options with a given CPU, RAM and disk configuration that you can start from, while others have you build your virtual machines from scratch where each item tacks on a particular price. Sadly, only a minority of the cloud vendors has simple pricing models that don't overwhelm you with these details.

On top of these issues, many providers change their prices frequently and without much notice, rendering all your comparative research obsolete. Granted, they are usually reducing prices, but still these changes are frustrating to monitor.

Next, despite claims that you pay for what you actually use, that isn't always the case, and isn't always consistently applied across a single vendor's entire cloud services line. Some providers charge for configured virtual machine (VM) or VPN instances whether or not they are actually up and running. Others set monthly usage maximums and charge for overages, just like your cell phone company when you exceed your data plan. That means more work for you to keep track of whether or not you are below your plan's cutoff level.

Finally, you must distinguish between capital or one-time costs and recurring monthly operating costs. Some enterprises are more adverse to one or the other, and it sometimes can be counter-intuitive. David Goodman is the CTO at the International Rescue Committee (IRC) in New York, NY. The non-profit organization supports more than 150 field offices all over the world, some in rather out-of-the

## Migrating your apps to the cloud

---

- 1 Build apps from the beginning with the cloud in mind**
- 2 Keep track of the monthly cloud computing bills, and understand how they are calculated**
- 3 Build or find appropriate tools to monitor your apps' uptime**
- 4 Understand network latencies and end-to-end performance**
- 5 Involve your users and make them part of the decision process**
- 6 Think about using the cloud for testing new apps too**

way places. “Some of our field offices are located in refugee camps that barely have any electricity, let alone an Internet connection.” The IRC can often get one-time funds to cover large IT purchases, but covering a monthly subscription for a cloud service can be a problem. “When we use the typical cloud subscription model, that causes us to add significantly to our monthly overhead, and that is a lot harder for us to handle.” As a result, he must be careful about what he moves into the cloud.

### DO YOU KNOW WHERE YOUR SERVER IS?

*“These are offices where servers could be sitting under plants or in broom closets, and we wanted to get them centralized and virtualized.”*

.....  
**Bill Gillis**, CIO, Beth Israel Deaconess Medical Center, Boston

### Latency issues

Once you understand the cost of your cloud app, the next biggest issue is getting your hands around latency issues. Network engineers have long studied the effects of latency on applications performance, reducing router hops and increasing router packet processing. No one knows this better than the IRC’s Goodman, who has to deal with two-second latencies across satellite networks and over lousy Internet links across the world. “No one is really building cloud apps that deal with these huge latencies,” he said. “And simulating and testing apps under these conditions is also really difficult. You really need plenty of Internet bandwidth for even the simplest cloud app.”

Even with all their latency challenges, the IRC is in the process of migrating from Sharepoint to a

combination of Box and Django. “We had so many problems with the network connections and bandwidth requirements for Sharepoint that we wanted to design an app that can provide a very lightweight front-end to a robust back-end with some sort of offline synchronization of documents,” Goodman said. “We liked that we could restrict the amount of bandwidth in the Box client to match our needs.”

Even if you have this bandwidth, the cloud complicates understanding network latency because of several factors that make its calculation a lot more unpredictable:

First, endpoints aren’t fixed like they are for most on-premise situations. Users can be located anywhere ranging from high-speed fiber connections to 3G mobile data networks. This makes latencies sometimes horrific and often unpredictable. Second, the cloud infrastructure may not be optimally connected to your own, and it may be difficult for IT managers to even calculate the built-in latencies of the cloud provider’s network. Third, applications are becoming more virtualized and distributed across large-scale computing infrastructures, such as Hadoop clusters of hundreds of machines. This introduces additional latencies.

Finally, IT may not be completely aware of the ultimate end users and application owners, nor have the right service-level agreements (SLAs) in place to enforce minimum latency standards. Many SLAs specify ping or traceroute transit times, but most modern applications use other protocols that don’t necessarily correlate with ping times. Many SLAs also don’t differentiate among outages on a server, a network card, a piece of the routing infrastructure, or a security event.

