

Infrastructure-as-a-Service

Lowering IT costs is just one of many benefits driving enterprises to IaaS.

Tim Sexton

2015

Executive Summary

Moving an entire infrastructure (or even a portion of it) to the cloud can yield significant benefits. Not only can an enterprise get the computing power, storage, or other services they need when, where, and how they need them, but they pay only for the capacity used. What's more, they get it in a secure, monitored, metered environment that promises 100 percent uptime.

During the past decade, large enterprises have focused on consolidating their infrastructure through server virtualization, reducing hundreds or even thousands of physical servers by a factor of ten. For many organizations, moving some or their entire infrastructure to the cloud is the next logical step.



Table of Contents

2	Infrastructure-as-a-Service
4	Benefits for Enterprises
5	Choosing the Right Solution
6	When Infrastructure-as-a-Service Makes Sense
7	Preparing for Migration
9	Bonus: 7 Criteria for Choosing an IaaS Provider

Infrastructure-as-a-Service

Today, many are making the choice to move to some form of Infrastructure-as-a-Service (IaaS) for help in managing data center computing resources. In a 2012 study, Enterprise Strategy Group found that 30 percent of enterprises (defined as organizations with 1,000 or more employees) currently use some form of IaaS, compared to 19 percent in 2011. The study also found that an additional 55 percent are considering or plan to use IaaS in the future, compared to 44 percent in 2011.

The top four reasons for moving to IaaS include its availability as a resource for testing and development, the ability to run production applications, the availability of additional resources to accommodate spikes in workload demands, and for use as temporary computer resource for time-limited projects.

IaaS Overview

At its core, IaaS is a way for organizations to get the hardware, storage, networking, and other services they need to run their operations without worrying about buying, managing, or maintaining the equipment. With this model, enterprises “rent” the equipment, paying only for the capacity and space they use.

An offsite provider owns, manages, and maintains the equipment, providing enough capacity to scale up or down to meet the needs of its customer organizations. Organizations access the resources (virtual machines that run their applications) via a web browser.

There are many setups for IaaS, from full-blown compute-as-a-service to partial infrastructure-as-a-service, such as storage, servers, web hosting, and disaster recovery. The amount of infrastructure an organization chooses to move to the cloud depends on many factors, including its line of business, the availability of IT staff and IT expertise, and cash flow constraints.

There are also different ways of implementing IaaS: in a private cloud, public cloud, or hybrid private/public cloud. Again, the path an organization chooses depends on several factors and usually comes down to balancing cost with security needs.

Although the features of different IaaS offerings and vendors vary, there are several that are fairly standard, including the following:

- **Dynamic infrastructure scaling:** The service provider is responsible for ensuring the customer’s infrastructure needs will always be met, even if they include peaks and valleys. With this infrastructure model, organizations can always be assured that will have the IT resources they need.
- **Guaranteed uptime:** IaaS providers guarantee high availability and 100 percent uptime. This is spelled out in the service-level agreement (SLA).

- **Automation of administrative tasks:** These tasks include deploying and managing virtual servers, managing virtual server pools, and monitoring ongoing activity such as disk utilization, network activity, active jobs, and member and guest activity.
- **Policy-based services:** This feature imposes enterprise-defined policies on the infrastructure services that specific user groups can access. Based on these policies, provisioning, and decommissioning resources becomes an automated process.
- **Elastic load balancing:** Based on requirements at any given time, resources are balanced and distributed automatically, appropriately, and efficiently across multiple virtual computers.
- **Customized machine images:** With the IaaS model, building and deploying machine images on which applications can run is fast, automated, and customized to an organization's specific needs. IaaS providers offer a host of virtual machine images to their customers, but organizations themselves can also develop their own or import images from the existing operational environment. That way, they know the images will meet their configuration, security, and compliance standards. It also vastly increases the speed at which new instances of the image can be deployed.
- **Preconfigured templates:** Qualified users can create predefined templates for various types of virtual machines. This allows other users to choose from a preapproved list of templates, standardizing and removing risk from the process of quickly spinning up virtual machines.
- **Access to the newest technology:** Because organizations don't own the equipment, they aren't stuck with technology that will eventually become outdated and need an upgrade. Instead, they will always have the most up-to-date features and the highest levels of efficiency.
- **Stringent security controls:** Depending on the solution and the vendor, an IaaS offering will be compliant with SSL (Secure Sockets Layer) or AES (Advanced Encryption Standard) encryption, employ a virtual private network (VPN) and provide role-based access control.

