

Cloud
Economies of
Scale and
Management

Keeping
Control Isn't
Easy

Cloud Pricing
Explained

The Future
of IaaS



Keeping Control Isn't Easy

Buying Infrastructure as a Service is a mouth-watering prospect for CIOs and IT departments eager to drive down costs and speed up provisioning. But once IT is outsourced, keeping control of it is a major challenge.

Keeping Control Isn't Easy

By Jo Maitland

Cloud Economies of Scale and Management

Keeping Control Isn't Easy

Cloud Pricing Explained

The Future of IaaS

Buying Infrastructure as a Service is a mouth-watering prospect for CIOs and IT departments eager to drive down costs and speed up provisioning. But once IT is outsourced, keeping control of it is a major challenge.

INFRASTRUCTURE AS A SERVICE—or the ability to rent servers and storage by the hour in a pay-as-you-go way—is not only real but also saving companies millions of dollars and fundamentally changing the way IT organizations operate and support the business.

In 2006, Amazon.com Inc. launched its Elastic Compute Cloud (EC2) Infrastructure as a Service (IaaS) and

unconsciously kick-started the utility computing market. Now companies buy servers as fast as you can buy a book, thanks to advances in virtualization technology, service-oriented architecture, secure multi-tenancy and the proliferation of broadband networks.

But watching this shift in IT is like witnessing a toddler taking his first steps: all smiles and confidence and, then, bam! He falls flat on his face. Buying Amazon Web Services is disarmingly simple, so watch out. This is only the beginning of a long shift toward IT as a Service, and there will be many false starts and painful bruises along the way.

That's not to say don't try it. Companies that have dipped their toes in have learned quickly. And as the model matures, they will be ready to take full advantage of cloud computing.

HOW ONE GROWING FIRM USES AMAZON'S EC2

When online investment community KaChing Group Inc. arrived on the scene about a year ago, the company

Cloud Economies of Scale and Management

Keeping Control Isn't Easy

Cloud Pricing Explained

The Future of IaaS

found it difficult to plan business and IT growth based on unpredictable traffic patterns. "We are not in the business of trying to figure out how to scale our architecture up and down depending on network traffic. That's what we have Xignite for," said Pascal-Louis Perez, the CIO of Redwood City, Calif.-based KaChing.

Xignite Inc., an on-demand financial market data provider, hosts its application servers on Amazon's EC2 Infrastructure as a Service platform. These servers in turn feed real-time financial data, such as stock quotes, Securities Exchange Commission (SEC) filings, currency rates and mutual fund performance, to customers such as KaChing, Forbes and Citigroup Inc.

KaChing's computation engine, which also resides on Amazon's cloud, takes this real-time data to update the fantasy stock portfolios of its 400,000 registered users. Over time the company, which is a registered SEC investment adviser, will become a place where users follow site member trading experts to make real hedge fund investments.

Unlike the monthly or even hourly contracts companies can sign up for with many on-demand service providers, Xignite customers typically sign an annual contract and pay a monthly fee for server capacity and network bandwidth. But given that customers such as KaChing often reside in the same cloud on Amazon's ecosystem, network bandwidth costs are minimal, said Leo Chan, the CTO

of Xignite in San Mateo, Calif.

"We don't like charging for [server capacity] overages. So when we see that happening, we sit down with the customer to elevate their contract to the next [capacity] tier," said Chan, who would not disclose the company's pricing.

On a small scale, charging per instance makes sense, in that you pay for the minimum capacity you need as you go. Experts argue that pricing models based on usage do not scale well for large businesses with heavy transaction volumes. But if you are a startup more concerned with gradual and affordable growth, such on-demand pricing models make sense.

"The pay-as-you-go [pricing] model allowed us to grow little by little and use incremental [Xignite] services without up-front infrastructure investments," KaChing's Perez said. "When we built the prototype, [our servers] were only doing a hundred calls a day to Xignite's service, and we had 1,000 users. We're now at 400,000 users, so you can imagine the amount of scaling we would have had to do for that."

CLOUD ECONOMIES OF SCALE AND MANAGEMENT

Infrastructure as a Service comes in many flavors. Users looking for more granular network control might find that Amazon Web Services, geared primarily toward building Web apps and Web communities such as Ka-

Ching, does not offer enough knobs and dials to control how the service is used.

