



# **VMware vCloud<sup>®</sup> Architecture Toolkit**

## **Introduction**

**Version 3.0**

**September 2012**

© 2012 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. This product is covered by one or more patents listed at <http://www.vmware.com/download/patents.html>.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

VMware, Inc.  
3401 Hillview Ave  
Palo Alto, CA 94304  
[www.vmware.com](http://www.vmware.com)

## Contents

1. Overview .....	5
2. Using the vCAT Documentation Set.....	6
2.1 Recommended Reading Order .....	9
3. Cloud Computing and VMware vCloud .....	11
3.1 VMware vCloud Requirements .....	11
3.2 VMware Alignment to Standards .....	12
3.3 vCloud Definitions .....	13
3.4 Solution Area to Technology Mapping .....	14
4. Journey to a Mature vCloud Implementation .....	17



## 1. Overview

A *reference architecture* is an architecture template solution that addresses one use case in a particular domain. The VMware vCloud® Architecture Toolkit (vCAT) provides modular components and documents for multiple use cases, including design considerations and design patterns to support architects, operators, and consumers of cloud computing solutions based on VMware technologies.

vCAT is vendor-agnostic, but it does share vendor details when providing implementation examples. Vendors provide information about the use of their products with vCloud, including integration with vCAT, on the VMware Solutions Exchange (<https://solutionexchange.vmware.com/store>).

vCAT design guidelines cover multiple use cases. Instead of referring to *best practices* (a term subject to misinterpretation because best practices depend on use cases and are subject to many variables, including change over time), vCAT provides *design guidelines*. Architects must determine which design guidelines apply to the requirements, constraints, and characteristics of their projects and chosen technologies. When using the toolkit, consider the use case that best applies to your situation, and choose the design guidelines that support your design implementation.

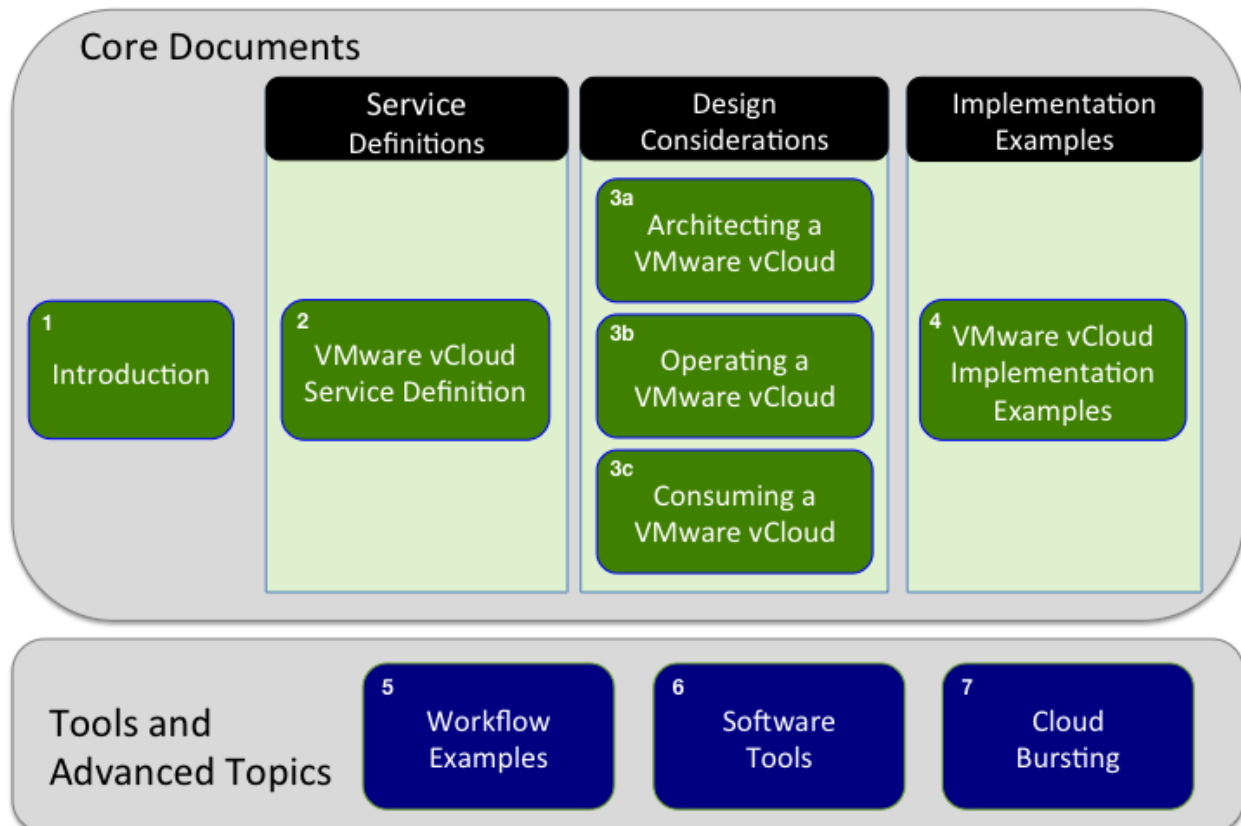
This document covers the following topics:

