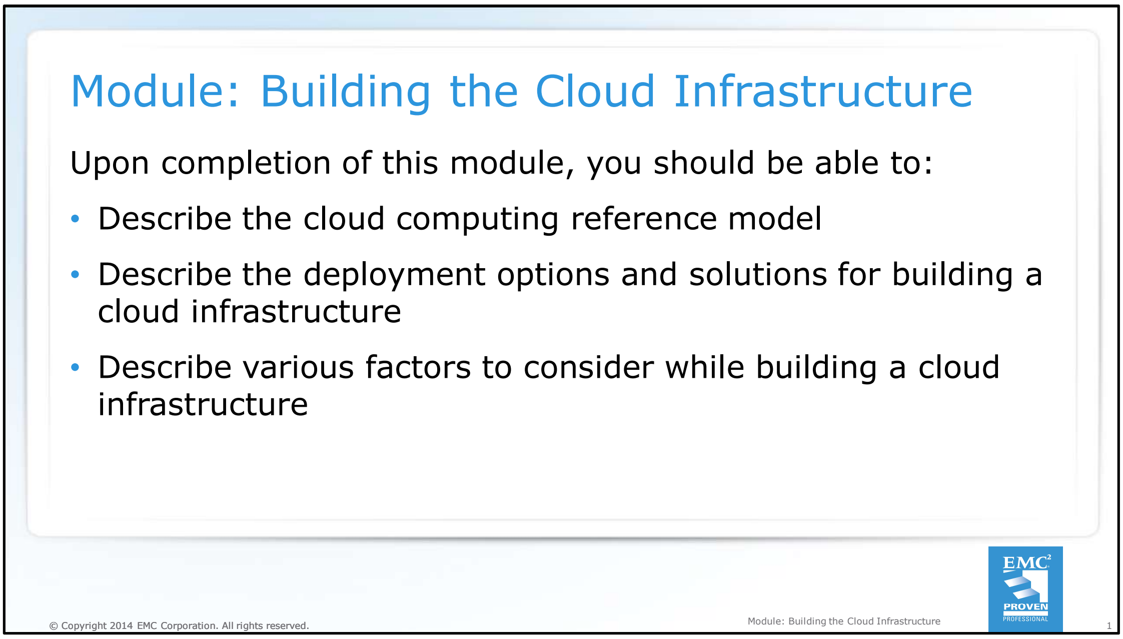

Module: Building the Cloud Infrastructure

Upon completion of this module, you should be able to:

- Describe the cloud computing reference model
- Describe the deployment options and solutions for building a cloud infrastructure
- Describe various factors to consider while building a cloud infrastructure

This module focuses on the cloud computing reference model, deployment options, and solutions for building a cloud infrastructure. The module also focuses on various factors that should be considered by a cloud service provider while deploying a cloud infrastructure.

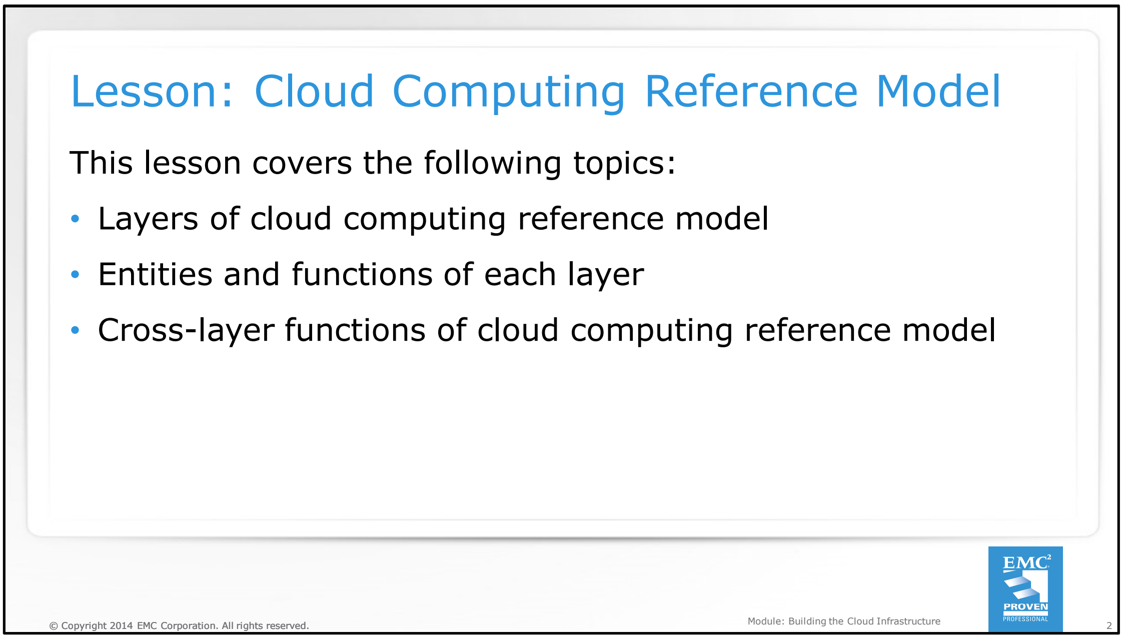

Lesson: Cloud Computing Reference Model

This lesson covers the following topics:

- Layers of cloud computing reference model
- Entities and functions of each layer
- Cross-layer functions of cloud computing reference model

This lesson covers the cloud computing reference model. It covers the entities and functions of the five layers of the model. It also covers the three cross-layer functions of the cloud computing reference model.

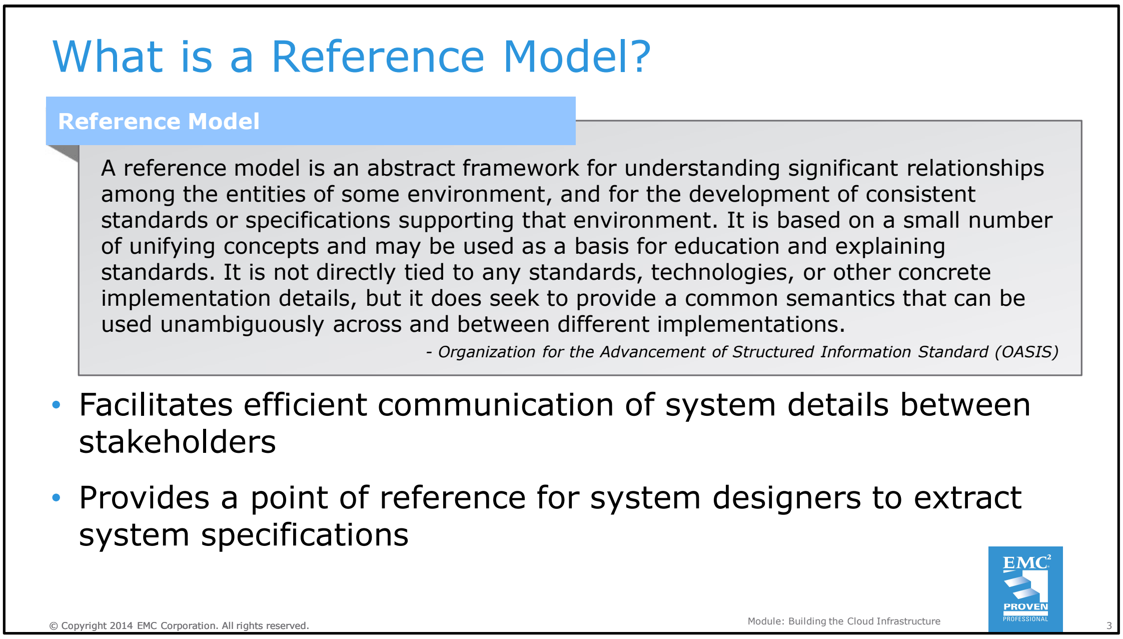

What is a Reference Model?

