



THE RULES OF CLOUD SOURCING

An overview of the cloud service provider
selection criteria and contracting options

nuvalo[™]



THE RULES OF CLOUD SOURCING

An overview of the cloud service provider
selection criteria and contracting options

April 2014

Seattle, WA



The Rules of Cloud Sourcing

The information contained in this document represents the current view of Nuvalo on the issues discussed as of the date of publication. Because Nuvalo must respond to changing market conditions, it should not be interpreted to be a commitment on a part of Nuvalo, and Nuvalo cannot guarantee the accuracy of any information presented after the date of publication.

This whitepaper is for informational purposes only. NUVALO MAKES NO WARRANTIES EXPRESSED, IMPLIED, OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user, without limiting the rights under copyright, no part of this document may be reproduced, stored in, or introduced into a retrieval system, or transmitted in any form, or by any means (electronic, mechanical, photocopying, recording, or otherwise) or for any purpose, without the express written permission of Nuvalo.

Nuvalo may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided, in any written license agreement from Nuvalo, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

Copyright © 2014 Nuvalo. All rights reserved.

**MANON BUETTNER**

Hello,

The topic of selection criteria for IaaS inspires me because in my practice I always begin explaining to my prospective clients, as our journey starts, the nuances of selecting the most appropriate cloud service provider. Naturally there are brand name options, which come immediately to mind that offer a sense of security to companies, regardless of size. With a closer look I find that often these providers are selected not for their relevance and fit, but simply because they are safer brands. In fact, many of these large market-leading brands are going to be the least agile and most expensive options available.

This report drills down into key IaaS provider selection criteria so you understand the technical underpinnings and strategic implications that can affect your business. It is intended to instill the confidence in IT decision makers and purchasing managers so that they can use this document to quickly answer basic questions about IaaS selection criteria, such as:

- What are the various IaaS models and how does each model function within our current business structure?
- How important are certain IaaS selection criteria for our business workflows?
- What are the ways for my business to source cloud services?

I hope you enjoy this whitepaper, and please feel free to contact me should you like to discuss this topic further.

Best Regards,

Manon Buettner

About the author

Manon Buettner is an accomplished infrastructure architect, data center and cloud analyst and practitioner with over fifteen years of experience designing, selling, and implementing cloud solutions, data center selection and migrations, and network enablement for mid-market and enterprise businesses, as well as for telecom and service providers.

Ms. Buettner founded Nuvalo in January 2010, parlaying her vast expertise with service providers into the role of trusted advisor to IT Executives. Leading the Nuvalo team, she has worked with some of the most notable Internet, media and entertainment, healthcare, FMCG/retail, and financial services companies to successfully deploy infrastructure and enterprise business applications with service providers throughout the country and abroad.

An expert voice on cloud computing, virtual data center, and related solutions, she is uniquely positioned to offer guidance and resources to support virtually any IT initiative.

Ms. Buettner works closely with dynamic enterprise IT teams to identify business and technical requirements, introduce relevant solutions, and manages complex colocation or managed cloud services searches and deployments.

Manon Buettner is a contributing writer for Channel Partners magazine and is a popular featured speaker at data center and network trade events. She is the founding co-chair of Cloud Girls, a vendor-neutral community of female technology advocates focused on cloud education.

President	Manon Buettner
Research Analyst	Jo Peterson
Research Analyst	Michelle Hyde
Research Analyst	Liana Antanovich
Design	Eben Greene

NUVALO

800 Bellevue Way NE #400
Bellevue, WA 98004
Tel. 877-9NUVALO
(877-968-8256)
Email: info@nuvalo.com
www.nuvalo.com

Contents

About the author	5
Introduction	7
1 What is cloud computing	8
2 Cloud computing models.....	10
2.1 Public cloud.....	11
2.2 Private could.....	11
2.3 Hybrid cloud	12
2.4 What is moving to the cloud?.....	13
2.4.1 Cloud native applications	13
2.4.2 E-business hosting.....	14
2.4.3 General business applications	14
2.4.4 Enterprise applications	14
2.4.5 Test and development	14
2.4.6 Batch computing	14
2.5 Cloud in corporate business and IT strategy	14
2.5.1 IT policies.....	15
2.5.2 Application impact.....	15
2.5.3 Automation.....	15
3 Six selection criteria for a cloud service provider.....	17
3.1 Performance.....	17
3.2 Security	19
3.3 Support	20
3.4 Pricing.....	21
3.5 Tenure.....	23
3.6 Governance	24
4 Four cloud sourcing options	25
4.1 Work directly with the hundreds of providers	26
4.2 Utilize the expertise of a cloud services broker (CSP)	27
4.3 Engage a large, reputable diversified consulting firm	28
4.4 Leverage a boutique consulting and sourcing firm	29
Conclusion	30
About Nuvalo	32

Introduction

In this whitepaper we will address common concerns that my clients have expressed when embarking on a cloud project. You will be guided to seeing past the cloud for your enterprise architecture – a key vision you'll want to develop before jumping into your selection and budgeting process. You will learn the tips and tools for evaluating a cloud provider across six key areas – performance, security, support, pricing, tenure, and governance. And finally, we will **review different sourcing models for cloud services.**

In working with CIO's, VP's of IT, and hands on application and data center managers, we repeatedly hear the same concerns when sourcing cloud services. Perhaps you can relate...

- "We are finding the sheer number of provider choices overwhelming, and wondering where to start."
- "We are not sure if we are talking to the best providers for OUR needs. We have started out on a cloud project and now feel the sands shift under our feet."
- "How do we keep up with the ever-changing landscape?"

Maybe you received pricing from two or three providers with an extreme variance in price. For example, on a recent project, a client requested pricing for cloud storage. We provided the same specifications to three providers, whose responses ranged from \$12K to \$50K! What? Could a range like this really be? Finally, are you straight up concerned that you don't know what you don't know and that you will choose the "wrong" provider?

If you are you thinking of taking a hybrid approach to cloud services, this adds another layer of complexity beyond pure public or private cloud. Before starting conversations with prospective cloud service providers, it's important to understand the interdependencies and the utilization metrics associated with the applications or environments co-marked for migration and the business drivers for moving them in the first place. Only then you can begin the search for the right cloud provider, focusing on six key areas: **performance, security, support, pricing, tenure, and governance.**

You are welcome to share this whitepaper with anyone you think it would benefit. For more ideas on how to select your IaaS model, begin the search for the right cloud provider, and execute a successful cloud migration, you can [request a quote](#), visit Nuvalo.com for more resources, or [join Nuvalo's free newsletter](#).

1 What is cloud computing

Cloud computing is the evolution of IT service delivery from a remote location, either over the Internet or an intranet, involving multi-user environments enabled by virtualization. Cloud architecture is underpinned and enabled by common off the shelf (COTS) hardware, on top of which virtual machines are provisioned and offered to consumers. Flexibility is an important component of cloud computing, but it is both a feature and a natural result of the main drivers of cloud computing: economies of scale, improved efficiency and speed. In simple terms, cloud computing is offering IT capabilities as services.

The cloud model of computing has evolved over time from basic dial-up Internet access from a single workstation to racks of servers with specific applications obtained from Internet sources and used on a billable service agreement. Today we have a full cloud environment where everything (applications, infrastructure and storage) is now available on a pay-for-usage service model.