- Using the vCAT Documentation Set.
- Cloud Computing and VMware vCloud.
- Journey to a Mature vCloud Implementation.

## 2. Using the vCAT Documentation Set

The vCloud Architecture Toolkit provides a set of documents to support the design of complex, integrated, reference architectures for architects, operators, and consumers. The documents are shown in Figure 1 and are briefly described in Table 1.

**Figure 1. VMware vCloud Architecture Toolkit Document Map**



**Table 1. vCAT Documents**

Document	Description	Audience
<i>VMware vCloud Architecture Toolkit Release Notes</i>	Information about the VMware Architecture Toolkit, toolkit packages, how to use the documentation center, and information about changes since the vCAT 2.x release.	All
<i>VMware vCloud Architecture Toolkit Introduction</i>	<p>This document covers the following topics:</p> <ul style="list-style-type: none"> <li>• A brief summary of vCAT documents.</li> <li>• Suggested reading order depending on audience/role.</li> <li>• Introduction to cloud computing and basic cloud computing requirements and definitions</li> <li>• List of references.</li> </ul>	All
<i>VMware vCloud Architecture Toolkit Service Definitions</i>	Service definition lifecycles, including specific considerations for private, public, and hybrid vCloud instances, and examples of service offerings designed to help you create service definitions that meet specific business objectives.	All
<i>VMware vCloud Architecture Toolkit Architecting a VMware vCloud</i>	Design considerations for architecting and building a VMware vCloud, including the basis for a reference architecture and guidance on requirements for implementing a VMware vCloud infrastructure.	Architects
<i>VMware vCloud Architecture Toolkit Operating a VMware vCloud</i>	Introduction to high-level operational areas and the evolution to support vCloud dynamics. Provides information about the operational procedures, roles, and responsibilities, setup, management, and monitoring of a vCloud. Also covers VMware management tools that support vCloud operations.	IT Operations
<i>VMware vCloud Architecture Toolkit Consuming a VMware vCloud</i>	<p>Answers to consumer questions, such as the following:</p> <ul style="list-style-type: none"> <li>• How do I handle the application lifecycle in a vCloud?</li> <li>• How do I protect my workloads?</li> <li>• How do I guarantee workload resource requirements are met?</li> </ul>	Consumers, end users
<i>VMware vCloud Architecture Toolkit Implementation Examples</i>	Provides examples of how to build a private or public enterprise vCloud.	Architects

Document	Description	Audience
<i>Workflow Examples</i>	Provides a description of useful scripts and workflows for VMware vCenter™ Orchestrator™. Other examples are given that use technologies such as PowerCLI. References to where these scripts can be found are included.	Architects, IT Operations
<i>VMware vCloud Architecture Toolkit Software Tools</i>	Information about software that can benefit architects and operators. Provides information about freely available technologies that have been created and used to assist in vCloud design, deployment, and operations. Also includes information about several powerful tools that are available only as part of a service engagement with VMware Professional Services or a VMware partner.	Architects, IT Operations
<i>VMware vCloud Architecture Toolkit Cloud Bursting</i>	Provides the theory behind auto-scaling an enterprise cloud environment by using multiple cloud locations, including those owned by an enterprise and/or a service provider. This theory leverages VMware technologies but applies to other cloud technologies as well. This material is based on VMware field experience with customers and service providers.	Architects

The following typographical conventions are used in all vCAT documents.