Reference Model

A reference model is an abstract framework for understanding significant relationships among the entities of some environment, and for the development of consistent standards or specifications supporting that environment. It is based on a small number of unifying concepts and may be used as a basis for education and explaining standards. It is not directly tied to any standards, technologies, or other concrete implementation details, but it does seek to provide a common semantics that can be used unambiguously across and between different implementations.

- Organization for the Advancement of Structured Information Standard (OASIS)

- Facilitates efficient communication of system details between stakeholders
- Provides a point of reference for system designers to extract system specifications

© Copyright 2014 EMC Corporation. All rights reserved.

Module: Building the Cloud Infrastructure

3

According to Organization for the Advancement of Structured Information Standards (OASIS), a reference model is an abstract framework for understanding the significant relationships among the entities of some environment, and for the development of consistent standards or specifications supporting that environment. A reference model is based on a small number of unifying concepts and may be used as a basis for education and explaining standards. A reference model is not directly tied to any standards, technologies, or other concrete implementation details, but it does seek to provide a common semantics that can be used unambiguously across and between different implementations.

Key goals of reference model are:

- Conveys fundamental principles and basic functionality of a system it represents
- Facilitates efficient communication of system details between stakeholders
- Provides a point of reference for system designers to extract system specifications
- Enhances an individual’s understanding of the representative system
- Documents the system for future reference and provides a means for collaboration

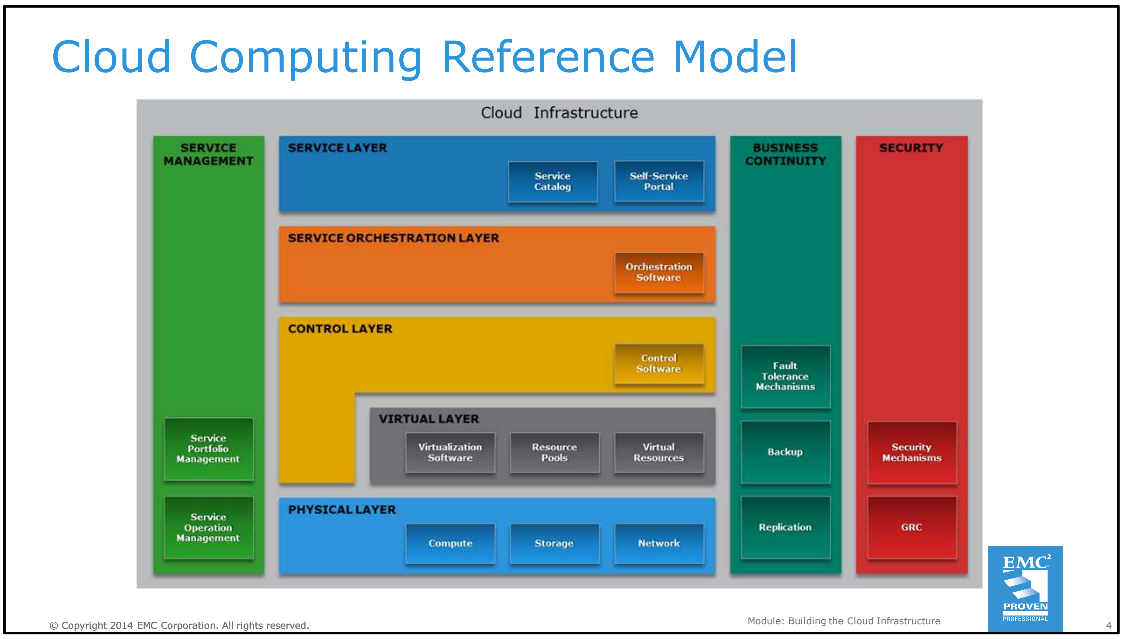

Cloud Computing Reference Model

Cloud Infrastructure

SERVICE MANAGEMENT

Service Portfolio Management

Service Operation Management

SERVICE LAYER

Service Catalog

Self-Service Portal

SERVICE ORCHESTRATION LAYER

Orchestration Software

CONTROL LAYER

Control Software

VIRTUAL LAYER

Virtualization Software

Resource Pools

Virtual Resources

PHYSICAL LAYER

Compute

Storage

Network

BUSINESS CONTINUITY

Fault Tolerance Mechanisms

Backup

Replication

SECURITY

Security Mechanisms

GRC

The cloud computing reference model is an abstract model that characterizes and standardizes the functions of a cloud computing environment by partitioning it into abstraction layers and cross-layer functions. This reference model groups the cloud computing functions and activities into five logical layers and three cross-layer functions.

The five layers are physical layer, virtual layer, control layer, service orchestration layer, and service layer. Each of these layers specifies various types of entities that may exist in a cloud computing environment, such as compute systems, network devices, storage devices, virtualization software, security mechanisms, control software, orchestration software, management software, and so on. It also describes the relationships among these entities.

The three cross-layer functions are business continuity, security, and service management. Business continuity and security functions specify various activities, tasks, and processes that are required to offer reliable and secure cloud services to the consumers. Service management function specifies various activities, tasks, and processes that enable the administrations of the cloud infrastructure and services to meet the provider's business requirements and consumer's expectations.

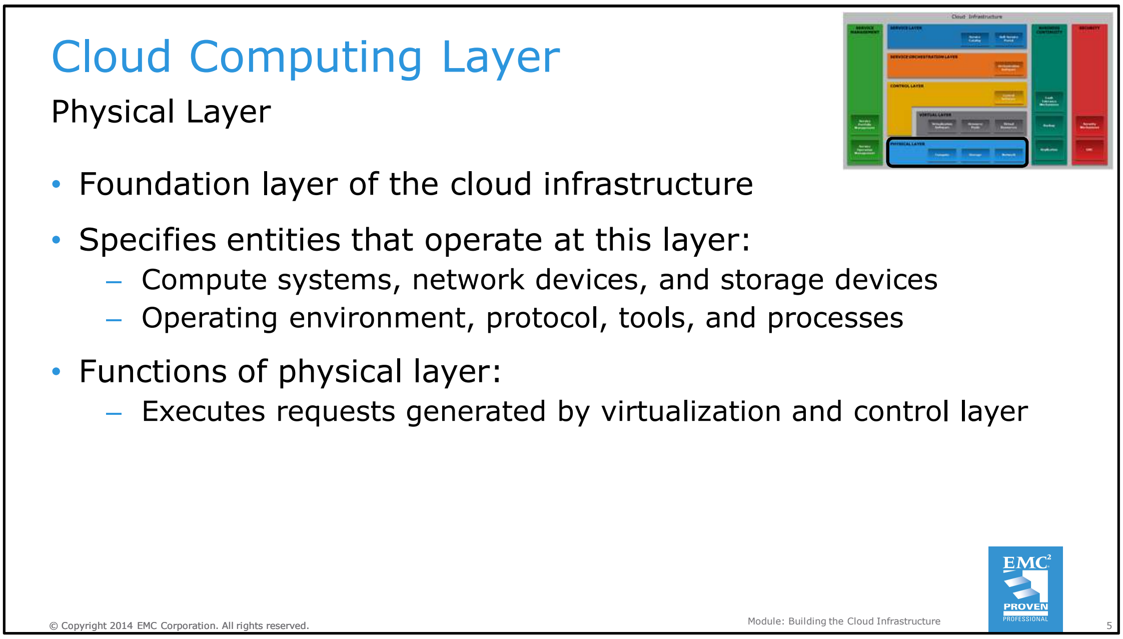

Cloud Computing Layer

Physical Layer

- Foundation layer of the cloud infrastructure
- Specifies entities that operate at this layer:
  - Compute systems, network devices, and storage devices
  - Operating environment, protocol, tools, and processes
- Functions of physical layer:
  - Executes requests generated by virtualization and control layer

Physical layer is the foundation layer of the cloud infrastructure. Physical layer specifies the physical entities that operate at this layer such as compute systems, networking devices, and storage devices. This layer also specifies the entities such as operating environment, protocols, tools, and processes that enable the physical entities of this layer to perform their functions and serve other layers of the cloud infrastructure. A key function of this layer is to execute the request generated from the virtualization layer or control layer. Examples of requests from the layers include storing data on the storage devices, performing communication among compute systems, executing programs on a compute systems, creating backup copy of data, or executing security policy to block an unauthorized activity.