Internet

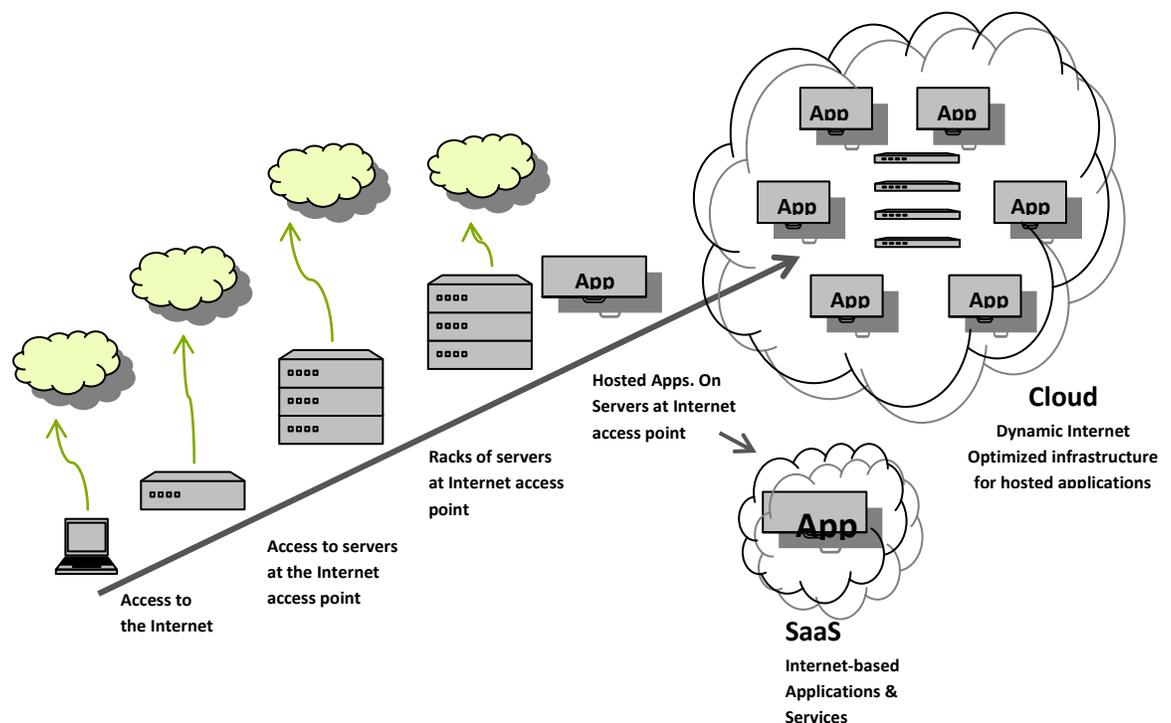


Figure 1 – Desktop and Internet evolution into Cloud

If we think of it in terms of the everyday utilities we consume, i.e., electricity, gas, telephone, television, etc., we see that all of these services are presented to the end-user in a simple way without the user needing to know how the services are provided. Similarly cloud computing offers computer application to developers and users this same view that simplifies and essentially ignores much of the details and inner workings. Just as with the utilities like electricity, the user pays for only what he needs or consumes without having to buy and maintain the whole power plant. That said, there is still a great deal of confusion as to what cloud computing really is.

	Cloud	Not cloud
Infrastructure as a Service	<ul style="list-style-type: none"> • Remote delivery • Multi-User • Scalability and load-balancing • Eliminates need for on-site facilities 	<ul style="list-style-type: none"> • Some would argue this is more on-demand delivery of virtual machines than pure cloud • Only some vendors offer truly automated flexibility, e.g. Amazon Auto-Scaling
Managed hosting	<ul style="list-style-type: none"> • Remote delivery • Eliminates need for on-site facilities • Multi-user 	<ul style="list-style-type: none"> • Provisioning and de-provisioning not automatic • Immediate scalability often not available
Virtualization	<ul style="list-style-type: none"> • Multi-user • Flexibility • Scalability 	<ul style="list-style-type: none"> • Provisioning and de-provisioning not automatic • Flexibility not infinite and subject to hardware constraints
Private cloud	<ul style="list-style-type: none"> • Multi-user • Remote delivery via intranet • Flexibility • Scalability 	<ul style="list-style-type: none"> • In most cases, it is on-demand virtualization • Billing mechanisms often not in place; reliance on chargeback • Few companies can achieve economies of scale necessary to make this work • Flexibility not infinite and subject to hardware constraints
File sharing and collaboration	<ul style="list-style-type: none"> • Remote delivery and access • Multi-user 	<ul style="list-style-type: none"> • Flexibility is limited by hardware and space constraints
Email	<ul style="list-style-type: none"> • Remote delivery and access • Multi-user • Flexibility • Scalability 	<ul style="list-style-type: none"> • Provisioning and de-provisioning not automatic • Immediate scalability often not available • Often more managed hosting than "cloud"
Voice services	<ul style="list-style-type: none"> • Remote delivery • Multi-user • Scalability 	<ul style="list-style-type: none"> • Many services require additional customization and are incapable of working "right out of the box"
Mainframe	<ul style="list-style-type: none"> • Multi-user • Flexibility 	<ul style="list-style-type: none"> • Provisioning and de-provisioning not automatic • Hardware constraints exist • Risk of a single point of failure within the infrastructure • Today, could be synonymous with private cloud

Figure 2 - Cloud vs. not cloud

2

Cloud computing models

2.1	Public cloud	11	
2.2	Private cloud	11	
2.3	Hybrid cloud	12	
2.4	What is moving to the cloud?	15	
2.5	Cloud in corporate business and IT strategy		16

With the entire buzz around cloud, you may wonder if you need to be in the cloud just to stay relevant. Our feeling is that no one should move to the cloud just to move to the cloud. First, understand how the cloud is going to enable your business objectives. Are you trying to accommodate seasonal variations in your business without 'building the church for Easter Sunday'? (Think online tax business in January - April, retailers on Cyber Monday, the floral business on Mother's Day, and weight loss on New Year's Day). Then, ask yourself how the cloud will enable those objectives. What are you looking to solve: scalability, capacity, elasticity, agility? Next, evaluate how mission critical your cloud project is. Cloud uses a range from test/development "sandbox" environments through elastic compute, which enables firms to meet demand during their peak times. Lastly, internalize and discuss how these factors drive your decision to leverage the cloud. The more mission critical your application is, the deeper is evaluation.

Even in the IaaS space, there are seemingly as many definitions of "cloud" as there are providers. Research firm [Vanson Bourne](#) suggests you only call an offering *cloud* if it's available on-demand and self-service, and if the pricing model is scalable, elastic, and pay-as-you-go. In many instances companies think they're buying "cloud" and often times paying regular rent on storage they are not using, or not using entirely. According to Vanson Bourne, that's not cloud.

Cloud computing falls into two general types, each still utilizing the three layers of cloud (PaaS, IaaS, and SaaS), plus a combination of a few types:

- Public cloud
- Private cloud
- Hybrid cloud

Let's review each type individually.

"Infrastructure as a service (IaaS) is the provision of on-demand compute resources coupled with associated storage and data center networking capabilities. IaaS consists of system infrastructure services that are the most basic and fundamental form of cloud computing service and parallel the infrastructure and data center initiatives in IT today." - Gartner

2.1 Public cloud

A public cloud is one based on the standard cloud computing model, in which a service provider such as Amazon, Microsoft or Intermedia makes resources, such as applications and storage, available to the general public over the Internet. Public cloud services may be free or offered on a pay-per-usage model. An enterprise (user) buys (rents) a specifically built for it tenant (piece of a public cloud) from a hoster or a telecommunications company, and adds an application layer on top of it. In general the user pays fees for the resources he uses similar to services from a utility.

