Planning the Migration of Enterprise Applications to the Cloud

A Guide to Your Migration Options: Private and Public Clouds, Application Evaluation Criteria, and Application Migration Best Practices

Introduction

Cloud computing—IT resources and services that are abstracted from the underlying infrastructure and provided on demand and at scale in a shared multitenant and elastic environment—represents a paradigm shift from which both enterprise IT and service providers can benefit.

According to the definition of cloud computing from the National Institute of Standards and Technology (NIST), IT services that are delivered as cloud services offer:

- A pay-as-you-go model with minimal or no initial costs
- Usage-based pricing, so that costs are based on actual usage
- Elasticity, so that users can dynamically consume more or less resources
- Location independence, high availability, and fault tolerance
- Ubiquitous access to services, where users can access services from any location using any form factor

A cloud can provide IT infrastructure services such as servers, storage, network and network services, or infrastructure as a service (IaaS); an application deployment platform with application services such as databases, or platform as a service (PaaS); or subscription-based software applications, or software as a service (SaaS).

Today, service providers, who already excel at provisioning, managing, and scaling services for multiple customers, are providing offerings based on IaaS in which the enterprise uses the pay-as-you-go compute infrastructure from the service provider. A cloud provided by a service provider is known as a public cloud.

In addition, some enterprises are choosing to build a private cloud—enterprise IT infrastructure services, managed by the enterprise, with cloud computing qualities: self-service, pay-as-you-go chargeback, on-demand provisioning, and the appearance of infinite scalability.

© 2010 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

Whether they are considering a private cloud or a public cloud as their services model, enterprises first must consider which of their many applications belong in the cloud and how to migrate them. In the public cloud, service providers also might want to assist enterprises in making the best possible decision about migrating applications. The applications chosen for migration and how you approach the application migration process can affect not only the ease and success of the migration itself but also the user's service experience.

Let us take a look at what is involved in application migration, as well as the business and technical factors behind a decision to migrate applications to a cloud model.

Application Migration and Clouds

Today enterprise data centers have a number of applications that are implementing on existing infrastructure with known scalability or resiliency issues. Although migration of an application to a cloud will not solve all intrinsic application scalability or resiliency issues, enterprises might be able to derive tangible financial and operational benefits in moving an application from legacy servers to a cloud.

Application migration is the process of redeploying an application, typically on newer platforms and infrastructure. The process involves the staging of the new environment before the actual cutover and requires coordination of IT teams at the time of cutover. If the migration is on a compatible platform, the application does not need to be recompiled.

In the case of the cloud, the application can be migrated from an existing data center to the target cloud. The target infrastructure can be a public, private, or hybrid cloud: that is, an environment that transparently combines multiple clouds, whether private or public. Additionally, the migration can involve a physical-to-virtual (P2V) migration if the existing application is not running on a virtualized platform.

To identify applications for migration to a cloud, it is necessary to first identify and understand the business and technical factors for the migration. Reducing costs and business agility are typical business factors for application migration to clouds. Cloud computing can provide significant cost savings because of the increased utilization resulting from the pooling of resources and the standardization and automation required for cloud services.

Cloud computing enables rapid delivery of IT services, which increases business efficiency. Ordering processes, which could have previously taken multiple weeks and involved multiple departments, are now reduced to a few mouse clicks. This increased IT efficiency translates to overall business efficiency and has the potential to unleash new innovations and opportunities.

On the operational side, manageability, performance, and scalability are the typical reasons why businesses consider cloud computing. By delegating the management of infrastructure and software platforms to a cloud service provider, customers can offload operational responsibilities to service providers. In an enterprise private cloud model, a common cloud IT management team can also help provide a similar service.

A cloud computing environment might offer increased resources, which can lead to performance improvements for certain applications. Applications that are designed to spread their workload across multiple servers will be able to benefit from automated scaling of resources to match the current demand. This is especially appealing for applications with unpredictable or cyclical usage patterns, because a cloud orchestrator can monitor usage and can dynamically scale resources up or down. This behavior, combined with the pay-by-usage characteristic of a cloud, can lead to significant financial savings.

Application Migration Options

After an application has been identified as a candidate for cloud migration, based on business and technical factors, it is necessary to consider for what type of cloud environment—SaaS, PaaS, or IaaS—the application is best suited.

Software as a Service

Based on the type of application, and if SaaS-based alternatives exist, it is worth considering if the SaaS alternatives can meet both business and technical needs. Such a change is no longer an application migration but more of a replacement of the existing application with a SaaS option. There might still be a need to migrate existing data to the new application.