**Table 2. Document Typographical Conventions**

<i>Emphasis</i>	Emphasis (italics) is used to emphasize information, introduce new terms, and for document and workflow names.
<code>Command</code>	System commands, file names, and registry keys use this font.
<code>Code</code>	This font is used for code snippets and scripts.
<b>User Interface</b>	UI objects such as tabs, buttons, and field labels are in bold text.
<a href="#">Hyperlink</a>	Blue, underlined text indicates an active link (URL).



---

**Note, Caution**

Notes contain information related to the topic that is of possible interest to the reader.

Caution is used to highlight important information that describes potential problems or actions that might cause unexpected results. A Caution alerts the user, and may indicate the possibility of significant data loss.

---

## 2.1 Recommended Reading Order

The documents can be read in the order shown in the document map or in the order recommended for a particular audience or role such as the following

- *vCloud providers* who provide the vCloud infrastructure and services. An *architect* has overall control over how a solution is designed and implemented in the environment.
- *vCloud operators* who are responsible for operation of the cloud. *Operators* are involved with the day-to-day running and administration of the vCloud environment. They need to understand operational procedures and how the vCloud components fit together.
- *vCloud consumers* who utilize cloud provider resources for application deployment. A consumer (organization or individual) is someone who consumes vCloud resources. Consumers want to run their workloads in the vCloud environment without concern for the underlying infrastructure or day-to-day administration.

The following table identifies the most relevant documents for each audience.

**Table 3. vCAT Audience Reading Guidelines**

	Architect	Admin/Operator	Consumer
Introduction	◆	◆	◆
Service Definitions	◆	◆	◆
Architecting	◆	◆	
Operating	◆	◆	
Consuming	◆		◆
Implementation Examples	◆		
Workflow Examples	◆	◆	
Software Tools	◆	◆	
Cloud Bursting	◆		

### 3. Cloud Computing and VMware vCloud

Cloud computing leverages the efficient pooling of an on-demand, self-managed, virtual infrastructure that is consumed as a service. VMware vCloud is the VMware solution for cloud computing that enables delivery of Infrastructure as a Service (IaaS).

#### 3.1 VMware vCloud Requirements

According to National Institute of Standards and Technology (NIST), the key components of a cloud are on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service. VMware aligns with the definition of *cloud* as elastic, lightweight entry and exit, available over Internet protocols, and running on a shared infrastructure.

A cloud always starts with a shared, virtual infrastructure. If any resource is dedicated to only one customer, you have a *managed hosting platform*, not a cloud infrastructure. Similarly, it is not considered a cloud if there are manual procedures that must be performed by the cloud administrator or service provider to provision cloud resources following a consumer request.

The VMware vCloud blueprint follows these basic NIST requirements as the foundation for an IaaS cloud:

- A cloud must be built on a *pooled, virtual infrastructure*. Pools include not only CPU and memory resources but also storage, networking, and associated services.
- The cloud should provide *application mobility between clouds*, allowing the consumer to enter and leave the cloud easily with existing workloads. The ability to use existing consumer tools to migrate workloads to or from the cloud is highly desirable. Mobility of workloads between clouds requires cross-cloud resource management.
- The cloud should be *open and interoperable*, allowing the consumption of cloud resources over open, Internet-standard protocols. Access to cloud resources does not require any other specific network protocols or clients.
- Cloud consumers should pay only for resources they consume or commit to consuming.
- The cloud should be a secure, trusted location for running cloud consumer workloads.
- Cloud consumers should have the option and the ability to protect their cloud-based workloads from data loss.
- Cloud consumers are not responsible for the maintenance of any part of the shared infrastructure and do not need to interact with the cloud provider to maintain the infrastructure. They are not responsible for storage and network maintenance, ongoing cloud infrastructure patches, or business continuity activities. The cloud should be available to run high-availability workloads, and any faults occurring in the cloud infrastructure should be transparent to cloud consumers as a result of *built-in availability, scalability, security, and performance guarantees*.

## 3.2 VMware Alignment to Standards

VMware continues to develop technologies that align with evolving cloud standards as defined by NIST and other global standards organizations.

vCloud solutions focus on the following areas:

- People – People who develop solutions, architect the design, operate the implementation, and consume the resources. (See *Operating a VMware vCloud* and *Consuming a VMware vCloud*.)
- Process – Processes for architects, operators, and consumers.
- Technology – Alignment with successful design, deployment, and integration considerations. VMware technologies address the relevant areas within the standards.