Sorting this all out isn’t easy, as Bill Gillis found out. He is the CIO for Beth Israel Deaconess Medical Center in Boston. They built a cloud-based electronic medical record system over six years ago that is hosted at their Internet provider and supports

# OWNER'S MANUAL



the clinical practices at their smaller doctor offices in their network. “These are offices where servers could be sitting under plants or in broom closets,” he said, “and we wanted to get them centralized and virtualized. We call it our accidental cloud, because we didn’t start out to build it that way. But these are very mission-critical apps: without them, a doctor’s office couldn’t see any patients.” Beth Israel owns all the hardware, software and other infrastructure, and each office connects to their app via just a Web browser and an SSL connection.

When they got started on their cloud project, they found a lot of resistance from the vendors of clinical care applications. “Healthcare IT is still very conservative,” he said. “Many of them told us that they didn’t support our cloud environment, even though they were selling retail cloud-based versions of their products.”

Their system usually performs very well under a wide variety of Internet connection methods, but when they need to debug an issue, they have to install specialized application and network performance management tools to track down latency and packet losses. “We just ship the practice office a monitoring box in the mail, and the office manager just plugs it in to the network. In a week, we can capture enough data to figure out where the problem is and get the right vendor to fix it.”

## Evaluating the cloud provider

While Beth Israel has a private cloud collocated at their Internet provider, they want to get out of that business and move their app to a cloud provider in the near future to leverage better support. They have found it isn’t easy to evaluate different cloud providers without some additional work, and ended up “cobbling together some home-grown tools for performance comparisons and aggregating various server log files. If we can satisfy our performance

requirements, and if the cloud vendor can provide the level of security that we need, then we can pick the best provider,” Gillis says.

But other enterprises have more complex requirements than just minimal packet loss. They have found that some cloud providers are cagey about describing their infrastructure, and some don’t even tell you where your apps are running or where their data centers are physically located. Some customers have had to sign non-disclosure agreements to obtain this information.

Several years ago, Boeing security engineer E.J. Jones was quoted discussing how he initially **designed a five-part checklist and graded each cloud provider**. The requirements included questions such as, “Can a provider tell us when, and how a failure has occurred?” and “Can they guarantee uptime?” More recently, The Boeing Co. has employed a third-party auditing firm that was allowed more detailed access to the cloud provider’s methods. Boeing recommends that when you are looking to evaluate each cloud provider you should have consistent IT controls in place. These should include standardized deliverables and touch points between your own IT organization and your cloud providers. You can see the level of their analysis with the infographic (See next page) that Boeing’s Stephen Whitlock, Chief Strategist, Information Security, shared at the **2012 Gartner Catalyst conference** in terms of evaluating each cloud vendor for storage, platform and hosting services.

## Getting better at provisioning

One of the more often touted benefits of the cloud is its ability to scale up and down as demands and needs change. But the dirty secret behind this ability is being able to automate the provisioning of the cloud-based VMs. This includes policy-based workload management and deployments, and real-time

## Some Security Characteristics

very secure/recommended for this use
  some vulnerabilities/use with caution
  very insecure/not recommended for this use

	Capability Maturity	Information Risk	Technical Control Effectiveness	Contract Control Effectiveness	Application Interoperability	Enterprise Integration
<b>SERVICES</b>	●	●	●	●	●	●
<b>APPLICATION</b>	○	●	○	○	●	●
<b>DEVELOPMENT</b>	○	○	○	○	●	○
<b>PLATFORM</b>	○	●	○	○	○	○
<b>STORAGE</b>	●	○	●	●	●	○
<b>HOSTING</b>	●	●	●	●	●	●

**SOURCE:** Stephen Whitlock, Chief Strategist, Information Security, The Boeing Co., This slide was part of a presentation at the 2012 Gartner Catalyst conference.

resource monitoring. Some cloud providers are using orchestration tools that automatically start or stop particular VMs in sequence so that a directory server can start ahead of a database server for example.

This is what The Motley Fool IT department realized early on. “You want to take the human element as much as possible out of your deployment and provisioning process,” said Chaudhry. “This helps you to minimize failures, and realize higher returns on your investment.” He suggests using various frameworks such as Chef and Puppet that can help with this automated configuration and provisioning management.

There are other tools that can correlate and analyze information from cloud infrastructure, application performance and other IT management tools in real-time. This information is used to map and dis-

play IT assets that deliver specific business services, calculate service quality, and identify which IT assets impact service quality and put it at risk. The idea is to model an end-to-end and real-time view of services across the enterprise, regardless of whether these services are on-premises or cloud-based.