SaaS removes the need to manage both the application and the infrastructure on which the application is deployed. This approach can be attractive, but certain criteria, such as service-level agreements (SLAs), data portability, and long-term costs, must be carefully evaluated when considering an SaaS deployment.

- SLAs: The SaaS vendor should provide a SLA for application overall availability, scalability, and performance, as well as provide clear polices and guidelines for application maintenance and upgrade windows.
- Data portability: The SaaS vendor should provide a way to allow customers to own and control their application data. SaaS customers should have the ability to export all application data that belongs to them, in a format that can be easily parsed and migrated to other internal or external applications.
- Long-term costs: The SaaS model can be financially attractive for a small business or enterprise that has not made any significant investment in its own data center, because of the low initial costs. As the number of users and the period of ownership increase, however, it is necessary to carefully compare longer-term recurring costs over time against amortized cost of ownership. As with other leasing-based financial models, the overall costs of ownership might be greater for a lease as time and usage increase.
- User management: Enterprise users are typically managed using directory services such as Microsoft's Active Directory or other Lightweight Directory Access Protocol (LDAP)–based services. As user accounts are added, deleted, or modified, most SaaS applications will not be automatically updated with these changes, thus creating additional work and potential security risks for IT administrators.
- Security: The SaaS application can contain sensitive corporate data when stored on the SaaS provider's infrastructure. You should require transparency in the service provider's security policies to be able to determine whether adequate security is provided, based on the nature of application data.

Platform as a Service

Platform as a service might be an option for migrating business applications that are based on standard application server software such as Java EE 5 or Microsoft's .NET platform.

In this model, the service provider manages the application platform software and might provide access to common application services such as SQL databases. The application platform might be shared by multiple applications belonging to different customers. How the application platform is mapped to the physical infrastructure is typically controlled by the cloud service provider.

The decision factors in such a migration will depend on the type and version of the application server used. Some PaaS environments might not support all features of the application server and might require application changes.

© 2010 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

All of the same criteria used for considering a SaaS deployment, including SLAs, data portability, long-term costs, user management, and security, should be considered for a PaaS migration.

- **SLAs:** A PaaS vendor should provide SLAs for application platform availability and performance. The cloud service provider should also provide clear policies and guidelines for maintenance and version management of the platform and policies for version compatibility for APIs between the platform and the application.
- **Data portability:** In a PaaS model, the application data is typically stored in a database provided by the cloud service provider. The customer must be able to export data in a format that can be migrated to other databases.
- **Long-term costs:** The financial model for a PaaS should be compared against those of an internal deployment of the infrastructure and the application server/platform using IaaS and deploying the application server using the cloud-based servers.
- **User management:** A PaaS application will require administrative and application user accounts. For both account types, customers should understand how the user management aligns with their existing directory services and user management processes.
- **Security:** In a PaaS deployment, the same application server might host applications from different customers. In such an environment, additional security is necessary to make sure that rogue applications are not able to exploit vulnerabilities in the platform software to affect other applications.

When evaluating PaaS, enterprises should also consider platform management and scalability.

- **Platform management:** Application servers provide management consoles and tools for monitoring and management of applications that are deployed on them. A PaaS deployment should allow customers to use similar tools to manage and tune their applications.
- **Platform scalability:** A PaaS environment might offer dynamic scaling (up or down) as an optional feature, based on the capabilities of the underlying application server. Dynamic scaling works well with applications where the user load is nondeterministic, such as for consumer Internet applications. If this feature is used, the PaaS provider should clearly indicate how the application will be scaled up or scaled down and how contention for resources will be handled.

Infrastructure as a Service

Migration of an application to an IaaS involves deploying the application on the cloud service provider's servers.

The initial step in making a decision to migrate to an IaaS model is to determine whether the cloud-based server hardware and operating system (OS) are compatible with the current server's hardware and OS. For example, if an application is running on an x86 server, the cloud servers must be able to implement x86 instructions. If the hardware is not compatible, the application might need to be recompiled or redeployed for the new platform. Similarly, if the OS is compatible, few changes will be required when the application is migrated.

Once again, most of the same criteria that are used for considering an SaaS or PaaS application migration, including SLAs, data portability, long-term costs, user management, and security, should be considered for the IaaS migration.

- **SLAs:** For an IaaS deployment, you need a SLA for the availability and performance of the server, network, and storage infrastructure. You also should receive information about the maintenance and management procedures for the infrastructure and how any potential downtime is handled.
- **Data portability:** In an IaaS deployment, the application might be using a database server that is also deployed in the cloud. In these cases, the cloud service provider must provide a way to replicate or migrate the block or file storage that is used by the database server.
- **Long-term costs:** The cost of an IaaS application should be compared against the cost of deploying that application on enterprise servers. In some cases, a public cloud IaaS deployment might have obvious benefits because of dynamic scaling and usage-based pricing. For always-on applications, based on the computing resources required, the longer term cost of ownership in a public cloud might actually be greater than the cost of ownership in a private cloud.