Service Provider Hosted Public Cloud is one based on the standard cloud computing model, in which a service provider such as Amazon makes resources, such as applications and storage, available to the general public over the Internet. By definition it is a multi-tenant public cloud.

The multiple tenant public cloud is a shared multi-tenant environment built on a highly efficient, automated, and virtualized infrastructure. Other key elements of the cloud include standardized application platforms provided as a service and a self-service portal that enables business groups to request and manage capacity for their applications.

2.2 Private cloud

The Private Cloud is a term used to describe offerings that emulate cloud computing on private networks (either enterprise's own on-premise, or in a dedicated rack in a hoster's datacenter). These, typically virtualization solutions, offer the ability to host applications on virtual machines in a company's own set of hosts. These provide the benefits of utility computing – shared hardware costs across company departments, the ability to recover from failure, and the ability to scale up or down depending upon demand.

The single tenant private cloud is an isolated environment built on dedicated hardware. For example, Microsoft Web Services is adding dedicated cloud instances to its Virtual Private Cloud (VPC) offerings that make cloud computing single-tenant by putting data on isolated hardware. Dedicated instances ensure that no other companies are running on the same physical host, which incurs an additional cost for the less optimized hardware usage.

What's important to note in this scenario is that all instances might not be running on the same dedicated hardware machine. The virtual machine (VM) instances might be spread out on different hardware to minimize hardware failure. Thus, this scenario becomes more expensive, in terms of hardware, because the hardware is dedicated to one single customer. Therefore, what we need to think about is the following: what is the additional cost that the host provider can charge a customer for this convenience? Today, Amazon charges its customers a \$10/hour fee for this convenience in order to recuperate the cost of not having this dedicated hardware available to other customer.

Why the interest in private clouds? IT departments have long sought utility-like IT environments where computing resources and applications can be provisioned with greater efficiency. They're using systems management software, cluster and grid technology, load balancing, and virtualization to do it.

The major criticism is that corporate users still have to buy, build, and manage their cloud environments. However, the positive aspects outweigh this simply because the private clouds can service multiple organizations within the company, provide more efficient use of resources, and allow much tighter control of sensitive information.

When to recommend enterprise private cloud for business productivity? The top reasons customers are building enterprise private clouds are:

- Elimination of redundancy
- Data sovereignty and security
- Single tenancy (dedicated service)
- Customer has specific or strict requirements

2.3 Hybrid cloud

A hybrid cloud is a cloud computing environment in which an organization provides and manages some resources in-house and has others provided externally. The hybrid cloud infrastructure is a composition of two or more clouds private or public that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds). Hybrid clouds address the needs of organizations that require periodic access to highly scalable computing resources beyond the capacity of the enterprise's existing data center infrastructure.

Organizations benefit from the enhanced efficiency and security of the private cloud, but are prepared for large and unexpected spikes in demand for workload specific applications. For example, an organization might use a public cloud service, such as Amazon Simple Storage Service (Amazon S3) for archived data but continue to maintain in-house storage for operational customer data. Another popular use is Amazon's Elastic MapReduce web service offering from Big Data analysis and processing, which can replace your existing expensive database and free up space and power costs.

Ideally, the hybrid approach allows a business to take advantage of the scalability and cost-effectiveness that a public cloud computing environment offers without exposing mission-critical applications and data to third-party vulnerabilities. This type of hybrid cloud is also referred to as hybrid IT.

Hybrid Cloud:
Cloud infrastructure composed of two or more clouds, private or public, that remain unique.

Hybrid cloud computing combines the benefits of public cloud resources where some IT functions are managed externally, while a defined percentage of a company's IT stays on-premise in a 'private' cloud.

The business benefits from the cost and flexibility advantages offered by public cloud computing and can apply that model to the data that it feels happy to manage externally. At the same time, the business is able to retain customer data (and other Intellectual Property) inside its own data center.

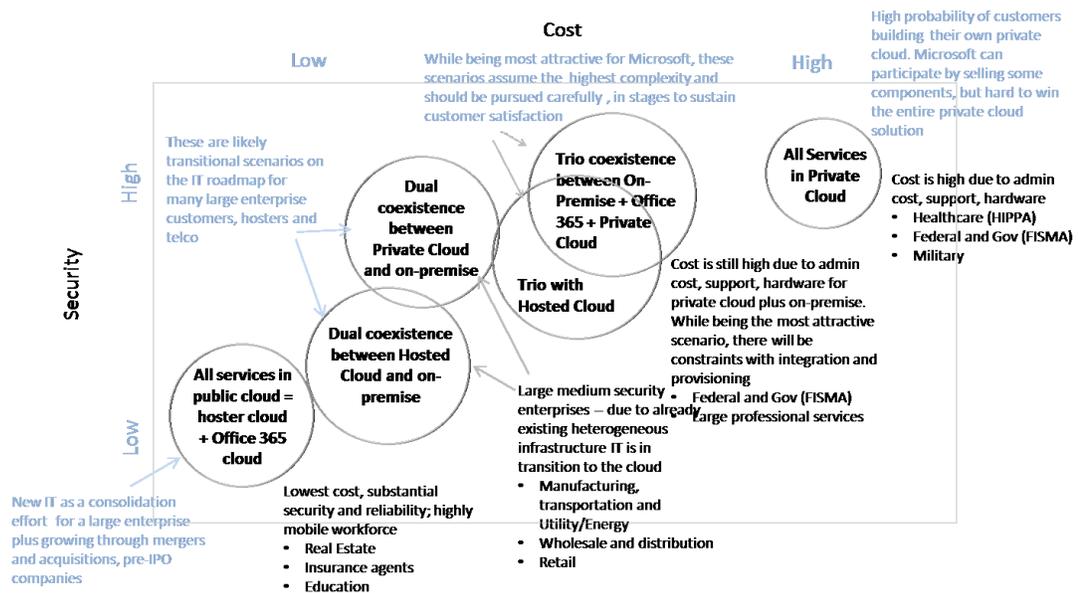


Figure 3 - Private cloud scenario attractiveness based on cost and security, with Microsoft Office 365 as an example

- Cloud-native applications
- E-business hosting
- General business applications
- Enterprise applications
- Test and development
- Batch computing

2.4 What is moving to the cloud?

The predominant question that businesses face when examining cloud migration is, "What are the applications or workloads that need to be brought to (or placed into) the cloud?" In this section we will identify the six most obvious candidates.

2.4.1 Cloud native applications

These are applications that are specifically architected to run in a cloud IaaS environment, using cloud transaction processing (TP) principles.

2.4.2 E-business hosting

These are e-marketing sites, e-commerce sites, SaaS applications, and similar modern websites and Web-based applications. They are usually Internet facing. They are designed to scale out and are resilient to infrastructure failure, but they might not use cloud transaction processing (TP) principles.

2.4.3 General business applications

These are the general-purpose workloads typically found in the internal data centers of most traditional businesses; the application users are usually located within the business. Many such workloads are small, and they are often not designed to scale out. They are usually architected with the assumption that the underlying infrastructure is reliable, but they are not necessarily mission-critical. Examples include intranet sites, collaboration applications such as Microsoft SharePoint, and many business process applications.

2.4.4 Enterprise applications

These are general-purpose workloads that are mission-critical, and may be complex, performance-sensitive or contain highly sensitive data; they are typical of a modest percentage of the workloads found in the internal data centers of most traditional businesses. These workloads are usually not designed to scale out, and may demand large virtual machine (VM) sizes. They are architected with the assumption that the underlying infrastructure is reliable and high-performance.

