

Define Your Private Cloud Strategy

Strategic Framework and Considerations for Charting Your Journey to the Cloud

Microsoft Services

Last Update: Thursday, 20 October 2011

Author:

Microsoft Services

© 2011 Microsoft Corporation. All rights reserved. This document is provided "as-is." Information and views expressed in this document, including URL and other Internet Web site references, may change without notice. You bear the risk of using it.

This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes. This document is confidential and proprietary to Microsoft. It is disclosed and can be used only pursuant to a non-disclosure agreement.

This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS SUMMARY.

Table of Contents

- 1 INTRODUCTION: GETTING STARTED ON THE JOURNEY..... 1**
- 2 UNDERSTANDING KEY BUSINESS DRIVERS..... 1**
 - 2.1 ASKING THE RIGHT QUESTIONS.....2
- 3 INTRODUCING THE CORE IO MATURITY MODEL – PRIVATE CLOUD VIEW..... 2**
 - 3.1 AN INCREMENTAL APPROACH.....3
 - 3.2 UNDERSTANDING THE IMPLICATIONS4
- 4 THE IMPORTANCE OF MANAGING EXPECTATIONS 5**
 - 4.1 SERVICE PROVIDERS AND SERVICE CONSUMERS.....6
 - 4.2 ALIGNING BUSINESS AND IT7
- 5 NEXT STEPS ON YOUR JOURNEY 7**
 - 5.1 NEXT STEPS7
 - 5.2 FURTHER READING8
- 6 APPENDIX: TABLE OF ATTRIBUTES: CORE IO MATURITY MODEL–PRIVATE CLOUD VIEW 9**

1 Introduction: Getting Started on the Journey

Whether you want to take your IT operations wholly to the public cloud, keep them on-premises, or get them hosted off-premises using a private cloud model, or choose to straddle the two poles in a hybrid configuration, you must start with utter clarity on the reason for doing so. This may seem like stating the obvious, but [as we've seen](#), it can be tempting to get caught up in arguments and discussions about compelling new technology and neglect to perform the required due diligence on fundamental questions of business need—*What are the key drivers for change?*—and strategic direction—*What is the right approach for achieving your objectives?*

These are not simple questions with simple answers. As we've learned from our experiences working with numerous large organizations around the world on private cloud migration projects of different scope and aim, the unique challenges of cloud computing require new ways of thinking, planning, and cross-group collaboration to attain common goals. We've also seen that success can be won early by those IT leaders who are able to frame a realistic strategy at the outset, rooted in a thorough appreciation for the capabilities and limitations of the organizations they lead, as well as for the interplay of IT perspectives and priorities and those of the business as a whole. What's required is, in a sense, a "cloud frame of mind."

To help you acquire the proper frame of mind for the task ahead of you, we outline here a framework for thinking through the various strategic considerations required in a private cloud deployment project.

2 Understanding Key Business Drivers

In the reference architecture we've developed to help our customers make the transition to a dynamic datacenter (many of these concepts you can read about in our [Private Cloud blog](#)), we describe business drivers for cloud computing in terms of the competing vertices of agility, cost, and quality:

Agility	Cost	Quality
<ul style="list-style-type: none"> • Reduce time to market: Implement new business solutions quickly to accelerate revenue growth. • Better enable the solution development life cycle: Speed up business solutions through better development and test, and a fast path to production. • Be more responsive to business change: Deliver quickly on new requirements for existing business solutions. 	<ul style="list-style-type: none"> • Reduce operational costs: Optimize daily operational costs like people, power, and space. • Reduce capital costs or move to annuity-based operational costs: Benefit from reduced IT physical assets and more pay-per-use services. • Make IT costs transparent: Service consumers better understand what they are paying for. 	<ul style="list-style-type: none"> • Consistently deliver to better defined service levels: Better service leads to increased customer satisfaction. • Ensure continuity of service: Minimize service interruption. • Ensure regulatory compliance: Manage the compliance requirements that may increase in complexity with online services.