In mid-2008, Preferred Hotel Group, a luxury hotel chain headquartered in Chicago, moved all its IT opera-

tions (including Microsoft SQL and Exchange servers, Citrix Systems desktops, file and print servers as well as various utility servers) to service provider Terremark Worldwide Inc. Prior to using Terremark, the hotel

Cloud Economies of Scale and Management

Keeping Control Isn't Easy

Cloud Pricing Explained

The Future of IaaS

Public vs. Private Cloud Infrastructure

INFRASTRUCTURE AS A SERVICE (IaaS) comes in two forms:

1. **Pay-per-use hosting of virtual servers at an external, or public, cloud service provider.**
2. **By operating an internal, or private, cloud, where your IT department offers virtual servers as a service.**

Cambridge, Mass.-based Forrester Research Inc. found the following indicators of likely buyer interest in, and adoption of, these two forms of IaaS:


- **About 25% of all enterprises plan to adopt IaaS via an external service provider.**
- **Firms are slightly less interested in internal clouds than they are in external IaaS.**
- **Large-business respondents report more awareness, interest and adoption of external IaaS than do small businesses. The figures are the same for internal clouds.**

These results are contrary to conventional wisdom regarding the initial demand for cloud services. Enterprises are leading the adoption, not small and medium sized businesses (SMBs). Moreover, they have different technology preferences and comfort levels with virtualization.

Forrester also believes that early adopters of IaaS service offerings are driven by the instant provisioning of servers and the pay-per-use pricing model. Further, unlike developer buyers, enterprise IT operations buyers may want to integrate their on-premise infrastructure with anything they deploy to a service provider, either temporarily or permanently. VMware, Citrix and others now push the idea of "bursting" from a private cloud into a public cloud for additional capacity as needed, but this is still at the very early stages, and neither company can provide a customer that is bridging from a private to a public cloud yet. —J.M.

Is the cloud a bright idea for your business?

➔ It's time you moved your IT infrastructure to a cloud computing solution. Let Layered Tech give you a **Virtualization Assessment**. It will give you a clear idea of whether cloud computing is the right approach for your business right now. The hardware and network resources at our world-class data centers allow us to offer you amazing levels of availability, processing power and scalability – without investing in any hardware yourself. We offer a wide range of technology agnostic solutions and five levels of support so we can design the solution that's just right for your enterprise and makes the most of your IT resources.



Email ebooksalestt@layeredtech.com, visit www.layeredtech.com/tta4 or call 1-866-584-6784 today.

5 LEVELS OF SUPERIORITY



chain hosted its infrastructure at Savvis Inc. but found that too expensive, prompting the switch to Terremark.

In between the move, the company calculated what it would cost to bring everything in-house. For a Dell blade server, a mid-sized storage area network (SAN), VMware software and the required power and cooling equipment, the hotel chain faced a \$250,000 up-front investment. Meanwhile, Terremark offered all the gear plus disaster recovery (DR) to a second site for \$16,000 up front and an additional \$12,000 per month. The Preferred Hotel Group currently has 36 virtual machines (VMs) and approximately 3.5 TB of storage at Terremark.

"It works out at a little more a month, but we don't have to worry about hardware maintenance and disruptive software upgrades," said Chad Swartz, an IT manager at Preferred Hotel Group. He says one of the biggest bonuses of using a cloud provider is getting your hands on top-notch IT gear. "They can afford a whole different class of SAN. I can look at Dell; they can look at Hitachi and a lot of other devices that make things run much faster than we could." Swartz noted that while disk prices have come down, the real expense of a SAN comes in managing it. "There's no assigning LUNs [logical unit numbers] in the cloud. It's so much simpler ... and we avoid that staffing cost," he said.

So if simplicity and cost are your primary concerns, why not use Amazon Web Services? Swartz said Amazon does not have all the controls that Terremark does. "Most cloud services are

"There's no assigning LUNs in the cloud."