Taken together or in bunches, IaaS features provide the ability to scale up and down quickly while paying only for the capacity actually used, which is important in today's fast-paced, competitive environment. For example, to stay on top of the competition, an enterprise may want to create batch processes that run every night to better understand its sales data.

But it may need ad hoc computing resources to do so. An IaaS agreement would give the organization the capacity it needs for nightly intelligence gathering without laying out significant budget.



Benefits for Enterprises

With IaaS, IT shops don't have to buy, maintain, or upgrade software, hardware or operating systems. They don't have to worry about network configurations. They can save big money, both up front and over time.

IaaS, like all cloud offerings, is a pay-as-you-go model. So organizations don't have to manage large cash outlays. Instead, they pay only for the resources they use, which is more cost-effective than the traditional method of paying set fees for services and equipment, even when they aren't being used.

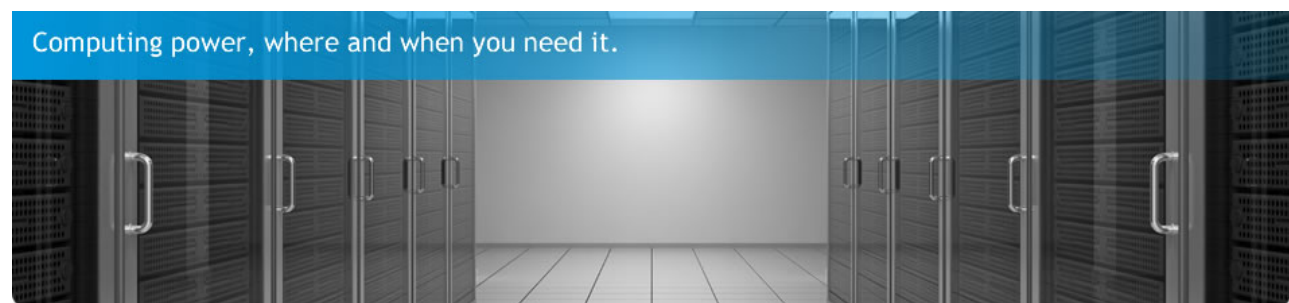
Either way, the IaaS model allows enterprises to meet capacity spikes without having to add infrastructure within the data center, and to reduce capacity when the spike ends.

The IaaS model also takes a lot of pressure off of IT staff, who generally are responsible not only for the hardware and the infrastructure, but also for handling requests from line-of-business owners, application developers, and engineers. In a traditional IT environment, IT staff often does not have the resources to respond to all such requests in a timely manner, forcing divisions to purchase externally.

With the IaaS model, requests can be filtered quickly through the IT department, maintaining internal policies and procedures and then sent to an IaaS vendor approved by the organization, recommends Ken Nunes, HopOne Chief Technology Officer.

IaaS is compatible with today's anytime, anywhere work model. Because everything is IaaS are web-based, administrators can manage the systems from anywhere, whether they are adding capacity, removing users, or accessing reports. Users can get the computing resources they need from wherever they happened to be, significantly improving the agility of the organization as a whole.

Finally, IaaS can solve the issue of inevitable system upgrades. Technology changes rapidly, and those changes usually bring higher capacities, greater reliability, and useful new features. Purchasing upgrades can be expensive, and organizations that rely on internal IT infrastructure must eventually upgrade their equipment (even virtualized servers), either to increase capacity or features or because a system has failed or is no longer supported by a vendor. With IaaS that never happens because the equipment isn't the organization's problem – it's the IaaS provider's responsibility.



IaaS VS. PaaS

If infrastructure as a Service is a way to provide virtual machines, servers, storage, and other infrastructure services via the cloud, what is Platform as a Service (PaaS)? There has long been confusion between the two, and with good reason.

Both offer computing via the cloud, but at different layers. The difference is this: The PaaS model provides both a computing platform, consisting of application building blocks such as databases and file sharing, along with applications that run on the platform, while IaaS provides the infrastructure building blocks only – storage, hardware, and networking components.

Clearly, PaaS and IaaS are closely related, because applications, platforms, and infrastructure can't work without each other. This has led to speculation that the two will eventually become one layer of service. By adopting some cloud infrastructure earlier rather than later, organizations will be better prepared.