© 2010 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

- **User management:** In an IaaS model, there might be up to three different user roles:
  - Server administrator
  - Application administrator
  - Application user

The user management procedures and tools for each of these roles should be evaluated.

- **Security:** In an IaaS deployment, virtual machines belonging to different customers might be implementing on shared physical infrastructure. When considering an application migration, the cloud service provider's security policies for virtual and physical isolation as well as compliance should be examined. The cloud service provider should allow auditing of security and compliance policies, using emerging technologies and specifications such as those of the CloudAudit forum.
- **Scalability:** Applications that are designed to scale out will benefit from dynamic scaling features in a cloud. Typically, these applications are multitiered and have request load-balancing features, such that a pool of stateless application servers can be dynamically scaled up or down. If you intend to use this feature, the cloud service provider should provide clear policies on how this type of scaling will function and how resource contention across customers and applications is handled.

Public or Private Clouds

Cloud service providers can offer advanced enterprise-grade capabilities, for security and performance, combined with network services such as load balancing and WAN optimization services. After you have evaluated these aspects of a cloud service provider's offerings, you need to consider whether it is always a good strategy to migrate applications to a public cloud. Not every application is suitable for migration to a public cloud. When deciding whether to choose a public or private cloud for an application migration, you need to examine the WAN traffic architecture, data security and management, legacy application integration, and security and compliance.

- **WAN traffic:** If an application workload is traffic intensive and communicates with other data center resources or applications, migration to a public cloud will not be optimal because of WAN bandwidth costs and potential performance effects.
- **Data security and management:** Applications might depend on a shared data center storage or other resources such as directory services for user, profile, and data management. When evaluating a migration, such dependencies need to be identified and addressed.
- **Legacy application integration:** Legacy business applications running on platforms such as a mainframe and AS400 that might be tightly integrated with other business applications might be a risk for migration to public cloud offers. Applications dependent on the legacy applications are definitely a good candidate for private cloud migration.
- **Security and compliance needs:** In a private cloud, an enterprise will have more direct influence over the infrastructure architecture and operational policies. This control can lead to efficiencies because of customization and optimization, which need to be evaluated against the additional costs.

© 2010 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

Application Evaluation Criteria

After a type of cloud environment and the cloud service delivery model are chosen, the candidate application needs to be further evaluated to make sure that it is feasible to migrate and also to prepare for the migration process itself.

Application Architectures

The application architecture will affect how an application can be migrated to cloud environments and sometimes whether an application is suitable for migration. The next sections discuss common architecture patterns and how they affect cloud migration.

Multitiered Applications

The majority of enterprise applications are built using multiple tiers to decouple the major functions and modules in the system. One such approach is to organize the application using three tiers as follows:

- A data management tier, which consists of relational or other database components
- A business logic tier, which uses application platform or containers, such as Java EE or Microsoft's .NET
- A presentation tier, which is responsible for interfacing with the user interfaces or other external systems, including managing state and data for presentation to these external systems

Applications that use a layered architecture approach have well-defined interfaces between these layers. Based on application usage patterns, it might be possible to migrate application tiers, or modules within a tier, separately. For example, in a web application, static content can be migrated to a content delivery network provider to allow parts of a website to load more quickly. In other cases, WAN bandwidth restrictions might prevent tiers from being separated, and all layers of the application will be migrated to a cloud. In either case, each layer and its major distributed components modules should be evaluated separately for how it should be sized and migrated to a cloud.

Application tiers might also have varying security and zoning requirements. For example, some application data might have to be secured behind a firewall.

Scale-Up and Scale-Out Architectures

A "scale-up" architecture is one where the application can benefit from more resources, such as CPU and memory, being added to a single server or node. In contrast, a "scale-out" architecture is one where an application scales by additional nodes being made available for the workload; that is, it scales horizontally.

Scale-out applications can take advantage of the pay-by-usage cost model of the cloud. When there are increased requests for an application, more nodes can be deployed to handle the increased load. When the requests slow down, the additional nodes can be powered off to reduce costs.

Today, it is not possible to dynamically scale up an application running on a single machine instance. This might change in the future, because virtualization systems are starting to support hot-plug features, where more memory and CPU can be added dynamically.