2.4.5 Test and development

These workloads are related to the development and testing of applications. They are assumed not to require high availability or high performance.

2.4.6 Batch computing

These workloads include high-performance computing (HPC), "big data" analytics and other workloads that require large amounts of capacity on demand. They do not require high availability, but may require high performance.

2.5 Cloud in corporate business and IT strategy

Initially, moving to the cloud was very ad hoc. You've heard the term "shadow IT", which implies that internal teams are creating their own environments (outside the purview of IT) in order to meet their deadlines and moving applications to a more agile environment. What's transpired as a result is the backlash of this instantaneous need. IT teams wrangle with the mushrooming growth and the attempt to institute controls for various generational versions and integrate them into corporate architectures which were never intended to be supported.

IT organizations now look to marry the consumer model of cloud - the immediacy of the public cloud - with the control and governance that they need to have over their environment.

Cloud is more deliberate now, not ad hoc

Adopting a cloud strategy is an instrumental part of the plan for changing up their business process, or how they will diversify their business. They can articulate how cloud will play a role in their infrastructure of choice. The result is that they will have more than one cloud of choice, for example Salesforce for CRM, Work Day for HR, etc. They will inherit clouds from the era of ad hoc, but they are also going to make strategic decisions about clouds that will be part of their corporate business and IT strategy. Cloud management is becoming increasingly important as IT teams will absolutely have to manage multiple clouds, establish IT policies and governance, and work hand-in-hand with development and operations teams.

2.5.1 IT policies

IT policies ultimately set the stage for the random acts of rogue people in the corporation lighting up clouds with ramifications. When you look back, IT used a traditional data center and had ITIL standards, or had an IBM shop with a governance model that said how you archived, what applications would be distributed, world access, things that were all second nature to those operating the data center. But when you have multiple clouds and organizations haven't spent time defining what the policies are, it's very hard to manage. This is why cloud management platforms like CSC's ServiceMesh, RightScale, and ScaleXtreme are rapidly gaining popularity.

2.5.2 Application impact

Let the workload be the guide. The application dictates the model, i.e. gold/silver/platinum, the architecture, and whether or not you need an active/active versus active/standby disaster recovery site. It is really a software-centric decision when you're choosing your cloud, so you want to have a way to build the application so you can run it anywhere. Think of creating a blueprint with agile development processes, which you can move to and among cloud service providers (CSPs). Before long, IT organizations will pre-select designated CSPs by workload type and deliver them "as a service".

2.5.3 Automation

The third point is really kind of interesting... we talk a lot about desktop management, virtual desktop, and unified desktop. When you look at the recent acquisition of Air Watch by VMware - we are practically jumping over desktop virtualization and looking at the mobile

desktop - and the ability to distribute applications to multiple devices from the cloud. This will be the next big wave of how we provide users within the enterprise the tools they need to do their job. A cloud management platform will play a role in being able to automate application release: you build the applications using blueprints, you manage how you run them in a multi-cloud environment with policy, and you automate how you release those applications to the enterprise.

Essentially, you use application “blueprints” to embed IT operations best practices for the governance and configuration of deployment environments that your development and test teams need to do their work. It’s all about increasing the pace and frequency of software releases to get new business functionality out to market faster.

Development and test get to use these from the beginning, and the environments are cloud-based and can be provided on-demand directly from their end users tools.

Now development no longer has to wait, or manually configure infrastructure and platforms on its own. Operations gets to drive environment consistency across all the SDLC stages and the chance for configuration related-defects slipping into production are reduced dramatically.

Every organization and IT department are different but the classes of policies are the same. If they can start with a platform they can nuance, that will give them a leg up.

3

Six selection criteria for a cloud service provider (CSP)

3.1	Performance	17
3.2	Security	18
3.3	Support	20
3.4	Pricing	21
3.5	Tenure	23
3.6	Governance	24

In this section we will consider the six selection criteria for a cloud service provider (CSP): performance, security, support, pricing, tenure, and governance.

3.1 Performance

Cloud providers use various cloud platform architectures, deliver different compute and network performance leading to differing levels of application and database performance results. Most CSPs can offer computing resources of varying sizes — from the smallest single-core instances to the largest multi-core mammoth-memory instances. What’s not so clear, however, is that storage IO can really vary from one cloud provider to another. And storage IO, not CPU, is often times the key determinant in your application performing exceedingly well or performing poorly in the cloud. Also, many providers’ customers suffer from “noisy neighbors” and variable compute and network performance due to flat networks.

The noisy neighbor problem is defined by the fact that Cloud in general is based upon shared resources. In this case the physical server upon which customer instances reside hosts multiple virtual machines. While it is generally easy and effective to partition CPU and RAM resources between virtual machines the disk subsystem is one which is extremely difficult to partition. In a noisy neighbor scenario one or more virtual machines on the physical host are consuming very large amounts of disk I/O resulting in very poor performance for the remaining virtual machines. To protect against this scenario, companies can gain significant advantages by utilizing CSPs who offer single tenant cloud servers or who are at least transparent about the contention ratios customers can expect when hosting on their infrastructure.

Many companies fail to baseline and/or define their utilization requirements in terms of granular metrics such as IOPS, RAM and bandwidth. Because they require more time and energy to gather, these specifications are guesstimated, which always means rounded up. Because different CSPs bill differently (e.g. RAM, resource pools, IOPS), it is crucial to define

IOPS (Input/Output Operations Per Second, pronounced eye-ops) is a common performance measurement used to benchmark computer storage devices like hard disk drives (HDD), solid state drives (SSD), and storage area networks (SAN). As with any benchmark, IOPS numbers published by storage device manufacturers do not guarantee real-world application performance.

-Wikipedia

these as granularly as possible. The result is a tighter CSP comparison and, of course, minimal overbuying. (Isn't matching spend to use a key value proposition for cloud adoption in the first place?)

Performance and usability questions to consider:

- (1) Infrastructure Questions:
 - a. Will they make available architecture diagrams and disaster recovery plans?
 - b. Are they using current generation, enterprise-grade hardware to host my data?
 - c. Will they reveal the names of the hardware vendors they utilize to provide the service (Cisco, Hewlett-Packard, NetApp, Fortinet, VMware, Microsoft, etc.)?
 - d. What commitments will they make regarding upgrades and refresh?
 - e. What reports are available to reflect my storage, compute and network usage?
- (2) VM and Hypervisor Questions:
 - f. How many VM modules will be used?
 - g. Do they offer flexible, easy-to-create server options to answer fluctuating capacity needs on-demand?
 - h. How many CPUs, GB of RAM or disk can a single VM have?
 - i. Can I leverage the same hostnames, IPs and MAC addresses they use internally?
- (3) General Questions:
 - j. Do they offer a range of services that can be mixed and matched to best fit the requirements of all my cloud-based applications?
 - k. Can they integrate hybrid solutions requiring connectivity between my private data center, colocation, managed hosting and cloud?
 - l. What is their UnixBench or iPerf benchmark scores compared to other providers?
 - m. Which offerings, if any, are being "white-labelled"?
 - n. Do they offer a range of connectivity offerings, from general Internet access to high-speed secure private networking? Is the connectivity provided directly from the provider or sourced via third party? Is there a service level agreement (SLA) for the connectivity and is that connectivity integrated into the cloud solution?
- (4) Portability Questions
 - o. Can they export VMs and easily migrate, or port, to an alternate provider in the future? Which hypervisor do they use? Is it compatible with the brand and version used in-house?
 - p. With whom do they partner or OEM for migration assistance? Is there a fee? What is the estimated time to migrate to/from the CSP?