Choosing the Right Solution

The term infrastructure is broad in meaning, so it's not surprising that subcategories of cloud-based services fall under the IaaS moniker. Some of the most popular IaaS solutions include the following:

- **Compute-as-a-Service:** One of the most ubiquitous IaaS offerings today, compute –as-a-service provides compute capacity that includes servers, operating system access, firewalls, routers, and load balancing on demand. These systems have management interfaces, and their capacity should be either shared or private.
- **Web hosting:** Many organizations rely on their websites for marketing and revenue, and any glitch in operations can mean a loss of business. Moving a website to an IaaS-based model ensures that the website won't get bogged down during peak traffic times – and that organizations won't have to overpay for capacity to manage those traffic spikes. What's more, loads will always be balanced, and uptime is guaranteed, thanks to SLAs. Other perks include offsite backup and fast connections for eliminating slow page and content downloads, no matter how much rich media a site contains.
- **Storage-as-a-Service:** Storage is one of those necessities that only grow over time. It can be a constant struggle to maintain enough storage capacity and manage it effectively. Storage-as-a-service goes a long way towards easing that burden. These solutions have interactive self-service portals that allow administrators to provision storage, transfer data to different tiers of storage, dispatch specific data sets to different media, and add or remove storage as needed. Tiers generally include fast storage for high I/O applications, standard storage for system disk and bulk storage for file serving. Cloud-based storage has the security controls to ensure that all data is stored securely in data center facilities, with extremely high availability.
- **Disaster recovery and backup as a service:** The idea behind moving disaster recovery to the cloud is to ensure that organizations have uninterrupted access to data and applications, regardless of emergencies, such as power outages, natural disasters, or system failures. These

solutions always include redundancy and automatic failover to ensure ongoing access, reducing downtime to nearly zero. Many solutions also employ continuous data protection (CDP), which allows for multiple versions of all data sets to be recovered.

- Servers-as-a-Service: Accessing servers in the cloud means that no matter what the project, or even if it's the busy season, there will always be enough compute power to go around. It's useful for one-time projects that require additional capacity, or for handling spikes in transactions. And because it's a service, enterprises can rest assured that they'll never be paying for more server capacity than they need.
- Networking-as-a-Service: This is the newest entrant in the IaaS category. The idea is to offer networking resources on demand in order to support virtual networks – resources such as firewalls, load balancing, and WAN acceleration services. Simply put, NaaS provides unified connectivity across storage, networking, and servers that changes to meet the demands of virtualized infrastructures.



When IaaS Makes Sense

All enterprises are looking for ways to cut costs, and the cloud provides an opportunity. The pay-as-you-go model means that organizations never pay for capacity not used. Plus, there are no more costs for upgrading hardware and other infrastructure, no more service contracts, and few if any virtualization licensing costs.

Organizations see cloud computing as a cost reduction strategy for many reasons, according to a study by Enterprise Strategy Group: There is no need to renegotiate contracts, no need to postpone projects for lack of capacity, and no need to dedicate staff to managing infrastructure.

Organizational models that require the capacity to quickly react to market changes are also good candidates for IaaS. Such organizations may need to stand up a system on short notice, create or troubleshoot a product quickly or get new employees up to speed in hours rather than days.

That might be possible with an in-house infrastructure, but not without some fast thinking around repurposing and repositioning equipment. With infrastructure in the cloud, the capacity is always there when needed.

It makes sense to consider IaaS when the data center is due for a complete refresh or overhaul. That's the time to consider if it makes sense for the organization to continue managing its own data center. Is that the best use of your salaried employees? If an organization is not in the IT business, chances are good that it may not be in its best interest to spend the money on new infrastructure.

Saving money is a popular reason for moving to IaaS – or any type of cloud computing. But depending on the situation, that may not always be the case. It's worth taking the time to do a complete cost/benefit analysis. That means considering every angle – the money an enterprise would save by not buying, maintaining, and upgrading its own equipment, as well as the costs of labor the organization is now paying to manage its internal infrastructure.

Those are the hard costs, and they are fairly easy to quantify. The bigger issue is determining the soft costs. For example, if an organization is anticipating significant growth but can't be sure how much computing capacity it will need, it's difficult to run the numbers.

However just knowing that a major growth stage is on the horizon is good information for a total-cost-of-ownership (TCO) calculation. Because IaaS can be scaled up or down quickly, it's not as important to know how much the organization will grow as it is to know that it will grow significantly.

But perhaps the most difficult costs to determine are the "missed opportunity" costs. In other words, if an enterprise can operate faster and better using an external IaaS provider, it may gain some type of advantage that it currently can't imagine.



Preparing for Migration

Although making the leap from a physical infrastructure to an IaaS environment may seem overwhelming, there is a logical way to structure the move. The first step is to virtualize the current environment if the enterprise has not done so already (generally, servers first, followed by applications, then desktops).