Standards are still evolving for private, public, community, hybrid, and other types of clouds. vCAT focuses on the most common core design areas. The technology is the same, but operations and vCloud resource consumption varies according to the type of vCloud, the type of vCloud provider, and specific consumer requirements.

- A *private* vCloud is operated by an organization and secured behind a firewall.
- A *public* vCloud is generally accessible to users on the Internet.
- A *community* vCloud is a specific public vCloud use case where access is limited to specified groups who share a common set of requirements.
- A *hybrid* vCloud is characterized by a connection between multiple vCloud instances. Typically, there is a bridge between two private vCloud instances that has a dedicated and secured connection. The underlying network resides behind an Internet-facing firewall.

As cloud computing continues to evolve, there will be many cloud definitions. The information in this toolkit is a valuable aid in support of your vCloud projects, regardless of your chosen definition.

### 3.3 vCloud Definitions

vCAT uses the terms *private vCloud*, *public vCloud*, and *hybrid vCloud*, based on a specific set of definitions provided by NIST. A private cloud is sometimes known as an *internal vCloud*.

#### 3.3.1 Private vCloud

A *private vCloud* operates on private networks, where resources are accessible behind the firewall by a single company. In many cases, all the tenants share one legal entity. For example, a university might offer IaaS to its medical and business schools, or a company might do the same for various groups or business units. The private vCloud can be managed by the enterprise and hosted on premise or operated on a dedicated infrastructure provided by a vCloud service provider or systems integrator. In any case, a private vCloud must conform to the organizational security constraints.

#### 3.3.2 Public vCloud

A *public vCloud* offers IT resources as a service through external service providers and is shared across multiple organizations or the Internet. This can be viewed as a vCloud infrastructure that is operated by one organization for use by multiple, legally separated organizations.

A public vCloud is provisioned for open access and might be owned, managed, and operated by one or more entities.

A public vCloud provider might also support a private, community, or hybrid vCloud.

#### 3.3.3 Community vCloud

A *Community vCloud* is a specific public vCloud use case where the cloud is shared, and typically owned, by a group of organizations with a common set of requirements. In many cases, the organizations also include some level of legal separation. Community vCloud resources are shared, with some parts under central control and other parts with defined autonomy. A vCloud built for government, education, or healthcare might be an example of a community vCloud.

A community vCloud can be offered by a traditional service provider, by a member of the community, or by a third-party vendor and hosted on one or more sites. It can be placed on-premise at one or more of the organizations' sites, off-premise at a vCloud provider site, or both on- and off-premise.

#### 3.3.4 Hybrid vCloud

A *hybrid vCloud* combines the benefits of the private and the public vCloud, with flexibility and choice of deployment methods.