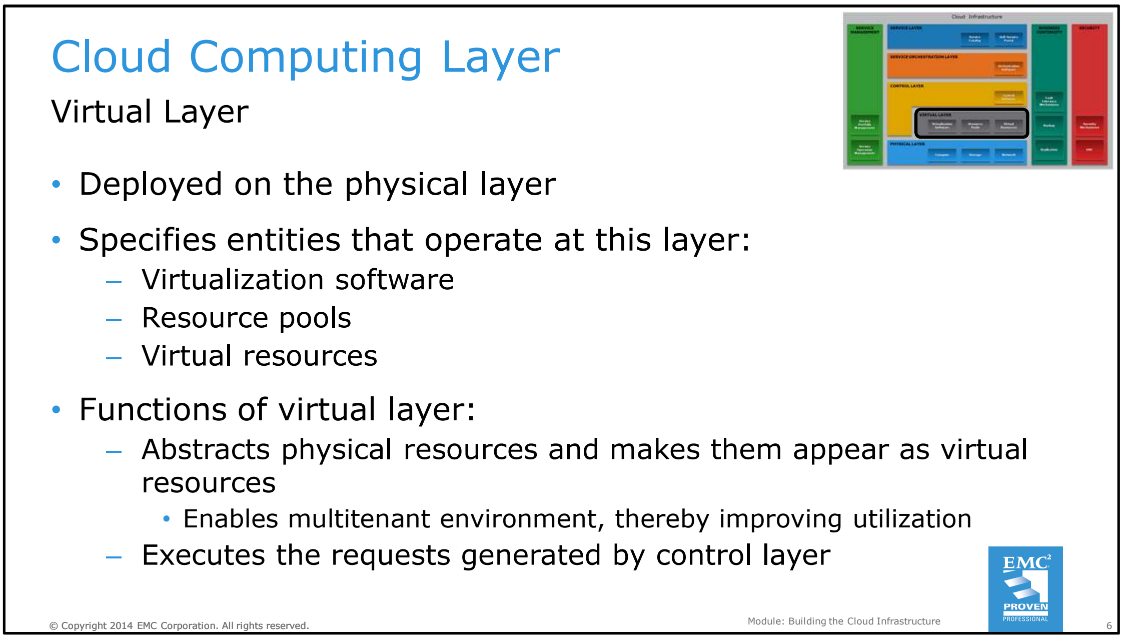

Cloud Infrastructure
SERVICE MANAGEMENT
SERVICE LAYER
Service Catalog
Self-Service Portal
SERVICE ORCHESTRATION LAYER
Orchestration Software
CONTROL LAYER
Control Software
VIRTUAL LAYER
Virtualization Software
Resource Pools
Virtual Resources
PHYSICAL LAYER
Compute
Storage
Network
BUSINESS CONTINUITY
Fault Tolerance Mechanisms
Backup
Replication
SECURITY
Security Mechanisms
GRC
Service Portfolio Management
Service Operation Management
EMC²
PROVEN
PROFESSIONAL

Cloud Computing Layer

Virtual Layer

- Deployed on the physical layer
- Specifies entities that operate at this layer:
  - Virtualization software
  - Resource pools
  - Virtual resources
- Functions of virtual layer:
  - Abstracts physical resources and makes them appear as virtual resources
    - Enables multitenant environment, thereby improving utilization
  - Executes the requests generated by control layer

Module: Building the Cloud Infrastructure 6

Virtual layer is deployed on the physical layer. It specifies the entities that operate at this layer such as virtualization software, resource pools, and virtual resources. A key function of this layer is to abstract physical resources, such as compute, storage, and network, and make them appear as virtual resources. Virtualization software deployed on compute systems, network devices, and storage devices perform the abstraction of the physical resources on which they are deployed. Abstracting the physical resources enables multitenant environment, thereby improving the utilization of the physical resources. Improved utilization of physical resources results in increased return-on-investment (ROI) on the infrastructure entities.

Virtualization software is also responsible for pooling physical resources from which virtual resources are created. Examples of virtual resources include virtual machines, LUN, and virtual network. The request to create resource pools and virtual resources is generated by the control layer. After receiving the request from the control layer, the virtual layer executes the requests. Apart from creating the resource pools and the virtual resources, virtualization software also support features that enable optimized resource utilization that further increases return-on-investment.

Other key functions of this layer include executing the requests generated by the control layer, and it also includes forwarding requests to the physical layer to get them executed. Examples of requests generated by the control layers include creating pools of resources and creating virtual resources.

Note: While deploying a cloud infrastructure, organization may choose not to deploy virtual layer. In such an environment, the control layer is deployed over the physical layer and it can directly request the physical layer to perform an operation. Further, it is also possible that part of the infrastructure is virtualized and rest is not virtualized.

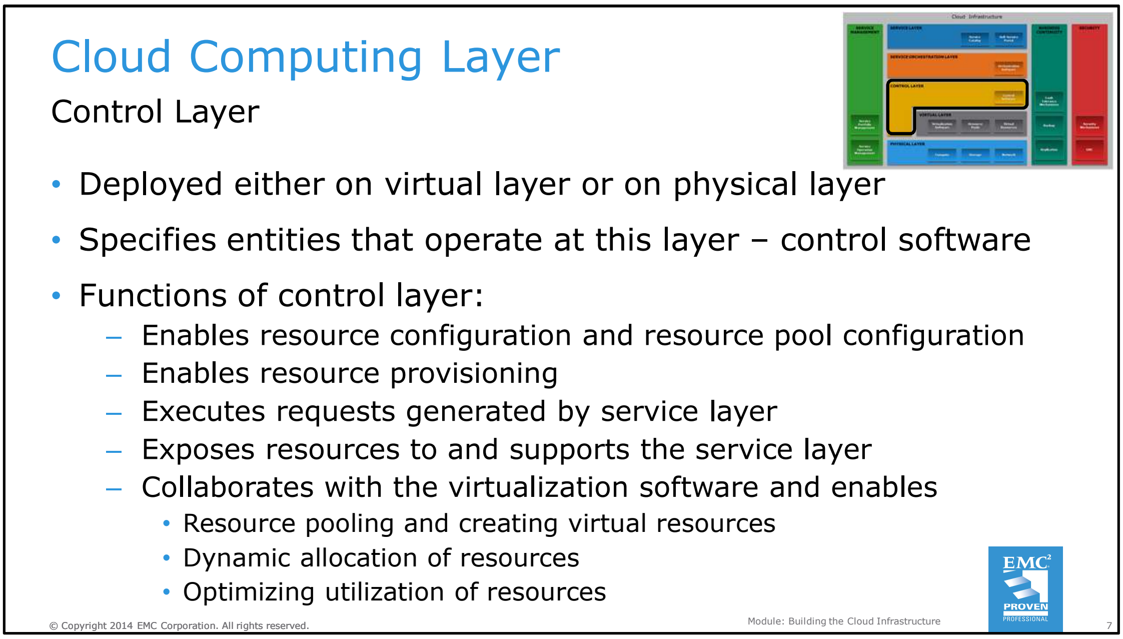

Cloud Computing Layer

Control Layer

- Deployed either on virtual layer or on physical layer
- Specifies entities that operate at this layer – control software
- Functions of control layer:
  - Enables resource configuration and resource pool configuration
  - Enables resource provisioning
  - Executes requests generated by service layer
  - Exposes resources to and supports the service layer
  - Collaborates with the virtualization software and enables
    - Resource pooling and creating virtual resources
    - Dynamic allocation of resources
    - Optimizing utilization of resources

© Copyright 2014 EMC Corporation. All rights reserved.