Where organizations locate their business need among these primary drivers and define their objectives as they consider their cloud computing options is a basic starting point in the process. For many in the IT shops of large organizations, whether private enterprises or public entities, the [private cloud is proving especially attractive](#), mainly for what it offers in terms of control over matters of security, data access, and regulatory compliance. Their primary interest in a private cloud architecture typically revolves around the very real pressures to cut costs without sacrificing too much control over essential data, core applications, or business-critical processes.

But according to a recent [Gartner survey of CIOs](#) worldwide, what's top of mind for a lot of them has less to do with the usual need to drive greater efficiencies and cut costs in IT than with the desire to increase enterprise growth and attract new business. This shows that IT leaders are beginning to think seriously about cloud computing as a way to turn IT into a key contributor to the positive side of the balance sheet, rather than a cost center that has to be constantly squeezed, prodded, and bled of every last inefficiency.

2.1 Asking the Right Questions

Cutting costs, increasing efficiency, and driving growth: these are all worthy objectives for implementing a cloud solution. Indeed, a cloud could conceivably deliver all of these benefits. But if we focus too much on all the things we could do with a cloud, we may fail to see what we should do—and *actually have the capability to do*—with a cloud. This is where a clear-eyed strategy for success comes in, which depends in large part on your ability to ask the right questions especially at the outset of your journey.

3 Introducing the *Core IO Maturity Model – Private Cloud View*

Once you've identified business needs and objectives, you need to understand fully the capabilities, limitations, and complexities of your current IT environment, which starts by performing an analysis of your technical and organizational maturity against the backdrop of the different capabilities of cloud computing. The next step is then to determine where you want to take your organization, assessing the prerequisites for the desired objective.

This process of change we describe in the *Core IO Maturity Model – Private Cloud View*. It is organized by infrastructure optimization stage (basic, standardized, rationalized, and dynamic), with the details of each stage described using an [industry-standard](#) definitional schema for cloud attributes (self-service, usage-based, service, elastic, pooled resources, and broad network access). As you progress in infrastructural and organizational maturity, your business will benefit from increasingly advanced private cloud-based scenarios.