—**CHAD SWARTZ**, IT manager,
Preferred Hotel Group

geared at building websites, not running infrastructure," he said. "Terremark provides more granular controls over the way servers are built ... the setup of DMZs [demilitarized zones], subnets and the ESX environment."

That said, Swartz would like even more control. Right now he has no way of seeing how much Internet bandwidth he uses or how many virtual machines are provisioned on the servers his company pays for, especially for DR purposes. "We have to trust that they are all on different servers for VMotion [VMware's feature for live migration of VMs], it's a good-faith thing, but I would prefer to see it," he said.

KEEPING CONTROL ISN'T EASY

Daryl Ford, the director of communications and infrastructure services at the University of Massachusetts

Cloud Economies of Scale and Management

Keeping Control Isn't Easy

Cloud Pricing Explained

The Future of IaaS

Daryl Ford, the director of communications and infrastructure services at the University of Massachusetts, Boston

Cloud Economies of Scale and Management

Keeping Control Isn't Easy

Cloud Pricing Explained

The Future of IaaS



Boston, is a big believer in cloud computing but knows what it's like when users run amuck in the cloud. At one time, a faculty member at the university quietly used a cloud service to back up 20 GB a night over the Internet. "It was bringing our network to its knees," Ford said. "We shut him down."

Ford says UMass has all the right policies concerning IT procurement and enforces them as best it can, but he said there will always be cracks in the system. "People could care less about policies. ... They want what they want when they want it and they don't involve IT." And this goes for cloud services in particular, as they are so

easy to buy, he said.

Pharmaceutical companies are another sector eager to use cloud infrastructure services to alleviate IT scaling issues, but concerns about security and compliance have prevented many of them from jumping in.

Almac Clinical Technologies has grown fast. In a matter of months, the size of its clinical trials has jumped from hundreds of people to tens of thousands. Getting the infrastructure deployed to support these trials in a timely way is almost impossible, according to Jack Kosowsky, the director of systems development at Almac. But before he can consider using a

service such as Amazon EC2, he said protecting the company's intellectual property in the cloud is "critical."

So security is a huge issue. The data associated with clinical trials is extremely sensitive. In the context of a multi-tenant environment, imagine if Merck got hold of Pfizer's data on a clinical trial. Thus pharmaceutical firms have stringent change-control requirements. "You can't turn a screwdriver in our data center without documenting it," said Kosowsky. The firm must show that its IT systems maintained the same state throughout the duration of a trial. Just applying a service pack to Windows requires a

Cloud Economies of Scale and Management

Keeping Control Isn't Easy

Cloud Pricing Explained

The Future of IaaS

Advice from the Trenches

PHYSICAL-TO-VIRTUAL migration is out there, but IT pros say that in many cases it's best to build servers in the cloud from scratch—and SQL servers in particular, as the drivers operate differently in the cloud from those in a traditional environment.

Data synchronization is a big deal and will likely force some downtime as companies do the initial copying of their data to the cloud. In many cases, companies physically ship data because they have too much to send over the wire. What does your service provider offer to make this transition easier? Until you are certain everything in the cloud runs as it should, keep your old environment up and running.

- **Request itemized bills that explain all the charges per contracted period. Beware of hidden extras, and question everything.**
- **Consider the implications of cloud lock-in. Ask your cloud provider which archiving and virtualization standards it supports and, if you wanted to move to a new cloud environment, the kind of services it offers.**
- **Find out whether your cloud provider has insurance for data privacy breaches and network interruptions as well as how much the insurance covers. —J.M.**

lot of thought and testing to ensure that the system can operate in a continuous state. The company's traceability matrix shows how even the slightest change affects every system in the IT department.

Before undertaking a trial with the firm, Almac's clients sometimes visit the company's data center to check out its security. "They take great comfort from a nicely laid-out data center with blinking lights. If we say [their data] is somewhere in the cloud, it's not going to give them that comfortable feeling," Kosowsky said.

Reliability is another deterrent from cloud services. "A daylong outage would be the end of cloud computing for clinical research," Kosowsky added.