Module: Building the Cloud Infrastructure

7

Control layer can be deployed either on the virtual layer or on the physical layer. It specifies the entities that operate at this layer such as control software. A key function of this layer includes executing the requests generated by the service layer in collaboration with the orchestration layer. Another key function of this layer includes forwarding requests to the virtual and/or physical layer to get them executed. Examples of requests generated by the service layer include creating service instance such as compute system instance for IaaS and application instance for SaaS.

The other key functions that are performed by control software are resource configuration, resource pool configuration, and resource provisioning. The control software in collaboration with the virtualization software enables resource pooling, dynamic allocation of resources, creating virtual resources, and optimizing utilization of resources. The control software initiates all the requests such as resource configuration, resource pooling, resource provisioning, and so on. These requests are passed on to the virtual layer or physical layer. In the absence of virtual layer, the requests generated by the control layer are passed on to the physical layer. In this case, these requests are fulfilled by the operating environment in collaboration with the control software.

This layer also exposes resources (physical and/or virtual) to and supports the service layer where cloud service interfaces are exposed to consumers.

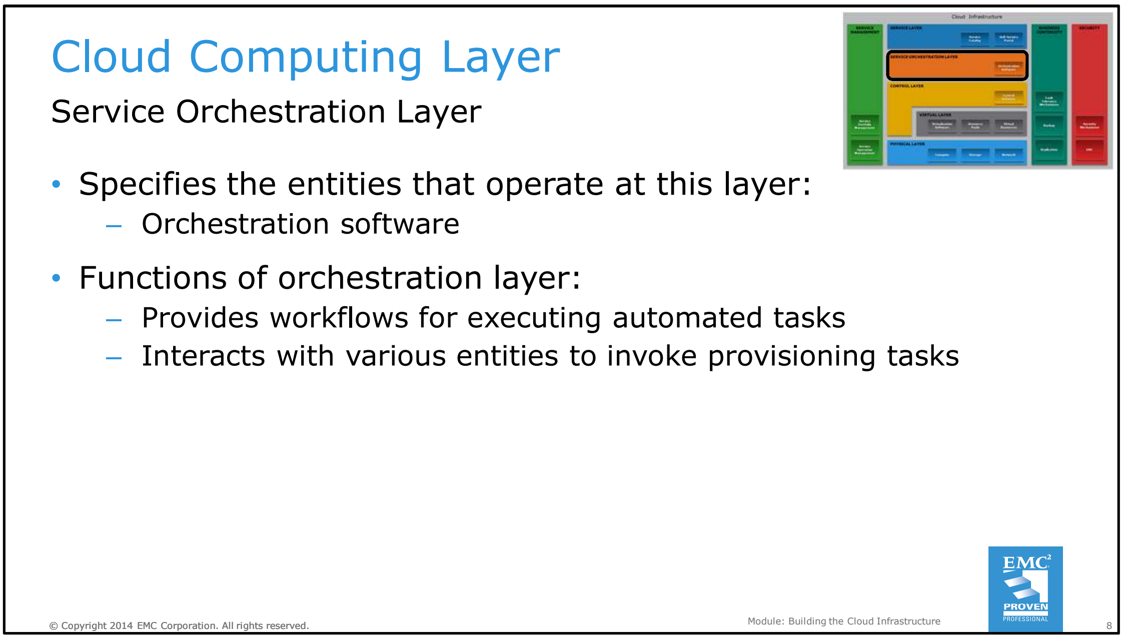

Cloud Computing Layer

Service Orchestration Layer

- Specifies the entities that operate at this layer:
  - Orchestration software
- Functions of orchestration layer:
  - Provides workflows for executing automated tasks
  - Interacts with various entities to invoke provisioning tasks

Service orchestration layer specifies the entities that can operate at this layer such as a orchestration software. A key function of this layer is to provide workflows for executing automated tasks to accomplish a desired outcome. Workflow refers to a series of inter-related tasks that perform a business operation. The orchestration software enables this automated arrangement, coordination, and management of the tasks. This helps to group and sequence tasks with dependencies among them into a single, automated workflow.

Associated with each service listed in the service catalog, there is an orchestration workflow defined. When a consumer selects a service from the service catalog, an associated workflow in the orchestration layer is triggered. Based on this workflow, the orchestration software interacts with various entities (from control layer, business continuity function, security function, and service management function) to invoke the provisioning tasks to be executed by the entities.

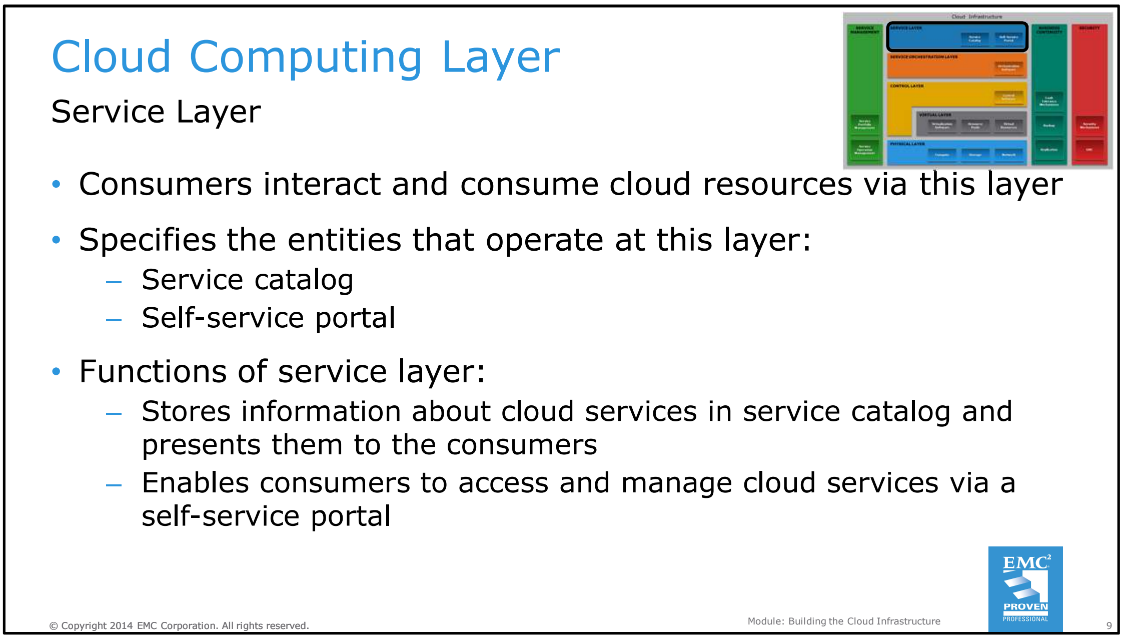

Cloud Computing Layer

Service Layer

- Consumers interact and consume cloud resources via this layer
- Specifies the entities that operate at this layer:
  - Service catalog
  - Self-service portal
- Functions of service layer:
  - Stores information about cloud services in service catalog and presents them to the consumers
  - Enables consumers to access and manage cloud services via a self-service portal

The service layer is accessible to the cloud consumers. This layer specifies the entities that can operate at this layer such as service catalog and self-service portal. A key function of this layer is to store and present the information about all the services offered to the cloud consumers in a service catalog. A service catalog is a database of information about the cloud services offered by a service provider. The service catalog includes a variety of information about the services, including description of the services, the types of services, cost, supported SLAs, security mechanisms, and so on.