- q. Will they charge an exorbitant fee to download the VMs?
- r. What bandwidth and/or mode of transport are available to download the VMs - Internet FTP, multi-threaded FTP/HTTP, ship a hard drive, etc.?
- s. Who owns the VMs and software licenses?

3.2 Security

Cloud security is at the forefront of the minds of those tasked with carrying out cloud initiatives, not to mention one of the predominant reasons cloud adoption has been deliberate with hybrid configurations being their first step. The majority of these debates fail to recognize that the actual physical presence of any device or server on the Internet is no longer a valid threat. A server physically located in a secure room within your office building is as vulnerable to attack by hackers and saboteurs as any Internet-connected PC. Physical walls, doors, and locks, no matter how strong, will no longer prevent your vital information from being compromised.

The safety of your data can only be guaranteed if the server is protected against Internet-based attacks. Your best defense is a professionally maintained system with firewalls and 24/7 monitoring by well-trained, experienced system administrators. In-house data centers do not typically employ security engineers, thus providing hackers with even greater opportunities to steal data or cause even greater harm. Many CSPs will offer advice and assistance to help you design and maintain a highly secure environment. They often times have their own security operations center or contract with a third party.

You may think that security, especially as it relates to public cloud offerings, is an excuse to shy from moving mission critical data to the cloud. But don't rule this idea out completely until you understand how prospective providers handle this area. Be sure to understand how the data is secured by asking the following questions, which encapsulate the key points. For many companies and applications, it is all about acceptable risk and risk tolerances.

Security questions to consider:

- (1) Does they offer a secure environment, both physically (data center) and logically?
- (2) How does their physical (data center) security compare to the security of your own IT environment?
- (3) Do they maintain a "defense-in-depth" solution that meets all of your internal security requirements?
- (4) Is the security built into the product offering, or is it something you have to design and bolt onto the solution?
- (5) How are the following topics addressed:
 - a. Encryption
 - b. Identity and access management (IAM) enforcing access rights and role-based permissions

- c. Backup and disaster recovery processes, procedures and tools
 - d. Data cleansing and sanitizing
- (6) How is encryption handled, specifically, "in transit" versus "at rest"?
 - (7) Is the security provided by and supported by a professional internal organization? If a third party is used, what is that organization's reputation?
 - (8) Can the security provided be customized and adopted to meet requirements unique to my enterprise?
 - (9) Do they meet all of the various business and regulatory compliance that the enterprise is legally required to implement?
 - (10) Can they supply documentation of their compliance/security processes and procedures?
 - (11) Can I be confident that their security measures are as stringent (or more) as what's implemented within my own data centers?
 - (12) Can they tell you with certainty where your business critical data will reside?

3.3 Support

Although many providers claim to offer "extraordinary" levels of support, CSPs have not been known to offer superlative customer support. And not all support or SLAs are created equally. You'll want to know how the provider stands behind their offerings. Do they offer comprehensive and well documented SLAs? Are they applicable day one or after some arbitrary time has elapsed? Are they clear and easy to understand? Can you enforce them easily or do you have to jump through hoops?

Support questions to consider:

- (1) What levels of support do you need both initially and over time?
- (2) How is a change order handled? Via self-service or paperwork? How quickly will changes be propagated?
- (3) Do you require a hands-off solution or more personal attention and hand-holding?
- (4) Which aspects of your environment do you want to manage, and what would you prefer to delegate?
- (5) Will I have phone access or email/chat only?
- (6) Describe your local team of cloud experts. How are you staffed after hours in terms of volume, location, and breadth of skills?
- (7) What is the default escalation path? What options exist to expedite escalation by defined level of severity?
- (8) Do they offer a global view of a client's entire cloud solution via customer portal? Are issue tickets viewable in real time? Can they demo the client portal for me?
- (9) Will they assign a single point of contact for all contracted services? (Some services may be white-labeled and directed elsewhere to troubleshoot.)

Key Decision Criteria:

- Performance/Scalability
- Security
- Support
- Pricing/Cost Controls
- Tenure
- Governance

-
- (10) Do they offer free technical support or only various tiers of paid assistance?
Subscription-based or ad hoc?
- (11) Under what circumstances will I speak with a call center agent or a technically astute person?
- (12) Are the support centers offshore or U.S.-based? There are benefits to each; however one will generally be preferred versus the other.

My point is that in order for this partnership to be successful, there should be a careful fit. Be sure your operations teams will easily mesh and the relationship is complimentary. A great way to validate a provider's support offering is to simply call the toll-free number, check their authentication policy, and ask a question or two.

3.4 Pricing

The popular driver behind most businesses moving to the cloud is to avoid overbuying infrastructure that is only used a fraction of the time.

Nearly all CSPs offer this elastic, utility-based pricing model, allowing you to pre-define (or reserve in the case of Disaster Recovery-as-a-Service) your base configuration and bursting parameters. However, this financial benefit may be lost when you are forced to buy packages that have pre-set CPU cores and RAMs tied to each other – often forcing you to over-buy certain resources.

Cloud customers should to be aware of the four main resources they are paying for: CPU, RAM, data storage and bandwidth. Many CSPs hide the total cost associated with the subscription-based pricing, choosing instead to focus on their “advertised” price. Also, upgrades may default to list price, often published online and “subject to change”. You must be aware of hidden costs and insist on transparency at the beginning of the relationship. Some of the most popular hidden fees include:

- Subscription charges – a monthly, pre-paid “package” for a specific level of resources.
- Software licensing fees – operating system, database software images are often left out of discussions
- Burst resource pricing above subscription level - Increased cost of additional bandwidth, CPU cores, memory, or storage
- I/O server-to-disk back-end network activity
- Amplified levels and types of customer support – phone and email support may not necessarily be included. Some providers even charge more for the escalation of a support issue.
- Data migration costs – what are the initial costs of modifying your application to fit the specific design requirements of a CSP and the cost of transferring your data to the cloud?

While it is human to name cost as a top three selection criteria, it has become nearly impossible to receive comparable responses from proposed CSPs. Some providers “bundle” resources, or offer resource pools, while others name hourly rates versus monthly.

Storage may be quoted using SSD, SAS or SATA, and depending on the provider’s infrastructure there can be a large variance in performance IOPS. Burst rates will vary and fail-over instances will have different recovery time objectives (RTO) and recovery point objectives (RPO).

Pricing questions to consider:

- (1) How is the solution priced?
 - a. If by resource pool, how are those defined? Can they be user-defined or are they pre-set as ‘use or lose’? Can spend be “trued-up” to align with usage on a (e.g.) quarterly basis? How can I right size to guard against over-buying?
 - b. Does it vary by compute capacity, network traffic, storage, additional services, flat fees and overages, etc.?
- (2) Is their pricing clear and understandable?
- (3) What is included and excluded from the complete solution?
- (4) Is support built into the pricing model, or is it a premium on top of the standard pricing?
- (5) Are essential components included or do I need an external party to provide?
 - a. Does the pricing deliver and satisfy the value I expect to realize from my cloud strategy? (Low pricing does not equate to high value if all other requirements are not fully satisfied.)
- (6) What access and view do I have into my environment?
- (7) Are their additional portals or metrics available beyond the standard offering?
- (8) Can I control who has access to utilize the cloud services?
 - a. What reports are available to manage use and internal cost center budgets?
 - b. Can spending and spending controls be implemented on a department or individual basis?
- (9) Do they offer and support an API to integrate the service into my operational model?