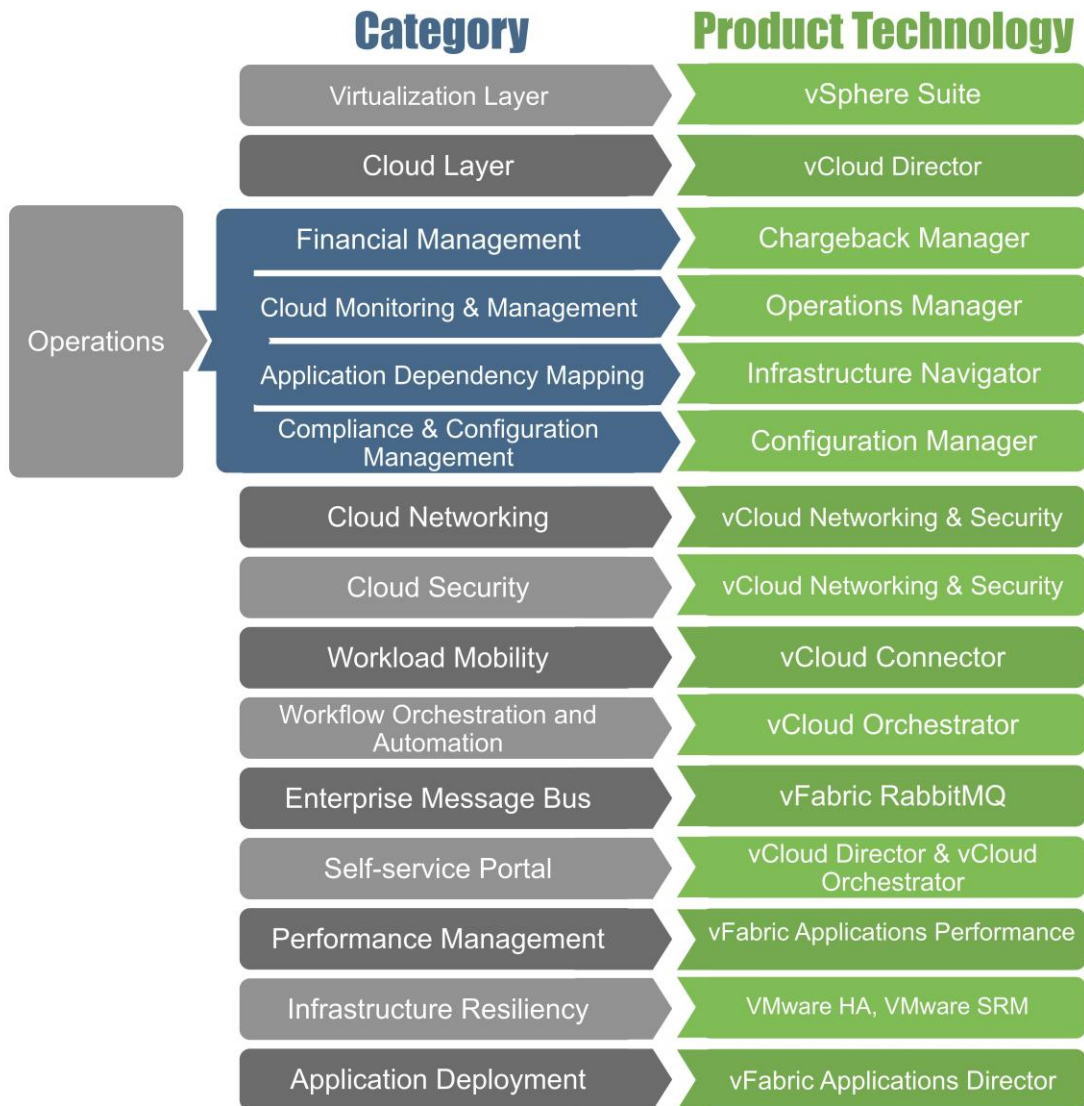
A hybrid vCloud consists of multiple, linked vCloud infrastructures. These distinct vCloud infrastructures can be private, community, or public, they but must meet a set of requirements defined by the providers and agreed to by the consumers. Connecting these vCloud instances requires data and application mobility as well as management.

When load-balancing between vCloud instances (*cloud bursting*), use a consistent monitoring and management approach when migrating an application or data workload.