Another key function of this layer is to enable cloud consumers to access and manage the cloud services via a self-service portal. A self-service portal displays the service catalog to the consumers. Consumers can use this web portal to request for cloud services. In addition to the service catalog, it also provides interface to access and manage the rented service instances. The provisioning and management requests are passed on to the orchestration layer, where the orchestration workflows—to fulfill the requests—are defined.

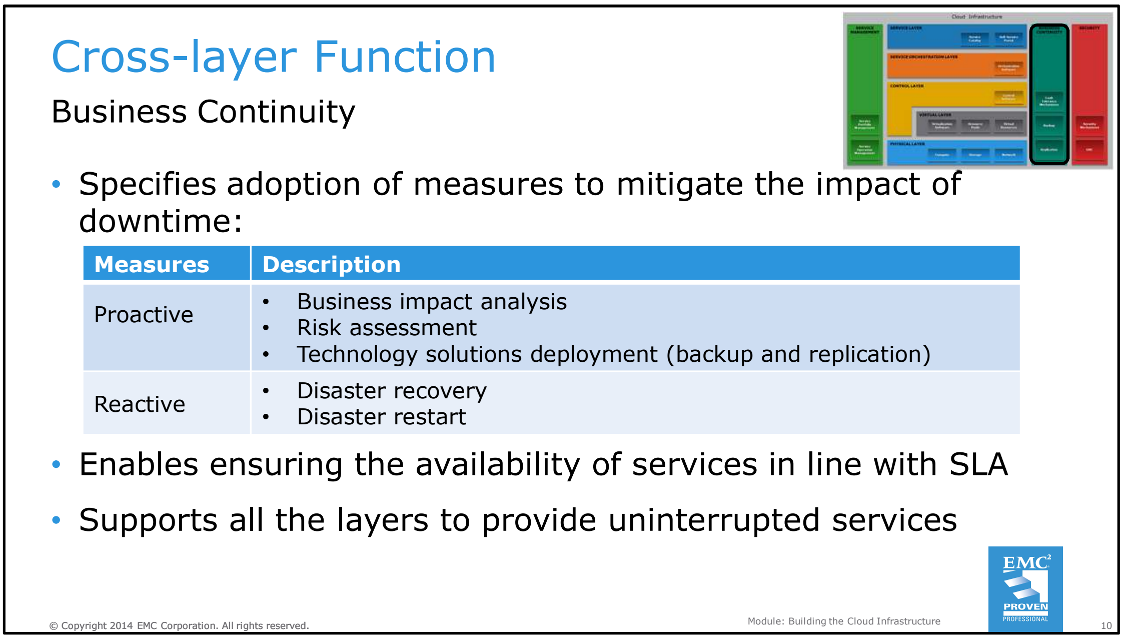

Cross-layer Function

Business Continuity

- Specifies adoption of measures to mitigate the impact of downtime:

Measures	Description
Proactive	• Business impact analysis • Risk assessment • Technology solutions deployment (backup and replication)
Reactive	• Disaster recovery • Disaster restart

- Enables ensuring the availability of services in line with SLA
- Supports all the layers to provide uninterrupted services

Business continuity (BC) cross-layer function specifies the adoption of proactive and reactive measures that enable a business to mitigate the impact of planned and unplanned downtime. Proactive measures include activities, tasks, processes such as business impact analysis, risk assessment, and technology solutions deployment (such as backup and replication). Reactive measures include activities, tasks, processes such as disaster recovery and disaster restart to be invoked in the event of a service failure. This function supports all the layers—physical, virtual, control, orchestration, and service—to provide uninterrupted services to the consumers. The BC cross-layer function of a cloud infrastructure enables a business to ensure the availability of services in line with the Service Level Agreement (SLA).

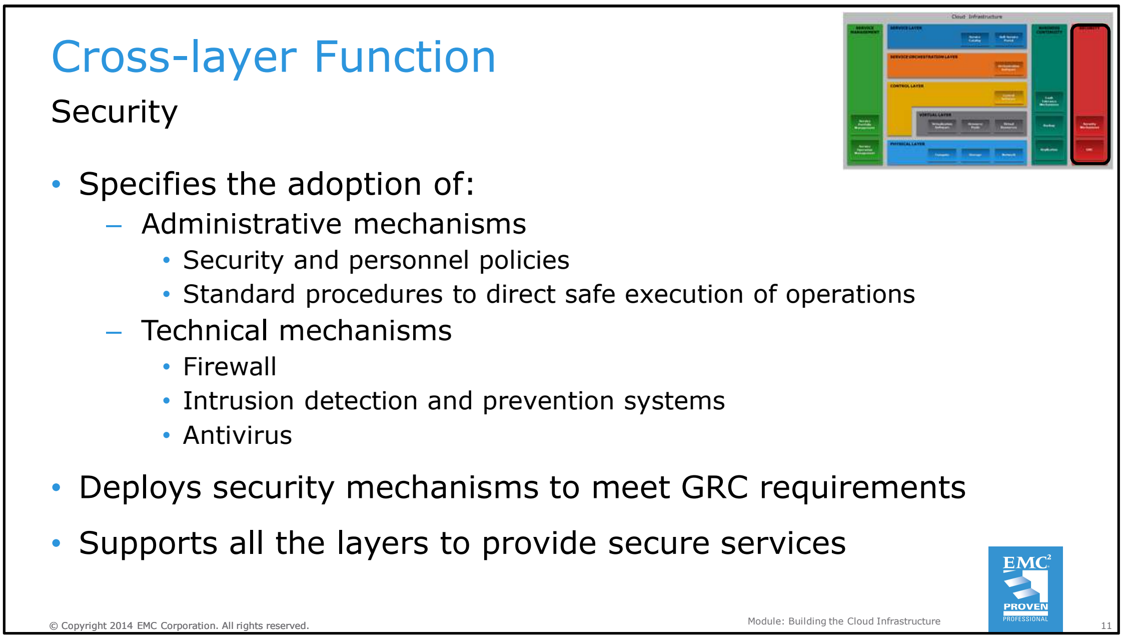

Cross-layer Function

Security

- Specifies the adoption of:
  - Administrative mechanisms
    - Security and personnel policies
    - Standard procedures to direct safe execution of operations
  - Technical mechanisms
    - Firewall
    - Intrusion detection and prevention systems
    - Antivirus
- Deploys security mechanisms to meet GRC requirements
- Supports all the layers to provide secure services

Security cross-layer function specifies the adoption of administrative and technical mechanism that can mitigate or minimize security threats and provide a secure cloud environment. Administrative mechanisms include security and personnel policies or standard procedures to direct the safe execution of various operations. Technical mechanisms are usually implemented through tools or devices deployed on the IT infrastructure. Examples of technical mechanisms include firewall, intrusion detection and prevention systems, antivirus, and so on.

Governance, risk, and compliance (GRC) specifies processes that help an organization ensure that their acts are ethically correct and in accordance with their risk appetite (the risk level an organization chooses to accept), internal policies, and external regulations. Security mechanisms should be deployed to meet the GRC requirements.

This cross-layer function supports all the layers—physical, virtual, control, orchestration, and service—to provide secure services to the consumers.

