

The Business Case for Cloud

An Esri® White Paper
April 2013



380 New York Street
Redlands, California 92373-8100 USA

909 793 2853
info@esri.com

esri.com

Copyright © 2013 Esri
All rights reserved.
Printed in the United States of America.

The information contained in this document is the exclusive property of Esri. This work is protected under United States copyright law and other international copyright treaties and conventions. No part of this work may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or by any information storage or retrieval system, except as expressly permitted in writing by Esri. All requests should be sent to Attention: Contracts and Legal Services Manager, Esri, 380 New York Street, Redlands, CA 92373-8100 USA.

The information contained in this document is subject to change without notice.

Esri, the Esri globe logo, ArcGIS.com, ArcGIS, esri.com, and @esri.com are trademarks, service marks, or registered marks of Esri in the United States, the European Community, or certain other jurisdictions. Other companies and products or services mentioned herein may be trademarks, service marks, or registered marks of their respective mark owners.

The Business Case for Cloud

An Esri White Paper

Contents	Page
Changes to the Computing Model	2
Procurement	2
Business Model	2
Accessibility	2
Technical	2
A Word on Virtual Servers	2
Background Information	2
Business Benefits	3
Service Models	3
Deployment Models	3
Considerations with Private and Hybrid Cloud	5
Conclusion	6

The Business Case for Cloud

"The next sea change is upon us. . . . The broad and rich foundation of the Internet will unleash a 'services wave' of applications and experiences available instantly. . . . This [new wave] will be very disruptive."

Bill Gates, to top managers and engineers¹

The expansive computing grid that we know as cloud computing has become the newly poured foundation for how businesses will deliver and consume products—as services. As new civilizations have built on top of the structures (and ruins) of older civilizations, cloud infrastructure has been laid over and built up from a rich history of previous computing substrates. Scratch below the surface of cloud computing and you'll find the era of the corporate data center, which required a cadre of technology specialists to provide the ongoing maintenance of the system. These entities evolved into storing single-purpose systems, requiring dedicated hardware and an application stack configured for maximum load, resulting in underutilized processing capacity and overweight electric consumption. Dig a little deeper and you'll uncover the mainframe, once housing the notable UNIVAC (for the few who could afford such mechanical behemoths). The mainframe transformed computing into modern-day engine rooms, to run the business for all types of organizations. The cloud takes these artifacts and taps into the ubiquitous, affordable, on-demand Internet connectivity that drives nearly every technological advance nowadays.

The concept of cloud computing, or on-demand technology, emerged with new workflow paradigms attributable to the arrival of Web 2.0. Cloud computing is defined by research and advisory company Gartner, Inc., as "a style of computing where massively scalable IT-enabled capabilities are delivered as a service to external customers using Internet technologies."² Research company Forrester defines the cloud as "a standardized, multitenant IT capability delivered via Internet technologies in a pay-per-use and self-service way."³

There are several permutations of the cloud computing definition. The preferred de facto definition of cloud computing includes the basic tenets that characterize cloud computing, as defined by federal technology agency National Institute of Standards and Technology (NIST) in elaborate detail.⁴ Characteristics include on-demand, commonly off-premises technology capabilities that are delivered as a service via the Internet. Consumers do not own the assets in the cloud computing model but consume them in a pay-per-use and self-service manner, in essence renting physical assets based on an organization's computing requirements and leveraging economies of scale to reduce capital and operational expense. Most commonly, these cloud services are owned and managed by a

¹ Gates, Bill, "Internet Software Services" memorandum, 30 October 2005, <http://scripting.com/disruption/mail.html>.

² Gartner, Inc., IT Glossary, "Cloud Computing," <http://www.gartner.com/it-glossary/cloud-computing>.

³ Forrester Research, *Is Cloud Computing Ready for the Enterprise?*, March 2008.

⁴ Mell, Peter, and Timothy Grance, *The NIST Definition of Cloud Computing*, NIST Special Publication 800-145, September 2011, 6–7, <http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf>.

third party, and offerings can range from data storage to web applications to other focused computing services.