### 3.4 Solution Area to Technology Mapping

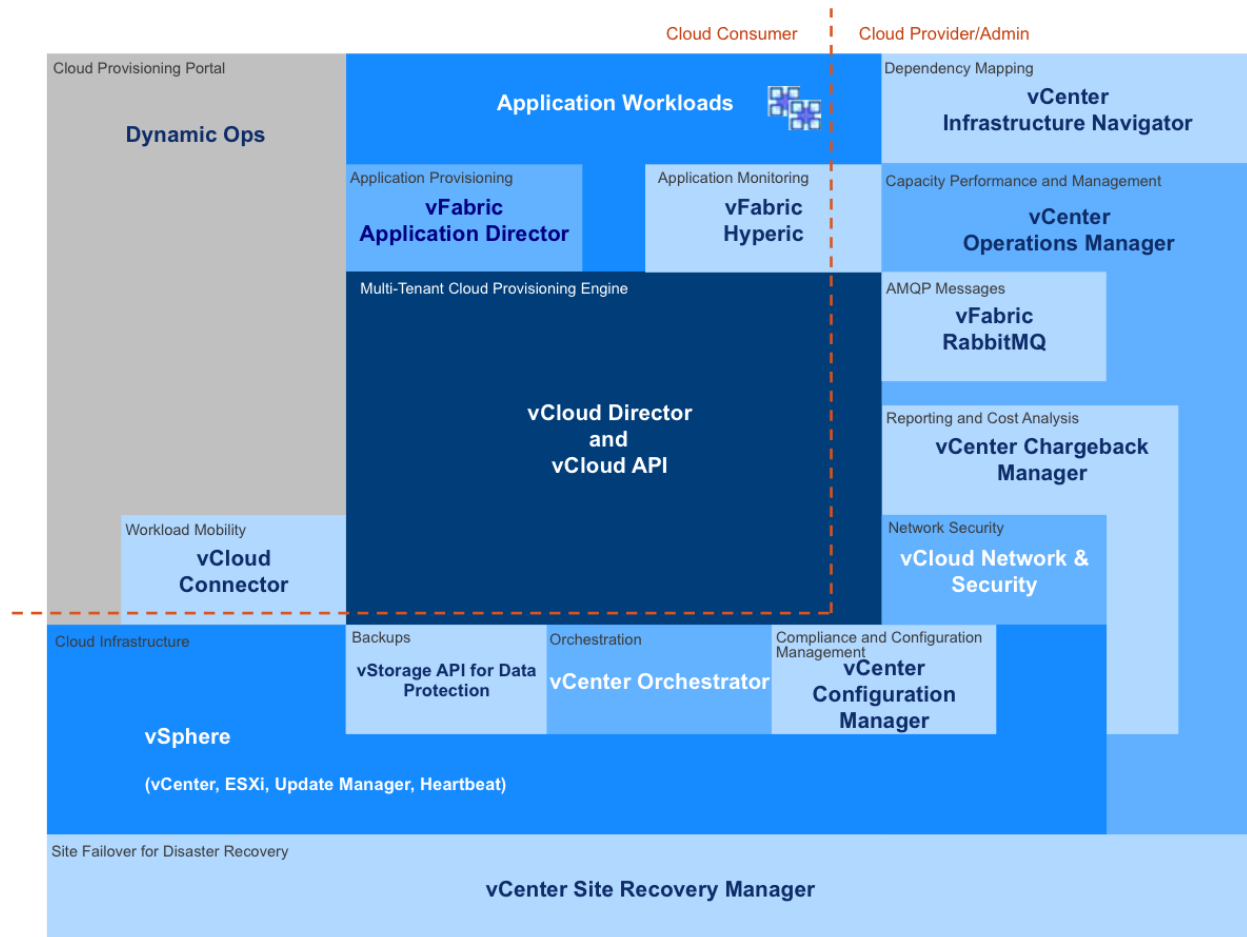
When considering various technology solutions for your vCloud architecture, evaluate the solution and operational requirements to provide justification for the proposed solution.

**Figure 2. Technology Areas**



The following figure provides additional details on the technologies covered in this vCAT release.

**Figure 3. vCloud Technology Areas**



**Note** Except for the gray components, components that touch each other are integrated.

### 3.4.1 VMware Professional Services

VMware offers professional services that align with vCloud use cases. These range from a proof of concept (POC) that might be used as a demonstration environment, to a production deployment that requires management, workflow automation, compliance enforcement, and validation. The following services are available.

- VMware vCloud POC Jumpstart Service – Provides knowledge transfer workshops and hands-on product installation, configuration, and use demonstrations for the vCloud solution.
- VMware vCloud Accelerator Service – Rapidly delivers a functioning VMware vCloud implementation suitable for deploying applications in a limited scale, pre-production environment. If all prerequisites are met, this service engagement can be completed in less than 30 business days.

- VMware vCloud Provisioning Accelerator Service – This is the vCloud Accelerator Service with the addition of VMware Service Manager™. Service Manager automates IT service management by providing a powerful workflow engine to drive IT service processes, enabling you to deliver and manage physical, virtual, and vCloud-based IT services in a standardized and scalable manner. Service Manager presents IT services to end users with a service catalog and automates the request process for provisioning vCloud resources.
- VMware vCloud Plan and Design Service – Provides a comprehensive architectural design for VMware vCloud that addresses the customer's unique business requirements and operational demands, helping to pave the way to vCloud computing. This service is designed for enterprises that have a well-established, vSphere-based virtualization strategy for production workloads and that are ready to take the next step toward building their production vCloud infrastructure.
- VMware Operational Readiness for Cloud Computing Service – In this four-to-six-week engagement, VMware consultants examine existing operational practices to evaluate performance across more than 150 attributes in five key areas. They uncover unknown or hidden barriers to success and highlight areas where additional focus on people or process can deliver increased productivity, streamline operations, and improve overall vCloud solution results.

Services can be combined or customized to meet your specific requirements.



## 4. Journey to a Mature vCloud Implementation

At every stage in the processes leading to a mature vCloud implementation, financial transparency, process maturity, organizational setup, and technology implementation are critical factors for success.