CLOUD PRICING EXPLAINED

Even if the cloud's control and security issues were suddenly resolved, the morass of cloud pricing would still be a challenge. Although cloud providers like to tout the simplicity of their services, IT managers have found that pricing cloud services is anything but simple. Pricing structures are based on a multitude of factors, from storage space needed to clock cycles used to monthly traffic allotments, and that's not all. Some service providers have additional charges hidden deep within their service-level agreements (SLAs).

To determine total pricing for a cloud service, users need to understand the individual service elements

that a provider bills for and how those charges are calculated. Does the provider, for example, bill based on traffic, storage space needed, server CPU time or a combination of these factors along with other elements?

IT managers have found that pricing cloud services is anything but simple.

Another critical factor in determining true costs comes down to the type of service needed. For some, that service may be little more than a hosted, dedicated server to run applications in the cloud. For others, the service may be cloud-based backup or business continuity or basic hosted storage.

Perhaps, the easiest way to break down pricing is to focus on the primary services offered. The majority of cloud service providers break down their services into three primary areas: servers in the cloud, storage in the cloud, and sites and applications in the cloud. Each is governed by its own formula for pricing.

Servers in the cloud come in two forms: virtual and physical. In other words, you can purchase time on a virtual server (where the physical hardware may be shared with others) or time on a dedicated server (where you

Cloud Economies of Scale and Management

Keeping Control Isn't Easy

Cloud Pricing Explained

The Future of IaaS

are the only “tenant” on that server).

TABLE 1 shows the pricing breakdowns.

The pricing comparisons are only part of the overall picture, each vendor listed above includes extra services and features for an additional charge. What’s more, prices can change depending on length of commitment, total bandwidth needs or total size of storage required, and in most cases, prices and packages are negotiable with a vendor’s sales staff.

Not all cloud service providers are created equal, and that becomes evident when you look more closely at the differences between providers and how they address customer needs. For the purposes of comparison, we picked three of the well-known cloud infrastructure providers:

- **GoGrid** includes load balancing at no extra charge on some of its server offerings and also includes the first 20 GB of storage for free.

- **RackSpace** uses a different billing mechanism and reduces per-gigabyte bandwidth charges as volume increases; the company also offers backup services at no charge on some of its virtual server bundles.

- **Amazon** offers sliding scales for rates on most of its services but charges for gets and puts to its storage services.

When you look at the simpler concept of cloud-based storage services, the differences in pricing structures and in how providers build bundles for

Cloud Economies of Scale and Management

Keeping Control Isn’t Easy

Cloud Pricing Explained

The Future of IaaS

TABLE 1: Basic Cloud Server Pricing for Virtual Servers

	VIRTUAL MACHINE INSTANCE	BANDWIDTH IN	BANDWIDTH OUT	BACKUP	SUPPORT
RackSpace Cloud 1,024 MB/40 GB	\$.06/hour or \$43.80/month	\$0.08/GB	\$0.22/GB	\$0.15/GB	Included
Amazon EC2— Small server Linux	\$325/year, plus \$0.03/hour	\$0.10/GB	\$0.17/GB	\$0.15/GB/ month	\$0.015/ instance/ hour
GoGrid 1,024 MB/ 60 GB	\$0.19/hour	Free	\$0.50/GB	\$0.15/GB	Included

SOURCE: FRANK OHLHORST

customer needs becomes more evident (see **TABLE 2**).

Once again, the pricing comparisons are only part of the picture. And like other cloud services, prices are affected by length of commitment, total bandwidth needs or total size of storage required. And as with other cloud services, there is plenty of room for negotiation.

And once you take a closer look, the differences between service providers take shape.

- **GoGrid** offers the first 20 GB of storage at no charge to customers using its hosted server services and offers its storage services only to hosted server customers. Typical of the market segment, GoGrid offers discounted per-gigabyte pricing as

customers buy larger storage amounts. GoGrid's storage service is offered as a mountable volume and does not yet offer a Web services application programming interface for gets and puts and other commands.