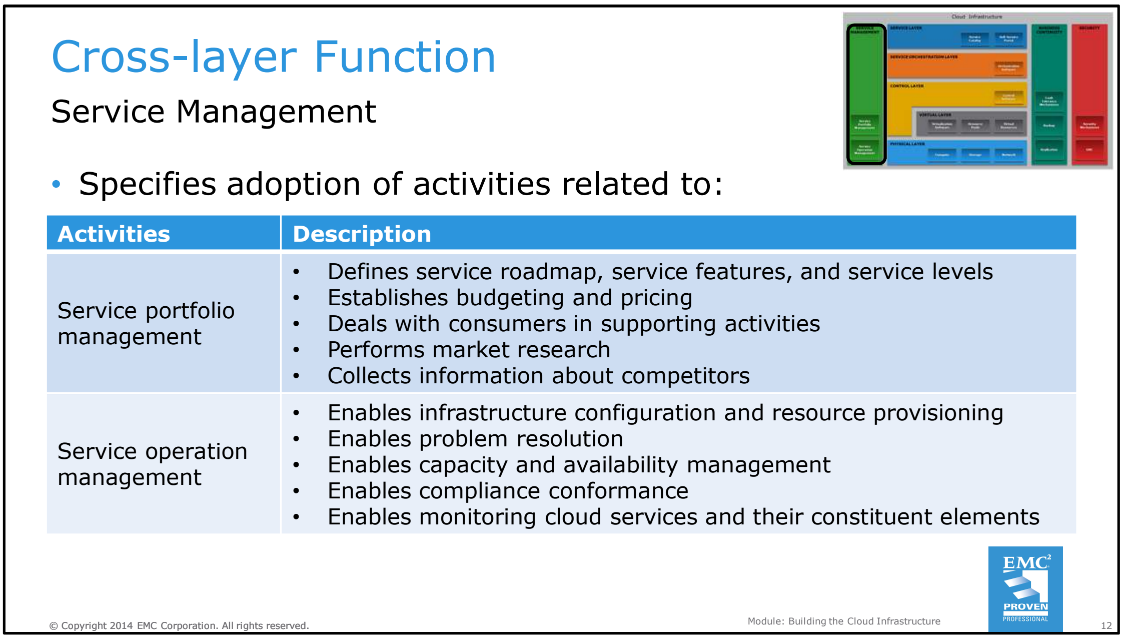

Cross-layer Function

Service Management

- Specifies adoption of activities related to:

Activities	Description
Service portfolio management	• Defines service roadmap, service features, and service levels • Establishes budgeting and pricing • Deals with consumers in supporting activities • Performs market research • Collects information about competitors
Service operation management	• Enables infrastructure configuration and resource provisioning • Enables problem resolution • Enables capacity and availability management • Enables compliance conformance • Enables monitoring cloud services and their constituent elements

Service management function specifies adoption of activities related to service portfolio management and service operation management. Adoption of these activities enables an organization to align the creation and delivery of cloud services to meet their business objectives and to meet the expectations of cloud service consumers.

Service portfolio management encompasses the set of business-related services that:

- Define the service roadmap, service features, and service levels
- Assess and prioritize where investments across the service portfolio are most needed
- Establish budgeting and pricing
- Deal with consumers in supporting activities such as taking orders, processing bills, and collecting payments

Service portfolio management also performs market research, measures service adoption, collects information about competitors, and analyzes feedback from consumers in order to quickly modify and align services according to consumer needs and market conditions.

Service operation management enables cloud administrators to manage cloud infrastructure and services. Service operation management tasks include handling of infrastructure configuration, resource provisioning, problem resolution, capacity, availability, and compliance conformance. All of these tasks enable ensuring that services and service levels are delivered as committed. Service operation management also includes monitoring cloud services and their constituent elements. This enables the provider to gather information related to resource consumption and bill generation. This function supports all the layers to perform monitoring, management, and reporting for the entities of the infrastructure.

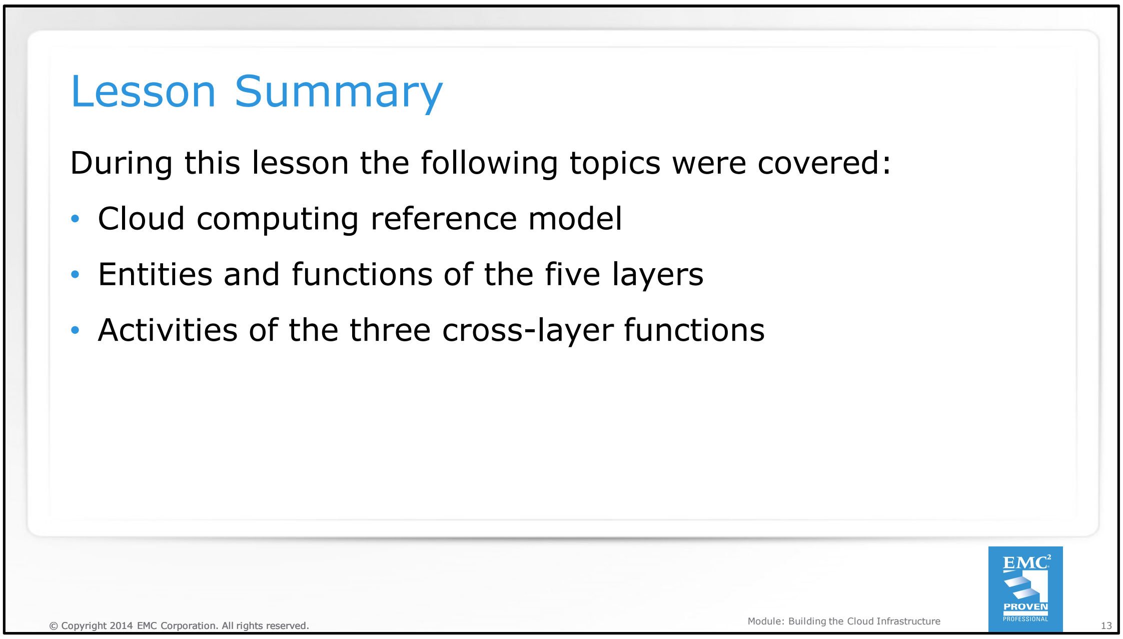

Lesson Summary

During this lesson the following topics were covered:

- Cloud computing reference model
- Entities and functions of the five layers
- Activities of the three cross-layer functions

This lesson covered the cloud computing reference model, and the entities and functions of the five layers. This lesson also covered the activities performed in the three cross-layer functions.