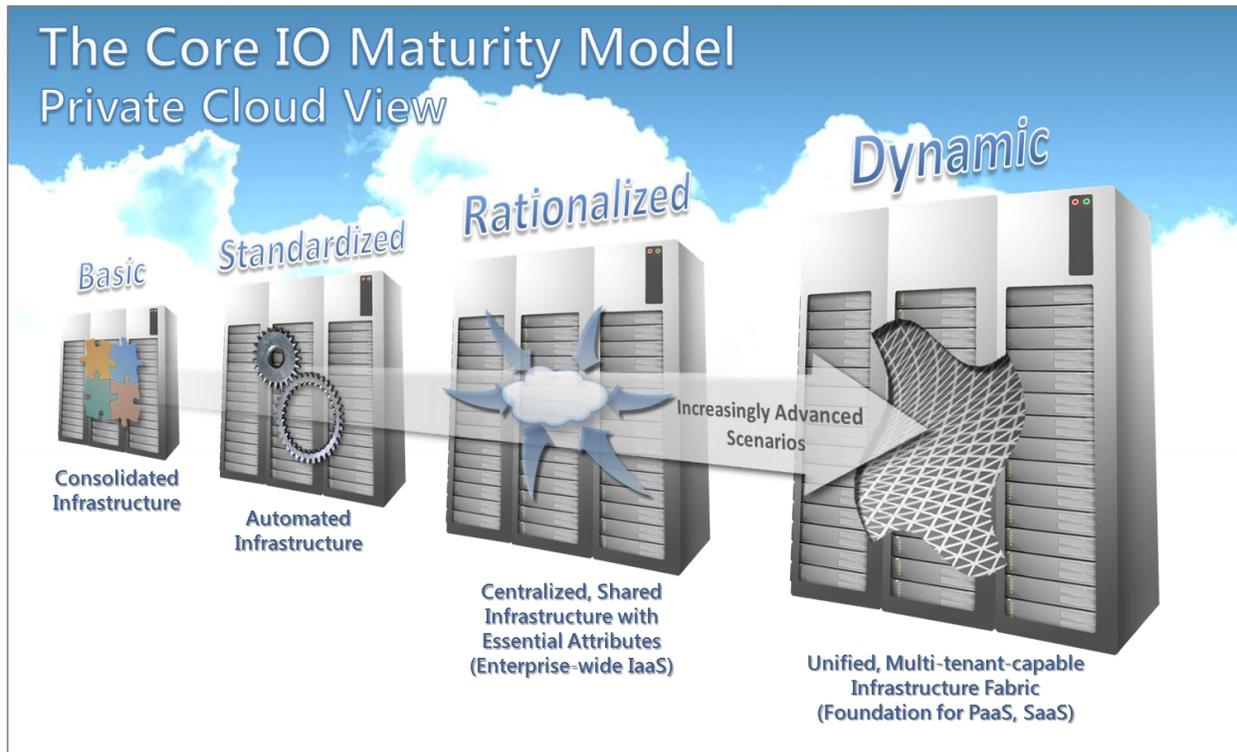


Figure 1. The Core IO Maturity Model-Private Cloud View.

3.1 An Incremental Approach

Many of the organizations we work with are starting at a relatively basic stage along the optimization continuum, where they've already managed to consolidate infrastructural resources for better cost-efficiencies through virtualization. If your organization fits this profile, a perfectly acceptable outcome might be to advance your business to the next stage by implementing more sophisticated infrastructure-level resource pooling (see the diagram above and the table of attributes for the *Core IO Maturity Model – Private Cloud View* [provided in the appendix](#)), which would achieve still greater cost savings as well as a measure of improved time to market.

Similarly, your current organizational capabilities may land you somewhere in the middle of the maturity model, with a relatively high degree of sophistication in those business areas you consider your top priorities, such as being able to respond to seasonal shifts in demand if you happen to be a major retailer. While your ultimate goal might be to bring into play platform as a service (PaaS) and software as a service (SaaS) architectures so you can leverage a larger set of hybrid cloud capabilities—such as anytime, anywhere access for your customers built on a highly unified set of compute, network, and storage resources—your near-term focus in the context of an infrastructure as a service (IaaS) model may be just in moving the dial specifically on automated provisioning and de-provisioning of resources.

It's in this approach—by making deliberate, incremental progress in the service of a longer-term strategy—that real IT transformation occurs on a private cloud model. The way forward is rooted in recognizing that change is an evolutionary process, where the investments you make in technology solutions must be evenly matched at each step by the maturity of your organization in managing them. Further, your strategy must be carefully applied in those areas where your business is likely to benefit most. Indeed, not all capabilities of a private cloud need to be—or should be—exploited. Greater elasticity, for example, may not be what your business fundamentally needs nor has the technical or organizational maturity to support. The task isn't to do it all, much less to do it all at once. The real task lies in balancing the potential goods of a private cloud solution against actual business needs, understanding your capabilities and limitations at each stage of the process, and charting a realistic, achievable course of action for getting it done.

3.2 *Understanding the Implications*

To be sure, considering capabilities and limitations in going from one stage to the next is just one vector of analysis. Any one of the capabilities or features of a private cloud solution with which you choose to align introduces additional considerations that may not be obvious at first. The objectives you choose for your private cloud will raise a number of questions about the various technical and organizational implications of implementing your solution.

Below are some examples of the kinds of questions you need to be able to ask in order to frame a comprehensive and realistic strategy for achieving your objectives. With reference to the maturity model diagram above, these are presented in the order of the fundamental private cloud attributes:

- **Self-service:** Allowing your users to provision the resources they need on-demand without human intervention may seem like an eminently desirable option, such as when it comes to partitioning new storage volumes or building out a new virtual machine cluster to support test scenarios, but very quickly you can run into questions of governance. How much control should you relinquish, and what are the potential consequences of offering a self-service model for common tasks? Consider the way SharePoint sites can proliferate in a large organization when left unchecked and the concomitant and increasingly costly burden on storage, network, and compute capacity.
- **Usage-based:** Pay-per-service, or “chargeback,” is one of the hallmarks of cloud computing, and if your private cloud strategy includes driving greater transparency of costs for the sake of better resource planning (among other potential objectives), you need to know the incentives you're trying to drive. Are you trying to reward good behavior and punish bad? In that case, a consumption-based tax might make sense, based on input/output operations per second (IOPS), for example. Or, you may wish to push more round-the-clock workloads to nighttime operations for the sake of power savings that support your company's “green” initiatives.
- **Elastic:** Being able to respond efficiently to fluctuations in resource usage—whether seasonal (as in the case of retail businesses), diurnal (as in the case of typical corporate email systems) or