- **RackSpace** tries to keep its storage services billing model as simple as possible. The company offers a sliding scale for storage purchases, where the price per-gigabyte decreases as total amounts ordered increases. What's more, if the file is more than 250 KB in size, the company doesn't charge for gets and puts. Amazon S3 does not charge for deletes and offers discounts as transaction volume and storage needs increase. The company also offers contract-based fixed pricing for those looking to stabilize their

Cloud Economies of Scale and Management

Keeping Control Isn't Easy

Cloud Pricing Explained

The Future of IaaS

TABLE 2: Basic Cloud Storage Pricing

	PRICE PER GB OF STORAGE	BANDWIDTH IN	BANDWIDTH OUT	PUT/POST/LIST REQUESTS	HEAD/GET DELETE REQUESTS
RackSpace Cloud Files	\$0.15/GB per month Unlimited files	\$0.08/GB per month	\$0.22/GB per month	\$0.01 per 500 requests	No charge
Amazon Simple Storage Services (S3)	\$0.15 per GB—first 50TB/month	\$0.10 /GB	\$0.17/GB per 1,000	\$0.012 per 1,000	\$0.012
GoGrid Cloud Storage	\$0.15/GB	N/C	N/C	N/C	N/C

SOURCE: FRANK OHLHORST

prices. For larger file transfers, the company recommends using its import and export services, which can reduce costs.

Ideally, pricing out cloud-based services should take little more than selecting features and determining storage and server computing needed. In reality, most IT managers will find that ideals don't apply here and will have to carefully consider the "hidden extras" as well as the standard charges to determine the cost for services. The bigger challenge is converting the techspeak of cloud services pricing into something that those holding the purse strings can understand and to bring those individuals into the loop to negotiate pricing and contract terms as well as SLAs.

The trick is to get all estimates in writing with clear, concise language that explains what the total charges will be per contracted period: Written confirmation is the only way to see an apples-to-apples comparison of pricing models and determine true budget amounts.

THE FUTURE OF IAAS

A recent study by the University of New South Wales in Australia, in collaboration with researchers at NICTA (National ICT Australia) and the Smart Services Cooperative Research

Centre, spent seven months stress-testing Amazon's EC2, Google's App Engine and Microsoft's Azure cloud computing services.

The analysis simulated 2,000 concurrent users connecting to services from each of the three providers, with researchers measuring response times and other performance metrics. The results were less than stellar, with response times varying widely depending on the time of day in which services were accessed. In addition, the study revealed a lack of monitoring tools to enable organizations to check whether the service has met their SLA.

It's obvious that cloud providers have their work cut out in terms of simplifying pricing models, beefing up security and providing SLAs that guarantee better reliability. The market is evolving fast, and today's dominant players could be history in two years. Just one catastrophic data privacy breach could extinguish a brand and potentially slow the market down for a decade. Hopefully, cloud providers are listening, and their customers will be willing to keep pushing for what they need. ■

Jo Maitland is an executive editor in the Data Center and Virtualization media group at TechTarget. This article features additional reporting by **Christina Torode**, a senior news writer in the CIO Media Group and TechTarget.

Editors' note: This chapter on Infrastructure as a Service is the fourth part of an e-book on cloud computing that also includes chapters on CIO strategies for the cloud, development for the cloud and Software as a Service.

[Cloud Economies of Scale and Management](#)

[Keeping Control Isn't Easy](#)

[Cloud Pricing Explained](#)

[The Future of IaaS](#)



- ▶ **Contact us today for a no-obligation 30-day trial on any of our range of cloud computing solutions. That's CHOICE, one of our 5 levels of superiority.**
- ▶ **Alleviate your migration costs with a no-obligation 30-day cloud solution trial. That's VALUE, one of our 5 levels of superiority.**
- ▶ **Envision your company's cloud computing future with a no-obligation Virtualization Assessment. That's VISION, one of our 5 levels of superiority.**

Why do infrastructure managers at leading edge Web 2.0 start ups, successful mid-sized enterprises and some of the world's largest consultancy and integration firms trust their business to Layered Tech? So they can move forward fast, with affordable, reliable and flexible hosted infrastructure in our secure, top-tier data centers. We help them cut capital expenses and operating costs, freeing up internal resources for their core business. And they know we are leaders who are always a step ahead in cloud computing. So when their business needs the ability to scale rapidly, they know we can provide on demand capacity, and the calibrated level of support they require. We are Layered Tech. Ask us more about our five levels of superiority today: Sales@layeredtech.com or 1-866-584-6784.