Changes to the Computing Model

There are four key differences that distinguish cloud computing from traditional computing models:

Procurement

Traditionally, individuals and organizations buy assets (computer and network components) and build the technical architecture to accommodate the hardware and software needed by staff and customers. In the cloud model, it is a service, commonly off premises, that is consumed. In essence, users rent these assets, leveraging the cloud provider's infrastructure or applications instead of buying their own. However, since they are using a dedicated cloud hosting provider, they also gain access to state-of-the-art systems without the capital investment. It is in the cloud provider's interest of continuing business success to provide high reliability, quick response times, and the flexibility to handle traffic fluctuations. This in turn allows cloud consumer organizations to reassess use of a previously fragmented on-premises infrastructure, which could subsequently drive down capital expenditures on equipment to be freed up to invest in core business and other uses.

Business Model

Unlike the traditional model, where one pays for fixed assets, overhead, and administration, organizations typically do not own assets in the cloud computing model and instead pay based only on the use of the service.

Accessibility

Deviating from traditional computing, cloud computing services are available only through standard Internet protocols—through any device that supports web access, whether that's mobile, tablet, desktop, server, laptop, or smartphone.

Technical

Instead of a single-use, dedicated, static system with underutilized capacity built for a maximum load (which indeed may never realistically happen), cloud computing supports dynamically scalable, "elastic" systems. The ability to dynamically scale up and/or down improves IT's ability to rapidly provision its systems based on traffic and demand. In other words, system capacity can be increased or reduced as needed. Another critical difference between traditional and cloud computing is that the latter supports multitenancy—where systems are configured such that they can be shared by many companies, business units, or individuals. Virtualization technology allows cloud providers to convert one server into many virtual machines, thereby eliminating client/server computing with single-purpose systems, maximizing hardware capacity, and allowing customers to leverage economies of scale. Virtualization is a common and integral part of cost optimization strategies for private and public cloud architectures.

A Word on Virtual Servers

Background Information

Virtual server technology has been traditionally used to reduce testing and development costs by allowing a user to create multiple "virtual" servers on a single physical server. A virtual server mimics the behavior and capabilities of a separate stand-alone physical computer. This gives the user the ability to test software applications running on multiple operating systems (OS), with each OS running on a different virtual server but all virtual servers running on a single physical server. Until recently, each of the virtual server environments was logically isolated from the others, but in reality, they physically shared the same hardware resources (including all available sockets and CPU cores of the physical server). In this regard, virtual servers were different from "logical" servers that segregate hardware resources.

Virtual server technology can also be used to run different applications, each of which is compatible with different OS versions, on the same server. For example, a single physical server may have one virtual server instance running Windows Server 2012 and its supporting applications, while another virtual server instance could be running Red Hat Linux with another enterprise application suite; in both instances, all hardware resources of the physical server host are shared among the virtual server instances.

Business Benefits

Server virtualization is getting a lot of attention because of the business benefits gained when organizations move to a virtualized environment, primarily optimization of their hardware assets. IT departments realize an increase in productivity through faster provisioning times for new servers, which in turn results in faster availability of resources for users. Virtualized servers can take advantage of the same server network tools that logical servers employ for business continuity and disaster recovery. Tremendous cost savings can also be realized through reducing and eliminating the need for new hardware and better utilization of the existing server hardware. Less hardware translates into lower energy costs and a reduction in required floor/rack space. Besides being beneficial for the environment, this translates into improving the external image of a company as being socially responsible and eco-friendly.

Service Models

Although there are many ways to take advantage of the services that are hosted in a cloud computing environment, there are at least three core cloud service delivery models:

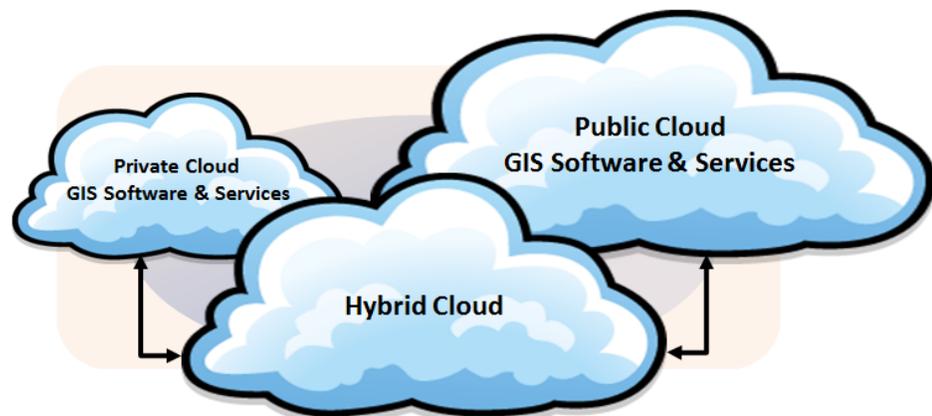
Software as a service (SaaS) typically refers to end-user applications that are delivered on demand, as a service, for example, Salesforce.com and ArcGIS.com. Since its first introduction back in the late 1990s, SaaS has grown considerably, moving into mainstream business solutions and ubiquity in technology portfolios.