By virtualizing as a first step, data and applications wind up separated from the physical machines on which they reside. This gives IT managers and executives a clearer understanding of which applications

and data are most important, along with the interdependencies among them. Through this process, it becomes clear which applications make sense to port to the IaaS environment, which is itself based on a virtual infrastructure.

At the same time, the virtualization process begins the cultural shift that will lead workers to access applications and data through virtual machines. And it will train IT staff to deal with services-based, shared-resources mindset – the same mindset and processes used in all cloud deployments.

The next step is to discover which apps and software versions are in use throughout the environment, which users have permission to use them, and the interdependencies among those applications. Once that information has been gathered, it's important to collect and measure data about each application's resource utilization levels – how much CPU power and memory does each app use?

IT shops will also need network and storage metrics, specifically as they describe throughput and latency. This data is crucial to determining how the applications will work in the new IaaS environment.

Application dependency mapping is next. This involved identifying dependencies between applications, as well as between applications and the data infrastructure. The idea is to find out how to architect the migration plan so that those dependencies aren't disrupted. It's a painstaking process, one that an outside consultant or a special-purpose mapping tool can help with.

It's also important to ensure that the cloud-based infrastructure is compatible with the organization's existing server hardware and operating systems. If the hardware isn't compatible, applications may have to be redeployed or recompiled for the new platform. If the operating systems are different, other changes may have to be made when the enterprise migrate its systems.

Finally, it's important to prepare the IT staff for its changing role. Although it might be tempting to think that the IT function will no longer be needed once an organization moves infrastructure to the cloud, that's not the case.

Instead, the IT staff's role will change from one of a hands-on, troubleshooting, in-the-weeds role to one that is more managerial. The IT manager now becomes a liaison to the IaaS provider, monitoring its work and making sure the organization gets its best value.


Infrastructure-as-a-Service makes a great deal of sense for many organizations. Agility and productivity are keys to surviving and thriving in a challenging economy. So anything that conserves cash and allows the enterprise to focus on its core business or mission is worth considering.

Choosing among the many options available today may seem an overwhelming task. But by asking the right questions, talking to the right vendors, and consulting with trusted advisors, enterprises are much more likely to make the right decision – a decision that will propel the organization forward securely, quickly, and profitably.

Choosing an IaaS Provider: Seven Criteria ... next page

Seven Criteria for Choosing an IaaS Provider

- I. **Robust Infrastructure** – Look for a fully managed IaaS cloud solution running in an enterprise-class data center with secure, reliable cloud services to support and deliver hosted software solutions.
- II. **Revenue Opportunities** – Ask what opportunities exist to generate additional revenue with services that create annuity income, including disaster recovery, application monitoring, and service desk offerings.
- III. **Best Practices** – Accept only a best-practices approach based on ITIL v3 standards; the provider should have 24x7x365 monitoring and management in place, as well as highly automated processes with a wide range of sophisticated software tools, all combined with real-world experience with software companies to deliver hosted solutions.
- IV. **Flexible Capacity** – Explore the cloud solution provider’s ability to deliver flexible capacity to accommodate fluctuations in customers’ production workloads, as well as on-demand capacity for training, demonstration, proof-of-concept or test/development.
- V. **Customizability** – Look for a consultative approach to architecting a cloud solution to the software solution’s requirements rather than just learning another cookie-cutter cloud offering. Alternatives such as legacy operating systems, hybrid solutions, and managed appliances should be offered.
- VI. **Serious SLAs** – Examine the provider’s service-level agreements. There should be a designated service-level manager overseeing stringent SLAs backed by stiff financial penalties to ensure target resolution times – not just target response times – are met.
- VII. **Industry Recognized** – Choose a provider with a proven, industry-recognized solution that has earned certifications and awards and that fares well in comparative rankings.



The Foundation of Internet Success®

Tim Sexton
Sales Manager
DCA3 • DCA2

DCA3	DCA2	SEA2
6621 Electronic Dr. # E Springfield, VA 22151 T: 571.282.6400 F: 571.282.6402	8201 Greensboro Dr. # LL McLean, VA 22102 T: 703.564.5400 F: 703.564.5402	3311 S. 120th Place # 150 Tukwila, WA 98168 T: 206.438.5900 F: 206.438.5902

Direct : 571.282.6321
Mobile: 571.426.9824
Fax : 202.318.0792
tsexton@hopone.net

www.hopone.net