© 2010 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

Geographic Access

Some applications are used from a within a single geographic location, while others might be used from multiple locations or even worldwide. For example, consumer-facing Internet applications often are used in multiple regions of the world in an unpredictable manner. Since enterprises also have multiple corporate locations, they might benefit from allowing applications to be distributed closer to points of access.

When migrating an application to a cloud, it is necessary to consider the locations from which the application will be accessed. Most global cloud providers will provide a mecha-nism to configure the expected access for an application by purchasing capacity in dif-ferent regions or geographic zones. In other cases, application access might need to be blocked or disallowed in a region, often for regulatory or security purposes.

When selecting a cloud, it is important to consider these factors and make sure that geo-graphic access can be monitored, controlled, and optimized.

Cisco Data Center Assessment Service for Application Dependency Mapping, formerly known as Application Dependency Mapping

Cisco Data Center Assessment Service for Application Dependency Mapping (Figure 1) identifies dependencies between applications and their dependencies on a shared data center infrastructure.

The data collection can be done in multiple passes or phases: first to identify all of the appli-cation's immediate dependencies and then to also to identify what other applications are dependent on the applications' dependencies. For example, if both application A and appli-cation B are using the same database server, this needs to be identified so that the migra-tion plan can include a combined move or can include steps to split the dependencies.

Cisco Data Center Assessment Service for Application Dependency Mapping tools can also produce server and application affinity maps, based on the observed traffic across applications, which is useful to get a sense of which applications might need to be migrated as a batch.

Figure 1. Sample Cisco Data Center Assessment Service for Application Dependency Mapping

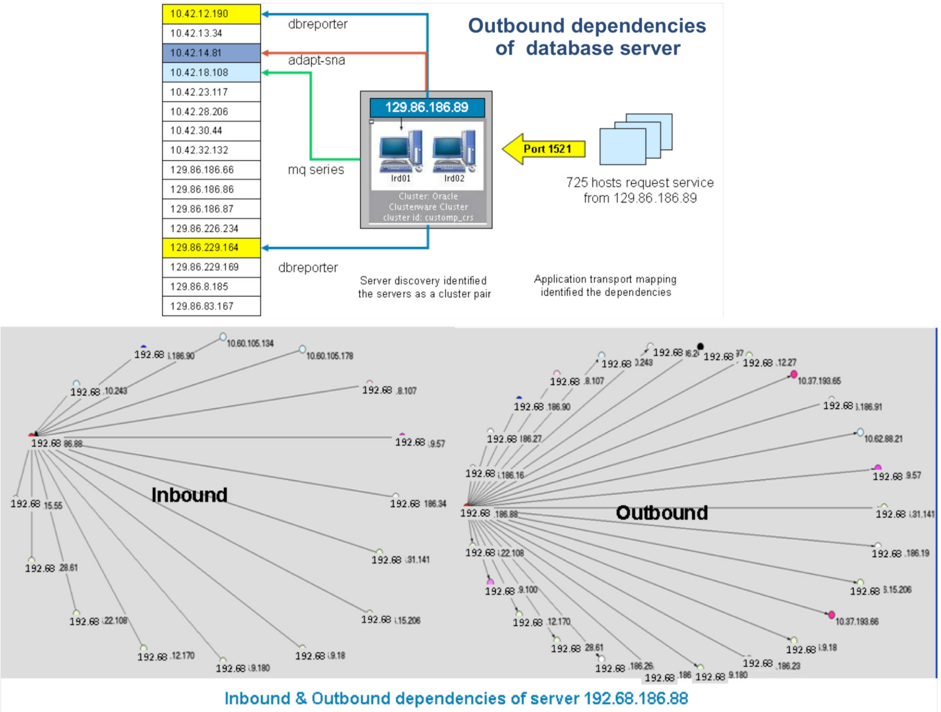

Application Profiling

Application profiling is used to measure and collect real usage data of an application before it is migrated. This data can help size the application deployment in a cloud. Ideally, application data should be collected for at least 10 to 15 days to allow capture of variances in daily and weekly usage patterns.

For each node on which the application runs, the following data should be collected:

- CPU usage
- Memory usage
- Storage data such as throughput, latency, and input/output operations per second (IOPS)
- Network data such as throughput, connections per second, and dropped connections

The node-level data can be used to estimate how many machines, and what type of machines, will be necessary when the application is migrated.

In addition to node-level statistics, it is also important to profile user activity, such as the total number of connected users, request and transaction rates, and request latencies. The usage data can also be used to build automated tests for the application to make sure of the same or an improved level of service after the application is migrated.

The node data, along with application usage data, can also provide an initial estimate of the costs of the cloud resources.