### Involvement of users from day one

Another aspect of understanding the cloud is how you need to work collaboratively with your end user departmental managers. They must be part of the process too. This is the philosophy used by Ravi Ravishanker, the CIO at Wellesley College. He has made it one of his goals to get his users involved to help him get out of the hardware business and migrate as much as he can into the cloud. They have been using



both Google Apps and Drive for more than a year to host their email and documents, and want to acquire more storage from Google than the 30GB that is allotted per user presently. “Some of our departments are putting their lab manuals and other documents on Google Drive. Its ease of collaboration and sharing makes it efficient for our users,” he said. “The proposed cost of purchasing additional storage in Google Drive when compared to a local SAN array and its backup locally makes it hard to justify the latter. We still support local storage for storing secure and sensitive data”

“All of our decisions are made with input from the various user departments, to make sure that they are vested in the decisions,” he said. “This way they

accept them more and see the value in the move towards cloud computing.”

As an example, the college’s art department managed and used high resolution images scanned from their slide collection on hardware running locally in their department. “We looked at

various solutions and decided to use an excellent cloud based service called Shared

Shelf from ArtStor, which had an excellent ROI that was functionally superior, far less maintenance and reasonably priced, resulting in high user satisfaction.”

The college has also migrated its learning management system and video storage to cloud-based providers as well as software used by health and counseling services too. “We do the integration and the user authentication pieces ourselves.”

## Monitoring performance

Certainly anyone who is moving their apps to the cloud should have some idea of performance to expect from their cloud apps, both in terms of actual statistics and also perceptions of uptime too. This is where real-time monitoring comes into play. It is increasingly seen as part of the package that a cloud hosting provider offers its customers. While there are many examples, the most noticeable one is the series of near-real-time dashboards from Facebook that show its power and water efficiencies for its data centers ([The dashboard for the Oregon facility can be found here.](#)).

The cloud changes the way IT responds to outages also. “If Google or Facebook had to answer to an angry customer every time an instance went down because of a 0.05% increase in failure rate, they might choose a different deployment strategy,” says Philip Molter, the CTO of Data Foundry, a major data center service provider in Austin TX.

Chaudhry agrees with this philosophy. “You have to factor in what happens when a machine goes down in one region in the cloud [and] how you will need hot standbys in other regions. You don’t have complete end-to-end control over your data center in the cloud, so you have to take [into] account infrastructure failures and build in redundancy and resiliency.” This also means that you need to be proactive in monitoring your cloud resources, where and how your services are being used, and loads. “The best companies have put a lot of effort into this,” he says.

### NOW WHERE'S THAT 'DISLIKE' BUTTON...

*“If Google or Facebook had to answer to an angry customer every time an instance went down because of a 0.05% increase in failure rate, they might choose a different deployment strategy”*

**Philip Molter, CTO, Data Foundry**



# OWNER'S IT MANUAL

Despite the numerous Google outages, the perception can be that an IT department is actually delivering better service, because users have the perception of control of their own apps. “Wellesley College users have been mostly pleased with the uptime on its various cloud apps. Most are more stable with much better uptimes than if we were running them internally,” says Ravishanker.

The college central IT department uses several open source third-party monitoring systems, as well as dashboards from several service providers to monitor Internet connections and services – since moving more apps to the cloud they have had to boost their bandwidth to a gigabit connection – and see how the bandwidth is being consumed by various services in near real time. “But the real monitor is hearing from users when they have issues,” he said.

Major public cloud providers are getting the message though, and documenting their management and monitoring APIs and working with third-party monitoring services to be more open and provide their customers with more details about their services.

## Clouds are for testing too

The cloud can even be used to support testing new apps and configurations out, which is what is being done at Wellesley College. “You can sign up for a cloud service with a credit card and get a VM up and running within three minutes. We can’t provision a VM in-house that quickly, and it makes some IT staff feel empowered by doing this. This means that they can install and check new academic software, they aren’t afraid to mess something up and we can wipe it clean, and they can start again,” says Ravishanker.

The Motley Fool’s Chaudhry also recommends making frequent deployments and updates of small changes to your cloud apps. “This minimizes the chances of large catastrophic failures. Ideally, you should aim for minor releases every week or so.”

---

### About the author:

*David Strom writes and speaks about networking and communications topics.*