VMware defines three stages on the journey to a mature vCloud: Standardize, Service Broker, and Strategic Partner. They are depicted in **Error! Reference source not found.** and described in the following sections.

**Figure 4. vCloud Journey States**



### 4.1.1 Stage 1 – Standardize

Stage 1 often coincides with a more mature server virtualization environment, and the focus is on creating a working vCloud solution with an on-demand service catalog that is directly accessible by end-users. The service catalog provides rapid deployment services for non-business-critical, development, and test environments as well as for externally sourced applications. Implementing the service catalog promotes cloud acceptance by business users as well as outlines a long-term vCloud implementation strategy with planning for operational and organizational change. The following capabilities are important for this stage:

- Financial Model and Measurement – Awareness and understanding of the costs of assets and underlying infrastructure capacity.
- People and Organization – Specialized but shared roles for managing virtualized environments.
  - No explicit virtualization Center of Excellence established.
  - See “Organizing for vCloud Operations” in *Operating a VMware vCloud* for information about the Center of Excellence.
- Process and Control.
  - IT processes are adapted for virtualization but are largely manual, with specific, customized inter-process integration.
  - Focus on limited, continuous improvement.
- Tools and Technology.
  - On-line, self-service capability for development and test provisioning.
  - On-line, self-service capability for Software as a Service (SaaS)-based applications.
  - Operational tools defined for virtualization environments.
  - Some business workloads run in a virtualized environment, whether internal or provided by third parties.

### 4.1.2 Stage 2 – Service Broker

Stage 2 is the first service-driven stage for a vCloud. At this stage, IT has transformed from traditional models and is focused on delivering business services within a vCloud environment. This represents a cultural shift within the organization. To be successful, it requires enhanced IT operational maturity, an optimized IT organizational structure, and supporting cloud management tools.

The term *Service Broker* implies that IT is organized at this stage to source internally and externally, such as adding external infrastructure capacity or providing access to vendor-based SaaS applications. The business is not necessarily aware of how the services are made available, but dramatically decreased development and provisioning times support business needs with increased quality of service and greater agility.

This stage focuses on:

- Alignment and buy-in from key business stakeholders.
- Creation of service governance, lifecycle and service design, and development processes.
- Providing service-based financial transparency.
- Automating and integrating tools and technology in internal and external systems.

Key capabilities for this stage include:

- Financial model and measurement:
  - Usage metering and cost showback.
  - Granular costing of underlying infrastructure assets.
  - Educating IT customers about paying for services as an operating expense.
  - Changing from project-based budgeting to demand-based budgeting.
- People and organization – Center of Excellence is established with dedicated, experienced, and knowledgeable staff.
- Process and control:
  - Fully integrated IT operational processes are adapted for virtualization and vCloud.
  - Agile-based service design and development processes are established.
  - Service-level financial transparency.
- Tools and technology:
  - Services are defined and offered through an online consumer portal for self-service access to the service catalog.
  - vCloud-level disaster recovery.
  - Blueprint and policy-driven service development and provisioning.
  - Purpose-built management tools for proactive vCloud operations.

### 4.1.3 Stage 3 – Strategic Partner

This stage is the final stage for a mature cloud. At this stage a highly efficient, scalable cloud with hybrid capability is available for an organization. IT is delivered as a service. There is automated, policy-driven governance and control across the vCloud environment with zero touch operations supported by predictive and self-healing operational tool capabilities. True application mobility and device-independent access is available. The vCloud is considered to be the de facto model within the organization. The term *strategic differentiator* implies that IT has changed roles and become a business differentiator by increasing agility, resulting in faster time to market, increasing efficiency, resulting in reduced costs, and increasing reliability, resulting in dramatically increased quality of service. The following are key capabilities for this stage:

- Financial model and measurement:
  - Usage-based pricing and chargeback for services provided to business customers.
  - Service demand-based budgeting.
  - Priced catalog of service offerings.
- People and organization – The Center of Excellence manages all elements of infrastructure, end-user, and application operations.
- Process and control:
  - Optimized, integrated, and fully automated IT processes enhance business agility and efficiency.
  - Continuous process, service, and performance improvements based on predictive capabilities.
- Tools and technology:
  - Full hybrid capabilities.
  - Tools that support single-pane-of-glass management across private and public vCloud environments.
  - Service-level disaster recovery.
  - Tools that support automated corrective actions for self-healing.