some other periodicity—can represent a major selling point for cloud computing, in no small part because of the perceived benefit of reducing a company’s carbon footprint. But it is important to consider first whether you really need a sophisticated system of automated provisioning and de-provisioning of servers to deal with fluctuations in demand. If significant and relatively unpredictable, then this capability may be appropriate. If the need is regular and predictable, straightforward automation may be sufficient for your purposes. Among other questions you need to ask: Which applications are priorities, and which can be pushed farther back in the queue?

- **Pooled resources:** Consolidating resources to save on infrastructure, platform, and/or software costs is a common goal for large-scale IT operations. If you’re in a large organization with several independent departments potentially with their own IT operations (a common model for governmental institutions), you are likely to encounter critical questions of process: How, for example, will wholly independent groups deal with the inherent limitations of shared infrastructure and services? Will standardized configurations come at the cost of the optimized systems to which they’ve grown accustomed? As you move forward in the process of pooling your resources to get the attendant benefits, you need to be cognizant of the likely tradeoffs in putting everyone on a standard set of services. It may very well be worth the cost to the organization as a whole, but it may not seem that way to those who lose capabilities or levels of service to which they’ve been accustomed.
- **Broad network access:** As you move out from behind the corporate firewall and away from tightly controlled client configurations and network access scenarios, there are several important considerations that will need to inform your strategy, beyond the obvious concerns over security, such as the nature and extent of supportability: What kinds of personal devices will you support and to what degree? How will mobile clients—smartphones, tablets, or other kinds of devices with third-party operating systems, applications, etc.—access network resources, and will you have the requisite bandwidth to service them? What forms of authentication will you support?

Whichever objectives you are aiming to achieve, the important point to bear in mind is that building a private cloud is a process for which there are numerous tactical and strategic considerations, for which the *Core IO Maturity Model* can provide a useful framework. A successful implementation hinges on your ability to think through all facets of the undertaking, clearly understanding the dependencies, tradeoffs, limitations, and opportunities of any particular strategy.

4 The Importance of Managing Expectations

The *Core IO Maturity Model – Private Cloud View* is in some sense a response to an underlying issue that we’ve encountered in our experiences working with large organizations on their cloud deployments. Often, we have found that the ambitions and expectations that organizations have for a cloud solution are out of line with the realities of the process and its potential outcomes. To help these organizations

set more appropriate expectations and frame strategies with realistic and achievable goals, the maturity model can be very useful. But its lessons must also find their way to the organization as a whole in order to ensure the alignment of business and IT on the goals of the project.

4.1 *Service Providers and Service Consumers*

A persistent perception of cloud computing is that it can very rapidly deliver significant return to a business with the potential for a minimum of effort and investment. It's a common expectation among consumers that the cloud should offer easy-to-implement, extremely scalable, well performing, and secure solutions. Internet-based email services have benefited consumers for years, and it's become commonplace for individuals to be able to collaborate, store personal data, or host applications large or small in the cloud. If consumers can benefit from all these on-demand services, why should it be any different for businesses? Why does it take a week or more to instantiate a new email account within a corporate firewall, when Hotmail or similar services in the cloud—which cost the consumer little or nothing—can be had almost instantly?