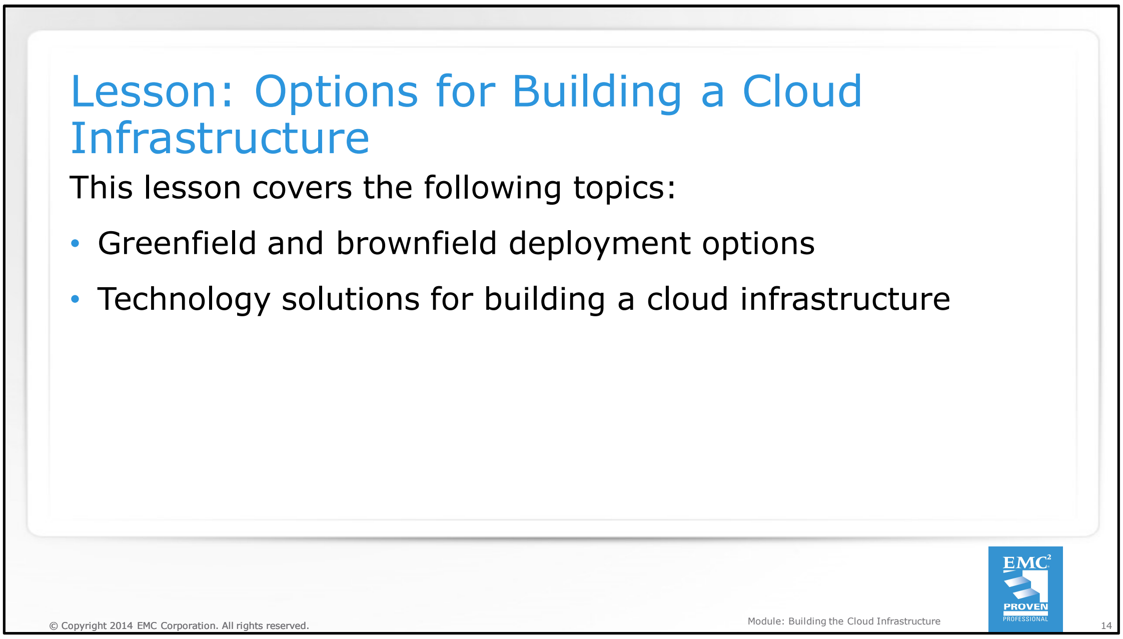

Lesson: Options for Building a Cloud Infrastructure

This lesson covers the following topics:

- Greenfield and brownfield deployment options
- Technology solutions for building a cloud infrastructure

This lesson covers the greenfield and the brownfield deployment options for building a cloud infrastructure. The lesson also covers the technology solutions for building a cloud infrastructure.

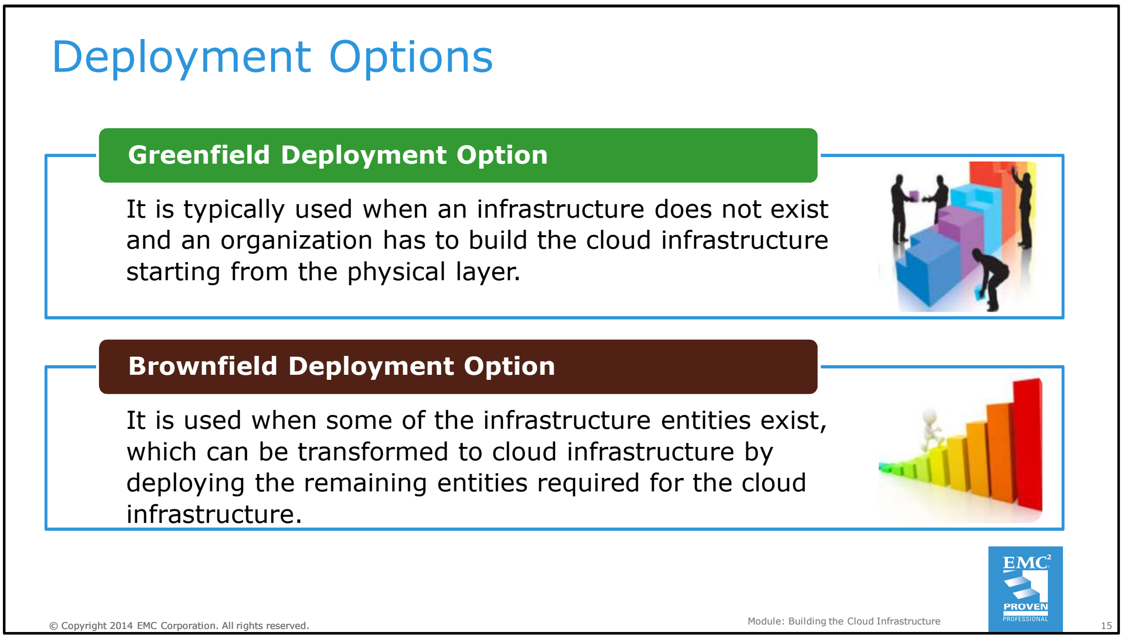

Deployment Options

Greenfield Deployment Option

It is typically used when an infrastructure does not exist and an organization has to build the cloud infrastructure starting from the physical layer.

Brownfield Deployment Option

It is used when some of the infrastructure entities exist, which can be transformed to cloud infrastructure by deploying the remaining entities required for the cloud infrastructure.

Module: Building the Cloud Infrastructure 15

Before building a cloud infrastructure, organizations must identify which deployment option is appropriate for them. There are two deployment options for building a cloud infrastructure and they are greenfield deployment option and brownfield deployment option. A greenfield deployment option is typically used when an infrastructure does not exist and an organization ha to build the cloud infrastructure starting from the physical layer. On the other hand, a brownfield deployment option is used when some of the infrastructure entities exist, which can be transformed to a cloud infrastructure by deploying the remaining entities required for the cloud infrastructure. For example, consider that an organization wants to use a brownfield deployment option to transform their existing data center, which has the physical, virtual, and control layers deployed. In such cases, the data center also has the business continuity, security, and service management in place. However, these three cross-layer functions are limited to a non-cloud environment. While transforming the existing data center to a cloud infrastructure, the organization will have to deploy the orchestration layer and the service layer. Further, the BC, security, and the service management functions will have to be transformed to support the cloud environment.

In both deployment options, apart from deploying the five layers and the three cross-layer functions, the organizations have to consider several factors that will enable them to deploy the cloud services that will meet the consumers' expectations. These factors are covered later in this module.

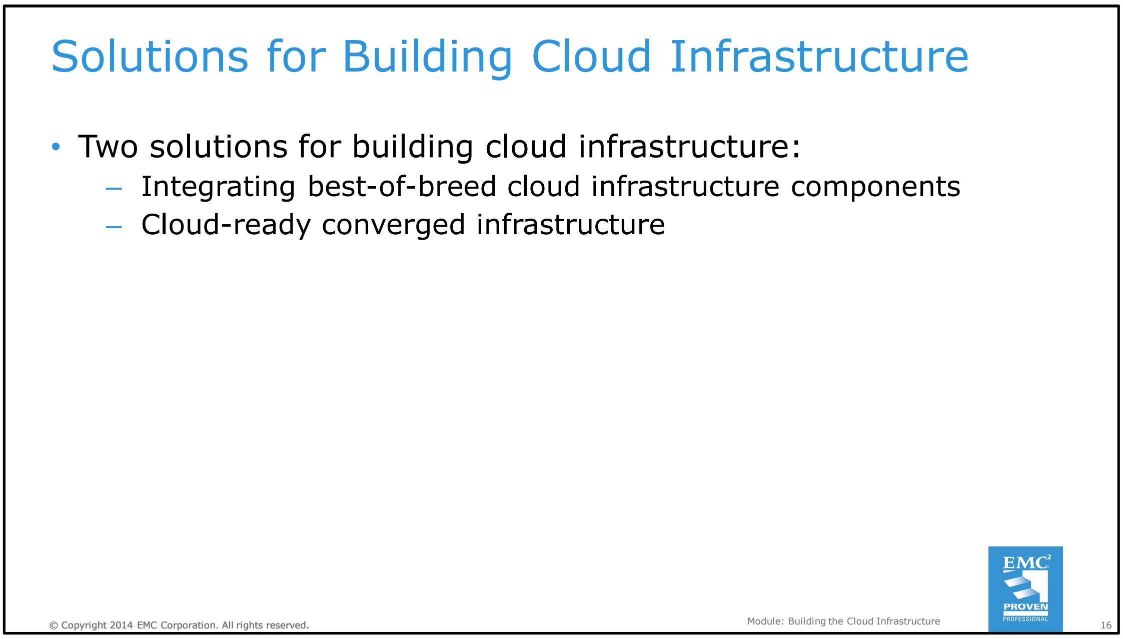

Solutions for Building Cloud Infrastructure

- Two solutions for building cloud infrastructure:
  - Integrating best-of-breed cloud infrastructure components
  - Cloud-ready converged infrastructure

There are two solutions for building a cloud infrastructure: by integrating best-of-breed cloud infrastructure components and by acquiring and implementing a cloud-ready converged infrastructure.