Only by thoroughly evaluating prospective service providers and vetting their ability to solve your unique business requirements, can you ensure a successful partnership. Decisions based solely on price or on limited understanding of the overall/long-term solution can be short sighted and lead to frustration and a great loss of time for all parties involved.

The normalization of provider responses, even with a strict and comprehensive list of technical requirements, is a great science. Not only you compare initial costs, I recommend modelling future growth. For example, be aware of selecting a CSP, which bills for RAM, if that is where you will grow most aggressively.

SATA or Serial ATA stands for Serial Advanced Technology Attachment and is the next generation computer bus interface from the parallel interface ATA. SAS, or Serial Attached SCSI, is a faster and historically more expensive interface. SATA and SAS refer to disk drive interfaces, whereas SSD (solid-state drive) refers to a particular kind of internal technology.

3.5 Tenure

With new pure-play cloud providers rapidly emerging, increasing and aggressive merger and acquisition rates within the cloud space, the tenure, reputation and track record of each provider should not be underestimated. Larger brand names, such as those named each year in Gartner's Magic Quadrant, offer an inherent level of trust and confidence – and at a steep price. Newer entrants are more agile and competitive both in terms of pricing and contract terms. These benefits may very well be worth the assumed risk.

You'll want to start by understanding how long the provider has been in the IT outsourcing business and if they have the financial resources to continue to deliver and upgrade services throughout the business lifecycle. Based on whether they are a public company or private equity backed, assumptions can be made about their mindset.

Business is global, is your provider? Does the service provider provide the seamless global footprint your business needs to grow? Can the data and environment be easily replicated globally? Will you be able to manage the infrastructure transparently, regardless of time or location?

Ask about their physical location(s) and proposed expansion. Physical location is very important for a few reasons – statute restrictions may regulate your use of customer data; Jurisdictional requirements will regulate the laws that apply in the event of a dispute; and environmental factors to assess natural risks such as hurricanes, earthquakes, etc.

Tenure questions to consider:

- (1) How long have they been in business offering the proposed services?
- (2) What is their financial position – private equity, backed, publicly traded?
 - a. What percentage of their revenue comes from IaaS and how is this trending?
 - b. How many customers are live on this platform? How similar are these customers to you and the application(s) you plan to host?
- (3) What metrics are available to prove data center and infrastructure uptime? Have them describe the details of any outages.
- (4) Do their processes and technologies stand up to recommendations from the IT industry's leading organizations? What are analysts saying?
- (5) What is their history and the track record of their partners – including, and quite possibly, most importantly – the data center operator with whom they lease space?

There is significant competition in the cloud computing industry and it is vital for providers to have and maintain excellent references for long-term viability. Be sure to ask prospective providers for references that can specifically highlight service reliability, quality of service, quality of support, and network performance. Review their website and ask them for case studies and analyst and media coverage examples. And be honest answering the question:

Under the U.S. Health Insurance Portability and Accountability Act of 1996, a HIPAA business associate agreement (BAA) is a contract between a HIPAA covered entity and a HIPAA business associate (BA).

It is an addendum to the Master Service Agreement (MSA) and provides additional detail and assurances related to how the service provider will protect a client's hosted personal health information (PHI).

Has vetting this provider been an easy process, or does it make you question the level of service and support that lies ahead?

3.6 Governance

The growing requirement for supervision by IT of every cloud services engagement within a company must be emphasized. Each service must be vetted by IT, with risky services filtered out according to leading analysts. IT will remain in possession of administrative control, own the decision making power over provider selection and maintain a catalog of services acquired by the company.

Governance is an important part of your SLA with a service provider. A service level agreement (SLA) is not a guarantee. Rather, it is a commitment from the provider to you that when things go wrong, and they will at some point, that they will respond in a timely manner and have skin in the game. By governance I also mean to include visibility, auditability, and an overall ability to ensure necessary compliancy.

Governance questions to consider:

- (1) How can I be confident that my data is wiped when you refresh your hardware?
- (2) Will the provider sign a Business Associate Agreement (BAA), if applicable?
- (3) If the standard SLA is inadequate, will the CSP work with you to customize something that will be acceptable to both parties?
- (4) In the event of an outage or problem, how are responsibilities shared between your company and the CSP?
- (5) What is the provider's level of transparency? Do they proactively inform consumers of SLA non-compliance and breaches?
- (6) What type of restitution is offered in the event of a disaster?
 - a. Does the language limit the provider's liability at the expense of your company's data?

You can assign privacy responsibility to vendors, but you can't delegate-away accountability through contracts. Be sure to map internal key controls to vendor responsibilities and instill governance mechanisms to manage vendor compliance and retain strong customer audit rights.

As far as reliance on information security standards frameworks, ISO 27001 / 27002 and SysTrust (SOC 2) are far more useful than SSAE 16's SOC 1. SOC 2 and 3 were actually written with the intention of being able to evaluate cloud and managed service providers. Finally, expose stealth subcontractors and multi-layered cloud suppliers, and include them in the privacy and compliance programs.

4 Four cloud sourcing options

4.1	Work directly with the hundreds of providers	25
4.2	Utilize the expertise of a cloud services broker (CSB)	26
4.3	Engage a large, expensive consulting firm	27
4.4	Leverage a boutique consulting and sourcing firm	29
	Conclusion	30

Cloud sourcing decisions are vastly based on a Customer's (Enterprise's) effort, time, and cost to internal IT department.

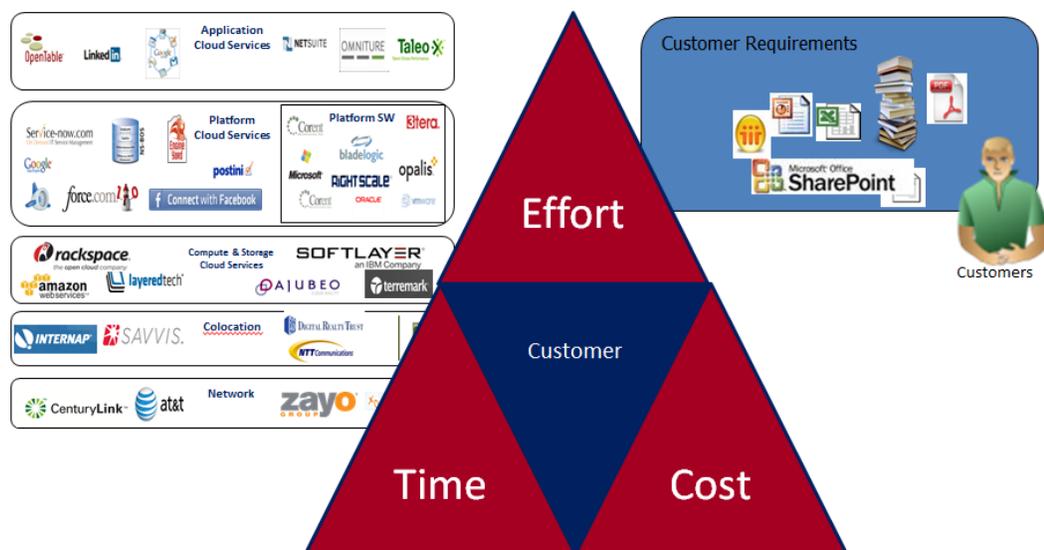


Figure 4 – Cloud sourcing ecosystem

There are four potential approaches to managing the identification, evaluation, negotiation and eventual onboarding of a proposed IaaS deployment:

1. Work directly with the hundreds of providers
2. Utilize the expertise of a cloud services broker (CSB)
3. Engage a large, reputable diversified consulting firm
4. Leverage a boutique consulting and sourcing firm

Companies may elect a combination; however, we will highlight the strengths and challenges of each of these four for the sake of purity.