The answer may or may not be self-evident, depending on your point of view:

- **The provider's perspective:** If you happen to work in IT, and are steeped in the technical and organizational challenges of delivering services to consumers within your organization, and understand in some measure the complexities of the systems involved in terms of both the technical and non-technical requirements imposed on them (such as data integrity, regulatory compliance, corporate information security, etc.), it will seem obvious that a public cloud-type solution would be highly problematic.
- **The consumer's perspective:** If, however, you happen to be a consumer of IT services, your perspective instead may be one of an inscrutably complex array of infrastructure components and software services that are often unreliable, unresponsive to change, and difficult to customize and control. And if you control budget, you may even question why you're spending so much money for such a disappointing return on investment. This point of view has the effect of putting IT, fairly or unfairly, in the challenging position of having to compete directly with large public cloud providers and offer services of comparable scale, reach, and efficiency.

This misalignment of perceptions and expectations between service providers and service consumers is telling, as it underscores the challenge of cloud adoption for large organizations in particular.

Consumers' personal experience of a public cloud has unfortunately contributed to the expectation that businesses can leap ahead, potentially in one step, from a traditional IT environment to one that is dynamic, responsive to changing business needs, highly available, and extremely cost-effective.

4.2 *Aligning Business and IT*

But the reality for most large organizations is that an incremental strategy is the only realistic path, given the technical and organizational complexity of current IT operations fed by years of layered investments that businesses are justifiably reluctant to abandon wholesale. And this is one of the key reasons why a private cloud model is proving especially compelling for those who are evaluating their cloud options in the context of enterprise IT: it takes into account existing investments and organizational aptitudes, permitting incremental approaches to adoption.

Be that as it may, the tension between unrealistic expectations and realities of cloud computing in an enterprise IT context can prove a challenge to resolve. Many IT leaders understand why an incremental approach is needed, but those outside IT—business leaders in particular—are less clear about the real implications of implementing a cloud solution. A sound strategy for achieving your objectives must also include an appropriate communications strategy for the sake of consistently setting and managing expectations for the organization as a whole. With the whole organization informed, from the board room to the front office, the hard work of defining and executing on your private cloud strategy is far more likely to achieve its objectives and set your business on the path to long-term success in the cloud.

5 Next Steps on Your Journey

For many corporations and government bodies seeking to improve business agility, reduce IT costs, and/or enhance service quality, a private cloud architecture—regardless of the service model—is the best choice. If your organization is evaluating a private cloud deployment, or even if it has already begun the process, there are many resources and programs available from Microsoft that will help you get and stay on track. Below are next steps you can take to deepen your knowledge of the various technical and organizational aspects of a private cloud deployment project, along with a set of recommended resources for further reading.

5.1 *Next Steps*

- Evaluate deploying your own private cloud using the [Hyper-V Cloud Deployment Guides](#), a set of best practices and tools designed by Microsoft Services to assist in deploying a private cloud based on Windows Server 2008 R2 Hyper-V and System Center datacenter solutions.
- Get on the [Microsoft Hyper-V Cloud Fast Track](#), which provides pre-validated private cloud configurations for compute, network, and storage, combined with hardware and software components available from select OEM partners.
- Review the [Datacenter Services](#) solution available from Microsoft Services, which encompasses four scenarios for IT transformation including Hyper-V Cloud services designed for private cloud deployment.

-
- [Contact Microsoft Services](#) to find out more about its Datacenter Services for private cloud or to engage its experienced consultants in your private cloud deployment project.

5.2 Further Reading

- [Are Your People Cloud-Ready? Why People and Process Trump Technology in the Transformation of IT](#)
- [Surviving Turbulent Times: Prioritizing IT Initiatives Using Business Architecture](#)
- [A Business Manifesto for Cutting Costs and Boosting Innovation](#)
- [Cloud Business Case: A View from the Field](#) [PDF]
- [IT Skills: Breaking Down the IT/Business Divide](#)
- [Cloud Computing: Architecting a Microsoft Private Cloud](#)
- [TechNet: Microsoft IT Showcase – Cloud Computing](#)
- [TechNet: Microsoft IT Enterprise Architecture and the Cloud](#)

6 Appendix: Table of Attributes: Core IO Maturity Model–Private Cloud View

A detailed description of the *Core IO Maturity Model – Private Cloud View* is provided in the table appended below for convenient reference. Note that bulleted items in **bold-face** represent those attributes that constitute the minimum requirements for a private cloud.