Platform as a service (PaaS) provides an application platform or middleware as a service on which developers can build and deploy custom applications. Common solutions provided in this tier range from APIs, database services, and security to integrated development environment (IDE) integration and identity, allowing developers to build applications and run them on the infrastructure typically owned and maintained by a cloud provider. Examples include Microsoft's Windows Azure and Esri's ArcGIS.com APIs.

Infrastructure as a service (IaaS) primarily encompasses those offerings that comprise the server-based hardware and technology that business processes would require for computing power, storage, or other infrastructure as a service. IaaS is also referred to as *compute as a service* or *storage as a service*. Amazon Elastic Compute Cloud (Amazon EC2) and ArcGIS® for Server on Amazon Web Services are examples.

Deployment Models

When cloud computing came to the forefront in 2008, conversation centered on the public cloud model, which leverages applications and infrastructure access on demand and is typically hosted off premises by a trusted cloud provider. Given the benefits of economies of scale, the public cloud can often be an economical option for organizations that aim to reduce their data management and application distribution costs. The public cloud is the most commonly referenced and assumed platform when discussing the topic of cloud computing.



However, many organizations and users of the traditional compute model are still not quite ready to jump into public cloud computing or are restricted from doing so by company policies or legal mandates. These organizations may be particularly concerned about their security, privacy, permissions, firewalls, authentication, and data protection requirements. In these situations, the core service tiers can be replicated within a private cloud environment, behind a firewall, and maintained within the parameters of the host organization. A private cloud model, according to NIST, is operated solely for an organization, can be managed by the organization or a third party, and may exist on or off premises.⁵

Adopting and building for a private cloud may make it possible for the company to meet its business requirements for additional security, enterprise integration, and specific customization, which a public cloud option would not entirely expose or allow.

Increasingly, however, exploiting a hybrid cloud deployment model has emerged as an attractive compromise between these two very different choices. No longer is it a decision involving either public or private cloud adoption. Organizations can take advantage of both models, balancing more or less services from either source based on the business rules of the organization. For example, if a public cloud service currently exists that provides a needed function for an organization, the company could run that public service against its data or information, with results that are company specific. A task or job could be initialized against a public cloud service, and the result of that job leveraged solely by the organization. In this manner, the public cloud-hosted application and infrastructure would be consumed only for a project-based time frame, without additional capital expense incurred by permanent procurement of hardware or software. In this hybrid model, cloud-based applications integrate with enterprise-based software infrastructure.

Another variation of an organization leveraging a hybrid cloud is the decision to host certain noncritical functions of an organization as services in a public cloud while at the same time keeping applications and services that are considered sensitive or proprietary on premises within a private cloud. Further, if the on-premises applications reach a

⁵ Mell and Grance, *The NIST Definition of Cloud Computing*, 3.

certain maximum load, leveraging a hybrid cloud would allow certain tasks to be drafted to the public cloud services—commonly referred to as "cloudbursting."

According to Gartner, through 2014, IT organizations will spend more money on private cloud computing investments than on offerings from public cloud providers.⁶ Additionally, a Gartner data center conference poll in November 2012 indicated that more than 50 percent of identity and access management deployments will be hybrid.⁷

Considerations with Private and Hybrid Cloud

The adoption of a private cloud has certain characteristics that set it apart from a public cloud option. Specifically, several of the benefits that are drivers to the public cloud are diluted or eliminated when a private cloud is built and/or adopted.