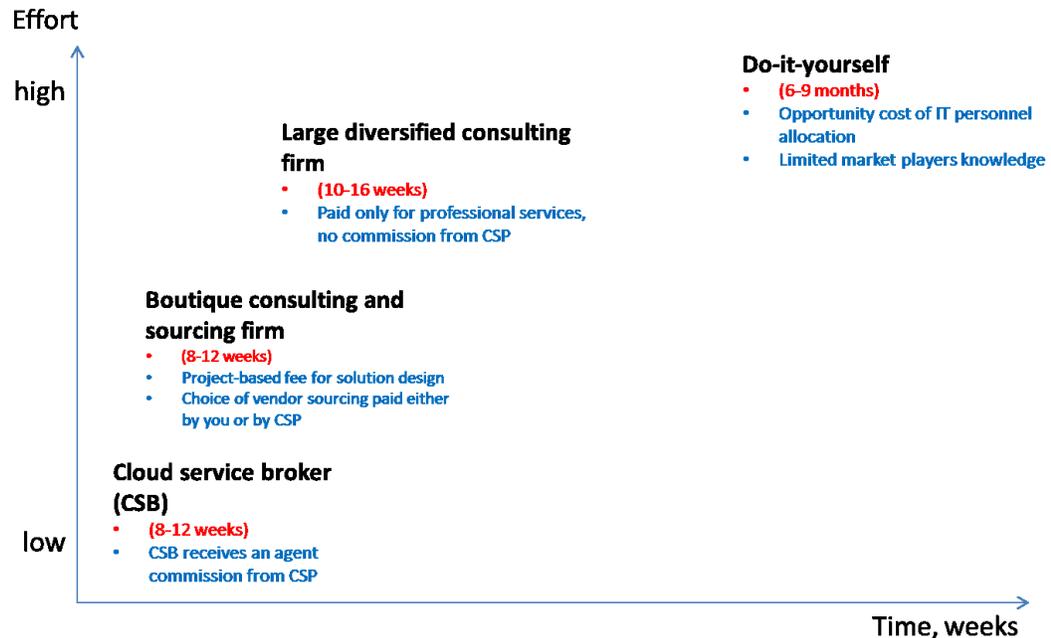


Figure 5 – Sourcing options attractiveness based on the effort required from internal IT department, time, and pricing option

4.1 Work directly with the hundreds of providers

It is as complex as it sounds. There are hundreds of providers out there, and it is NOT feasible for a mid-size enterprise IT team, which is not monitoring constantly the service provider market to take this on, evaluate all players, and short list and select who truly the most valuable and relevant CSPs are for several reasons. These include not knowing the “who’s”, time, frustration, normalizing and comparing data, and a few more reasons, too, because your organization might simply not qualify to work directly with a CSP. Also, CSP’s salespeople don’t necessarily enjoy asking buyers (repeatedly and in various ways) about their current percent of utilization, how much RAM they’ll need, and desired recovery point objective (RPO), for example. This breeds mutual frustration and delay.

The reality is that most sourcing is happening without the assistance of a consultant, broker, agent, or reseller. And what you get is either an instance that is spun up on a public cloud like Amazon Web Services with no competitive analysis, or a cursory search of brand names initiated through website inquiries launches individual and repeat calls, meetings, whiteboard sessions, portal demos, lunches and (if you are lucky!) sporting game invitations.

4.2 Utilize the expertise of a cloud services broker (CSB)

Instead of sitting across the table from each prospective provider sales team (and again later with their management), the use of cloud services broker allows you to share your business initiatives, project and stakeholder requirements, provider selection criteria, decision and implementation timelines and budget **only once**. It is then the broker's job to probe for and fine tune the requirements, invite only the most relevant CSPs to participate in the evaluation, and serve as the general program manager for the organization which consists of such tasks as: lead liaison for correspondence, scheduling and deadlines, coupled with normalizing vendor responses into a competitive matrix outlined with industry objectivity and perspective in order to aid the decision making process.

Arriving at a clear illustration of the optimal provider, pricing, and terms with ease and expediency is what the CSB offers. A CSB knows what the industry is bearing, what negotiation tactics are suitable for your pricing requirements, and can clearly identify key terms and conditions language to assure long-term leverage and protection.

In the case of most CSBs, compensation comes directly from the provider as a percentage of the business they bring. This allows the CSB to offer such services and expertise to their clients at a significantly reduced fee or often at no charge. The CSB has in place commission and compensation rates either directly or via a Master Agency or "distributor" of IT services. By having these provider-direct and Master Agency relationships, both the CSB and the provider are incented to work seamlessly together – the CSB owns the client relationship which the provider wants access to in this extremely competitive landscape, and the CSB is incented to have a flawless client implementation to assure of compensation after installation and for the life of the client. Because this synergy is in place, it lends comfort to the client buyer and gives validation to this lesser-known option.

There is a common misconception that *'where a broker is involved, the buyer is paying a higher rate'* in order to cover the commission. While there may be exceptions, the reality is quite the contrary. The commission margin comes directly out of the supplier's cost of goods sold (COGS), so not only will the CSB aid in driving pricing down, they also imply a threat of competition and can intimate knowledge of similar contracts. The CSB, an independent contractor, does not receive a salary, benefits, office space, laptop, parking pass and expense account as compared with a direct sales representative. Hence the cost of acquisition is minimal by comparison, especially considering the transient nature of sales people on the direct sales side. Some providers today are experiencing more than 45% of their annual sales from indirect channel sales efforts, so emerging cloud providers with little staff and brand recognition, their focus efforts (and money) are increasingly on educating and growing their channel mindshare.

I feel it's worth mentioning that with the rapid increase in popularity of the term 'Cloud Service Broker' companies should be certain to engage an individual or firm with proven managed and cloud services experience in alignment with the specific project. Very few CSBs

- A Cloud Services Broker gets to know you & your business
- They find out what you want in your ideal cloud relationship
- They have a stable of eligible candidates from which to choose
- They bring you together with the best candidates and help you find your cloud match!

.....

exist in the market today, which possess the infrastructure and operations knowledge and perspective to offer truly valuable guidance with vendor selection and contract negotiation. Make sure to qualify potential CSB candidates regarding their particular level of experience and expertise, and solicit references that can speak to their insight, responsiveness, and professionalism.

4.3 Engage a large, reputable diversified consulting firm

Fortune 500 companies on up in size will generally be inclined to engage a large, reputable consulting firm for a sizeable project-based fee in order to assist with such things as capacity planning, total cost of ownership (TCO) analysis, cloud readiness assessment, and RFP creation and management. These firms combine capabilities across management consulting, technology and business process outsourcing to provide comprehensive portfolio of services.

They have practice areas to address their clients' needs by industry and by enterprise function. Industry business services address the business and performance issues that are unique to a given industry. Function business services address issues and processes in specific enterprise functions, such as procurement or human resources, which are applicable across industries. Their offerings will span multiple growth platforms, such as strategy, digital, technology and operations.

As one might imagine, these engagements are incredibly time-intensive to negotiate and complete, generally spanning six to twelve months. There are teams assigned to work with business units and collect technical requirements from storage, security, network, compute, compliance, finance and operations leads. A project lead, or program manager, will compile the collective findings into a PowerPoint deck to present options for management's consideration and long-term planning.