	Basic	Standardized	Rationalized	Dynamic
Self-Service	<ul style="list-style-type: none"> Scripted processes used to deploy applications consistently 	<ul style="list-style-type: none"> A defined, end-to-end server deployment process that includes application-compatibility testing. 	<ul style="list-style-type: none"> An automated, end-to-end server deployment process with task sequencing 	<ul style="list-style-type: none"> An automated, self-service deployment portal for requesting new IT services for deployment across virtual or private cloud environments
Usage-based	<ul style="list-style-type: none"> No chargeback process 	<ul style="list-style-type: none"> Limited process in place to assign costs for IT services back to business groups. 	<ul style="list-style-type: none"> Per-project chargeback based on hardware & software acquisition costs 	<ul style="list-style-type: none"> Real-time monitoring with chargeback to business groups using IT services Allocation- or utilization-based reporting using consumer tenant-service model Resource usage monitored, controlled & reported, providing transparency to both provider & consumer Ability to integrate directly with financial billing systems

	Basic	Standardized	Rationalized	Dynamic
Elastic	<ul style="list-style-type: none"> ▪ No defined set of core standard images ▪ No capacity planning process for servers ▪ No high-availability strategy 	<ul style="list-style-type: none"> ▪ A defined set of core standard images ▪ Manual capacity management processes for consolidation of business, services, resources, and servers 	<ul style="list-style-type: none"> ▪ Images maintained at current update levels and available for deployment in physical and virtual environments ▪ Virtualization & management used to move applications & services dynamically when issues arise with datacenter resource ▪ Monitoring of current resource needs in real time, utilization is compared to existing capacity ▪ Use-analysis tools used to predict impact of change; multiple levels of service-availability clustering or load-balancing 	<ul style="list-style-type: none"> ▪ OS & workload images maintained in centralized library for provisioning new physical & virtual systems on demand for datacenter services ▪ Dynamic, on-demand, real-time server provisioning & de-provisioning ▪ Application models & virtualization allow services to consume capacity on demand dynamically ▪ Services available during complete site outage by means of geo-clustering, cloud resources & automation
Pooled Resources	<ul style="list-style-type: none"> ▪ No server virtualization or server virtualization used only in test environments ▪ Virtualized layers not offered as service within company 	<ul style="list-style-type: none"> ▪ Organization actively uses virtualization to consolidate servers for production workloads ▪ Virtualized server pool offered as services 	<ul style="list-style-type: none"> ▪ IT system-aware monitoring, reporting in place with basic remediation ▪ Consolidated views of all management tools, consistent reporting across heterogeneous systems ▪ Consolidated view of and management process for heterogeneous virtual environments ▪ Majority of production server resources virtualized ▪ Virtualized server pool offered as a service 	<ul style="list-style-type: none"> ▪ Virtualization used to manage resource allocation dynamically for workloads, including moving services from server to server based on resource needs or business rules ▪ Storage allocated dynamically from available pool of physical space, based on capacity required and within limits set by policy quotas ▪ Majority of production server resources virtualized ▪ Virtualized server pool offered as a service ▪ IT system-aware monitoring & reporting with full remediation and automated disaster recovery ▪ Service performance & health monitoring provide advanced management of datacenter resource usage

	Basic	Standardized	Rationalized	Dynamic
Broad Network Access	<ul style="list-style-type: none"> No network security, or it is limited to basic perimeter firewall; remote client access limited or non-existent 	<ul style="list-style-type: none"> Multiple vendor products used for firewall, IPS, Web security, gateway anti-virus, and URL filtering Secure, standardized remote access available to end users across organization 	<ul style="list-style-type: none"> Integrated perimeter firewall, IPS, Web security, gateway anti-virus, and URL filtering deployed with support for server and domain isolation Secure remote access integrated with quarantine for compliance with corporate policy 	<ul style="list-style-type: none"> Network security automated & proactive, with centralized alerting & reporting to meet network protection service-level agreements Network security, alerts & compliance integrated with all other company tools to provide a complete, company-wide scorecard & threat assessment including both private and public cloud environments Secure, remote access nearly always available with bidirectional connectivity Policy-based access to networks & applications