By example, if an organization's private cloud exists on premises, the maintenance and overhead costs will continue as with any previous centralized computing data center. In fact, system and architecture administration will need to include service catalogs, virtualization and virtualization management, self-service administration, and possible metering or implementation of tools and applications for measuring usage. Applications must ensure scalability based on demand and be selected judiciously based on usage, automation, and other key criteria for the organization or business. Further, applications and middleware must ensure the latest security patches, upgrades, or updates. And an infrastructure management team is still required, to perform backups, apply patches, and so on. As such, operational expenses must continue to be budgeted for private cloud success.

Likewise, capital costs are not eliminated with private cloud implementations. Physical infrastructure will need to be maintained and kept to the level of sophistication that meets business requirements for traffic response and storage and task load. Initially, the up-front costs and investments in a private cloud may be heavy—front-loaded.

If an organization decides to build and maintain a private cloud, particularly on premises, both capital and operational expenses must be factored into the return on investment (ROI) calculations. Economies of scale that may have been realized through massively shared infrastructure in an established public cloud provider's data center cannot be leveraged. Further, operating a private cloud brokerage becomes part of the business, whether or not it is a core competency of the company hosting the private cloud.

However, with services that are managed on premises in a private cloud, organizations know exactly where the host and the data are located, meeting organizational or legal requirements. Compliance concerns are evident and include business continuity/disaster recovery, logs and audit trails, and specific compliance requirements (Sarbanes-Oxley, HIPAA, PCI). The typical vendor lock-in that can be experienced in a public cloud are eliminated with an on-premises private cloud but are replaced by other characteristics, as previously outlined.

⁶ Gartner, Inc., "Five Cloud Computing Trends That Will Affect Your Cloud Strategy through 2015," 10 February 2012, Document ID G00230221.

⁷ Gartner, Inc., "Predicts 2013: Cloud and Services Security," 28 November 2012, Document ID G00245775.

Indeed, although the onus is on the organization to scale its private cloud and continually reinvest in the on-demand IT architecture and services, given judicious selection of services, the opportunity for increased business agility is unmistakable. In addition, leveraging public cloud services in a hybrid architecture can mitigate some of the capital spending of a pure on-premises private cloud system.

Additionally, multigenerational work forces add an additional challenge to successful business operations and workflows in the twenty-first century. The demographics of organizational staff can span multiple age groups, each with its own expectation of effective worker productivity tools. The "digital natives," for example, having grown up with technology, collaborate extensively and approach real-time information at a different pace than "digital immigrants." Team productivity may comprise multiple people in different geographies, on different operating systems, working with different languages, commonly using on-demand applications. As such, cloud services that are available 24/7, accessed from any browser on any device, regardless of time or time zone, can provide faster, easier access for these "native" workers to do their jobs, allowing competitive differentiation for the organization and, likewise, the retention and attraction of valuable and talented staff.

Conclusion

With the potential for unlimited scalability, availability, and reliability that comes from taking advantage of cloud services, opportunities emerge. Since servers can be launched when needed for predicted content load, the downtime of the traditional on-premises infrastructure process is avoided, allowing faster deployment times and variable capacity.

As business challenges evolve, on-demand services are meeting these tasks, with many potential advantages to the cloud consumer and the business organization. The disruptive power of cloud computing is evident in the growing trend of organizations migrating their IT infrastructures to private clouds as well as taking advantage of public cloud infrastructures. Hybrid clouds allow organizations to get the best of both worlds while accommodating diverse requirements. Hybrid clouds may also enable a smoother transition into the utility model offered in public clouds. Whether the hybrid cloud model is the solution or simply a transition model, organizations are finding more incentives to move some of or all their business into the cloud.



Esri inspires and enables people to positively impact their future through a deeper, geographic understanding of the changing world around them.

Governments, industry leaders, academics, and nongovernmental organizations trust us to connect them with the analytic knowledge they need to make the critical decisions that shape the planet. For more than 40 years, Esri has cultivated collaborative relationships with partners who share our commitment to solving earth's most pressing challenges with geographic expertise and rational resolve. Today, we believe that geography is at the heart of a more resilient and sustainable future. Creating responsible products and solutions drives our passion for improving quality of life everywhere.



Contact Esri

380 New York Street
Redlands, California 92373-8100 USA

1 800 447 9778
T 909 793 2853
F 909 793 5953
info@esri.com
esri.com

Offices worldwide
esri.com/locations