This option is fantastic for larger organizations who take comfort in working with an established and trusted brand, prefer to be led through the cloud planning and financial modeling phases, and receive a formal and professional "CxO-ready" business case. The cost ranges will typically be anything from \$30,000 to \$500,000, and these consulting firms are not receiving any outside compensation from the eventual CSP placement.

4.4 Leverage a boutique consulting and sourcing firm

Boutique consulting and sourcing firms who specialize in recommending cloud service providers will generally be following a formal business process methodology. They have established a vast ecosystem and can call upon third-party vendors, tools, and professional services firms to ensure a holistic approach to addressing client solutions.

This described firm sits in between a pure CSB and a large multi-discipline, national or international consulting practice, such as Accenture or PricewaterhouseCoopers, which stands ready to assist in any stage of the business process lifecycle – base lining, requirements gathering, RFP management, and implementation.

Where companies already have their requirements clearly defined, and which can jump straight into vendor selection and contract negotiation, a CSB will be a great fit. In my experience, however, the need to conduct a thorough baseline of infrastructure, application and/or IT service audit and needs assessment is in order.

**Building a business case and cloud strategy before seeking a solution
is the foundational key to success in this effort**

What are your business drivers? Which applications are cloud-ready? What are the inter-dependencies (and latency restrictions) to consider? What is the preferred balance of operational (Opex) versus capital (Capex) budget? Have these been explored within your team, and throughout your organization? By taking the time to evaluate the unique challenges within an organization, the firm is able to gather much deeper insight into their business *and cultural* requirements and offer the executive management team a strong business case for choosing a particular type of CSP (i.e. pure play, hybrid, or ITO model) and beyond simply the match-making aspect.

What leverage or benefits does the company get when they sign a contract through a consultant or broker versus directly?

Companies working with CSBs or boutique consulting and sourcing firms realize far more negotiating leverage of pricing and contract terms. Not only have these entities been involved with the negotiation of hundreds of hosting contracts, they work with an impressive number of service providers and will share market intelligence with your team regarding 'what you can get if you only know to ask.' The service providers also know that when this 'CSP advisor' is involved, they are forced to be competitive and no-nonsense, or else face the risk of elimination.

Typical engagements offered by boutique consulting and sourcing firms may include:

- Analyzing costs of maintaining or updating constrained data center facilities to provide a baseline cost of ownership
- Evaluating and identifying potentially excessive data center spend
- Assessing current and projected business resiliency or BC/DR options
- Overseeing infrastructure consolidation due to mergers and acquisitions
- Delineating core versus non-core IT resources to maximize talent pool
- Recommending Service Provider partners to increase business agility

Consulting services are typically project-based and can be secured as needed to complement your skills and resources. Specifically, and as required to support executive business cases, consecutive engagements may be necessary, which include: Baseline, Needs Assessment, Requirements Definition, RFP Management/Vendor Selection, Contract Negotiating and eventually On-Boarding. Too often companies are pressured to begin with Vendor Selection before properly defining their requirements, and which eventually back-fires when they realize the chosen provider had not been, indeed, the correct fit. By outsourcing the Baseline, Needs Assessment and Requirements Definition to a consultant, companies gain the benefit of the research and documentation without repurposing headcount and gain much objectivity to the equation. When you take into consideration the consultant's industry expertise, methodology and seasoned templates, you'll find their fee well justified.

Conclusion

I founded Nuvalo with a mission to educate and guide companies through the noisy and chaotic cloud and datacenter marketplace. Having come from the Service Provider side, I've met with hundreds of companies who hurriedly negotiated bad contracts with little concern for fit, terms, and any long-term strategy. There are absolutely insider "tricks of the trade" and as an advisor to companies of all sizes, I walk away from each engagement knowing that my client secured the most competitive price and contract protection. Conversely, my client has every confidence that the most ideal provider partner was chosen.

When choosing a CSP, I encourage you to favor one that uses open standard-based components and APIs, a standard third party virtualization layer (e.g. VMware), also offers its customers good porting tools, demonstrates efficient portability, and includes post-term migration assistance obligations in the contract.

Conduct financial due diligence on your short listed providers, and take into consideration what their 3-5 year goals may be based on whether they are a publicly traded company, private equity backed, or operated by a telecommunications company.

Regardless of whether you choose to go it alone or call an outside advisor/broker, I recommend the following best practice for sourcing cloud services.

"Gartner sees a growing need for what it calls cloud services brokerages (CSBs) — a new breed of middlemen that sit between the corporate users of cloud computing services and the cloud vendors themselves. These brokers would vet and/or approve cloud services before they get deployed."

-- 'Gartner to IT: Get a Grip on Cloud Services or Else', GigaOm, April 2012

1. Prepare a formal RFP document and submit to 4-5 service providers, and assume at least one will “no bid”.
2. Define and agree on the characteristics that best define your ideal provider, such as competitive price, quality, industry recognized leadership, geographic footprint, financial health, willingness to negotiate fair terms & conditions, reputation, organizational culture, scope of services, process maturity.
3. Consult with outside advisors to conduct due diligence to vet candidates before the RFP is issued.
4. Check with other organizations.
5. Review your short-listed providers’ financial health (including “for sale” signs).
6. Select best bid(s) applying an apples-to-apples comparison through multiple down-select rounds (requires very specific RFP).

Keep in mind when negotiating term length that CSPs may try to make a term more attractive, but ample options exist for month-to-month (or even down to the minute) billing. Anything longer than 1-year should be scrutinized. Pre-negotiate the early termination fees, revenue or resource minimum commitments, and renewal options.

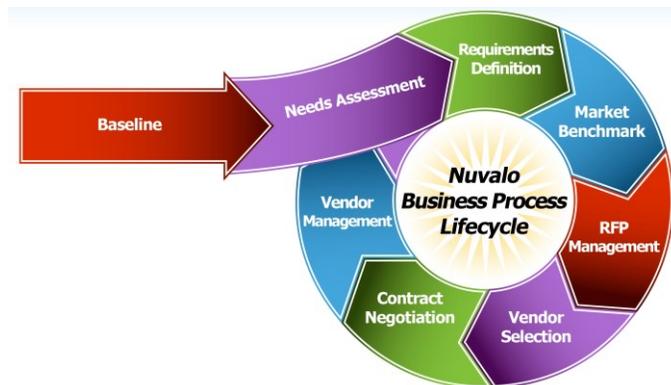
Approach the cloud with a long-term vision. Take a holistic approach to defining your road map, and remember that by taking the time to analyze and clearly define your architecture, utilization and business goals, you will be rewarded a hundred fold.

About Nuvalo

Nuvalo is a premier Seattle-based IT consulting firm specializing in the design and sourcing of hybrid colocation and IaaS deployments around the world. We build business cases and manage your vendor selection, ensuring you have all relevant data points and perspective before contracting for cloud, data center and network services.

Nuvalo takes pride in its ability to put insider-level knowledge of best practices and the provider landscape to work, ensuring each client's specific business needs are met by working with the right providers and solutions, under the best of terms for each client.

We have relationships with 400+ providers and stay abreast of emerging players, mergers, and industry news, which ensures our clients have the most current competitive landscape.



Contact Us

800 Bellevue Way NE #400

Bellevue, WA 98004

Tel. 877-9NUVALO

877-968-8256

Email: info@nuvalo.com

www.nuvalo.com