Application Migration Planning and Testing

The final step in preparation for the migration of an application to the cloud is to develop a work plan and tests for the actual migration. Based on the type of application and its business continuity requirements, the amount of planning required can vary. If downtime is not accept-able or needs to be minimized, the application should be migrated in phases, with both the existing and migrated applications available for some period of time. After initial tests, users can also be migrated in batches, and cloud capacity can be increased over time.

The application data collected in previous steps can be used for creating test suites and simulating different user types and loads. Having a test plan and automated tests will help make sure of a successful migration.

Designing the Solution—The Case for Professional Services

Evaluating applications, designing migration plans, and migrating applications to a targeted cloud computing model or models is a demanding task. It requires detailed application migration process design experience and deep understanding of cloud computing models as well as detailed knowledge of how the applications interact with both each other, the external (cloud) environment, and the underlying infrastructure. It also requires experience integrating IT systems and cloud management, and—being a significant project in its own right—a structured approach to program management.

Cisco Services has significant expertise and experience in each of these areas, backed up by multiple years of helping customers transform their data centers. Cisco Services offers a range of Cloud Enablement Services, from Cisco Data Center Strategy and Analysis Service for Cloud, formerly known as Cloud Strategy, to Cisco Data Center Virtualization Design Services for Cloud, formerly known as Cloud Planning and Design, to Cisco Data Center Build Services for Cloud, formerly known as Cloud Implementation. Application migration is an integral component in each of these services, and detailed application methodologies have already been developed by the cloud computing experts in Cisco Services. Cisco Services, therefore, are ideally placed to help you design and develop a state-of-the-art application migration approach for your cloud computing deployment.

© 2010 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

Summary

There are many different aspects to consider when selecting, evaluating, and planning the migration of a data center application to a cloud. The process begins with analysis of the factors for the application and a comparison of these factors to different types of cloud computing environments. Your next steps are to analyze and measure application details that help in building a migration plan, as well as a plan for testing each phase of the migration. This process (Figure 2) often is iterative, because data might be uncovered that leads to the reevaluation of the results in prior phases. Although no single approach will work for all applications, the best practices recommended here will help in determining application suitability and performing a smooth migration. Finally, Cisco Services are ideally placed to help you devise and develop an advanced application migration approach that will help justify your investment in cloud computing.

Figure 2. Application Migration Process

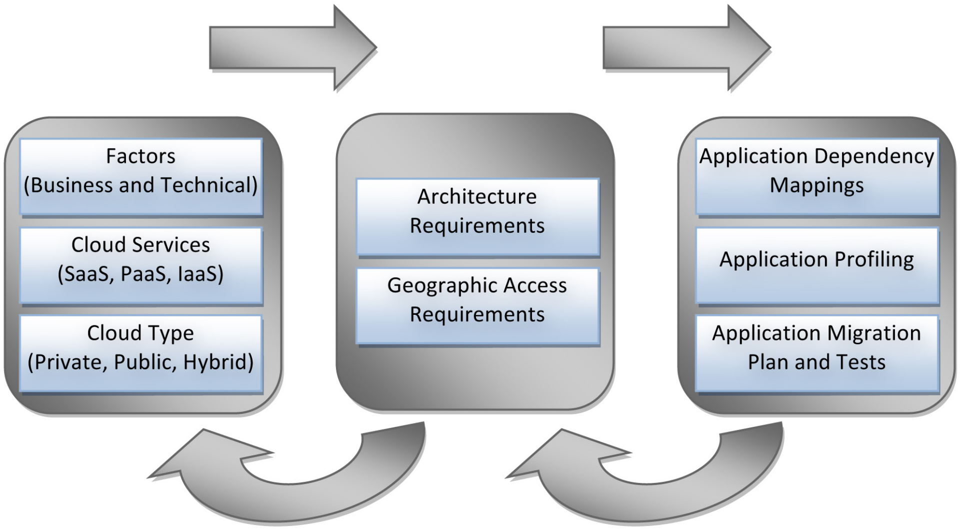
Factors (Business and Technical)
Cloud Services (SaaS, PaaS, IaaS)
Cloud Type (Private, Public, Hybrid)
Architecture Requirements
Geographic Access Requirements
Application Dependency Mappings
Application Profiling
Application Migration Plan and Tests

For more information about Cisco® cloud computing, visit www.cisco.com/go/cloud enablement.

For more information on Cisco Data Center Assessment Service for Application Dependency Mapping, visit www.cisco.com/en/US/products/ps10437/services_segment_service_home.html.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

Printed in the USA

C11-618438-00 08/10