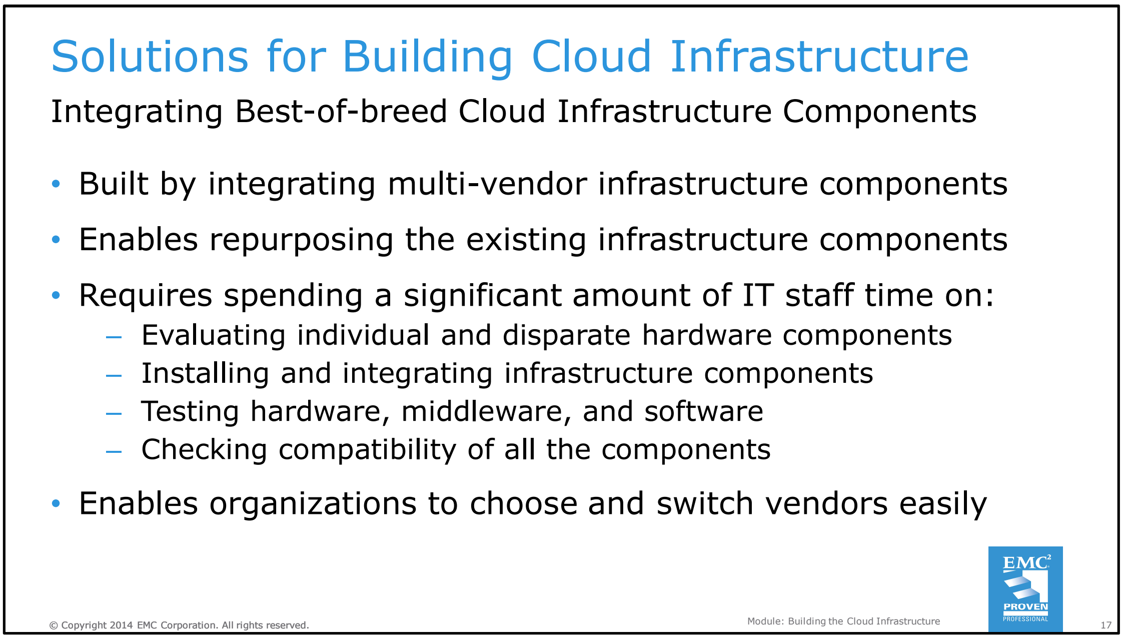

Solutions for Building Cloud Infrastructure

Integrating Best-of-breed Cloud Infrastructure Components

- Built by integrating multi-vendor infrastructure components
- Enables repurposing the existing infrastructure components
- Requires spending a significant amount of IT staff time on:
  - Evaluating individual and disparate hardware components
  - Installing and integrating infrastructure components
  - Testing hardware, middleware, and software
  - Checking compatibility of all the components
- Enables organizations to choose and switch vendors easily

Module: Building the Cloud Infrastructure 17

In an integrated best-of-breed cloud infrastructure components solution, organizations have the flexibility to use and integrate the infrastructure components from different vendors. This solution allows organizations to design their cloud infrastructure by repurposing their existing infrastructure components (in a brownfield deployment option), providing a cost advantage for this solution. This solution enables organizations to select a vendor of their choice for infrastructure components. This solution also enables an organization to easily switch a vendor if the vendor is unable to provide the committed support and not meet the SLAs.

When this method is used to build a cloud infrastructure, an organization may have to spend a significant amount of IT staff time evaluating individual, disparate hardware components, installing hardware, and integrating compute, storage, and network components. The IT staff may also have to spend effort integrating and testing hardware, middleware, and software. They also need to check the compatibility of all the components to ensure that the combined components interoperate and function as expected. This may delay the deployment of cloud services. Further, scaling of such an infrastructure takes longer because each component that is scaled requires integration with the existing infrastructure and testing for compatibility. Finally, this solution requires acquiring cloud infrastructure management tools and deploying them on the infrastructure.

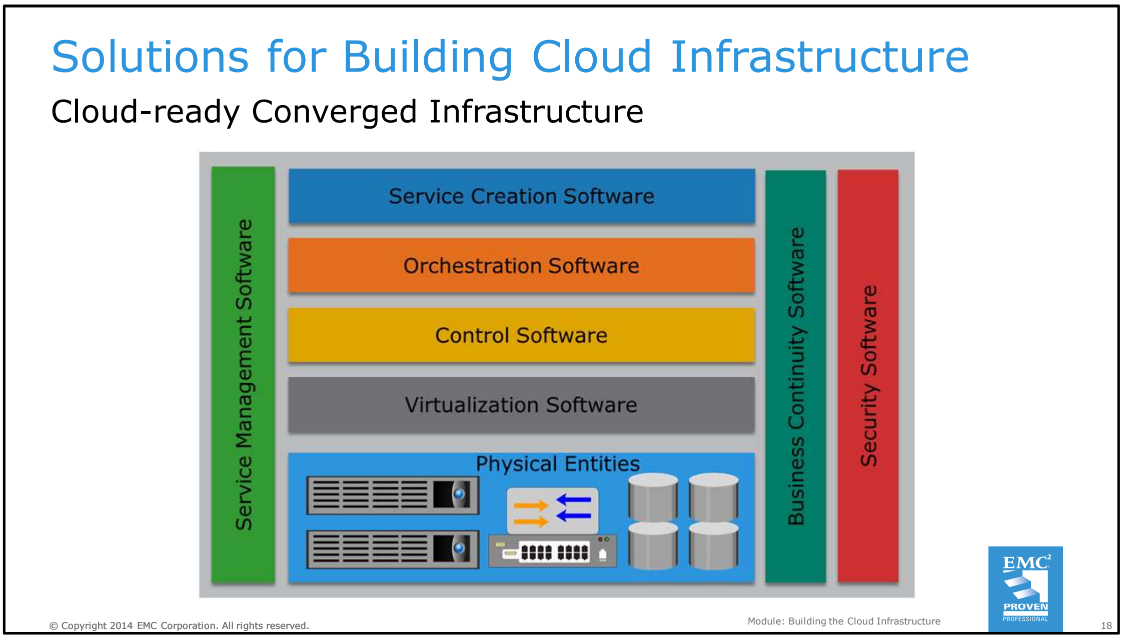

Solutions for Building Cloud Infrastructure

Cloud-ready Converged Infrastructure

Service Management Software

Service Creation Software

Orchestration Software

Control Software

Virtualization Software

Physical Entities

Business Continuity Software

Security Software

© Copyright 2014 EMC Corporation. All rights reserved.

Module: Building the Cloud Infrastructure

18

A *cloud ready converged infrastructure* solution provides a modular design that combines compute, storage, network, virtualization, and management components into a single package. This package is a self-contained unit that can be utilized to deploy cloud services, or can be aggregated with additional packages to support the demand for more capacity and performance. The package is pre-configured, reducing the time to deploy cloud services. Further, in addition to integrating various components into a package, this solution offers single management software capable of managing all hardware and software within the package.

A cloud-ready converged infrastructure solution has built-in capabilities that provide secured multi-tenancy. However, additional security mechanisms should be deployed to prevent external attacks. The solution is capable of managing and mitigating failure scenarios in hardware, software, and cloud services.

A potential area of concern regarding cloud-ready converged infrastructure solutions is the lack of flexibility to use infrastructure components from different vendors. Some vendors may provide organizations with the flexibility to choose multi-vendor infrastructure components such as network devices, compute systems, and hypervisors for this solution.

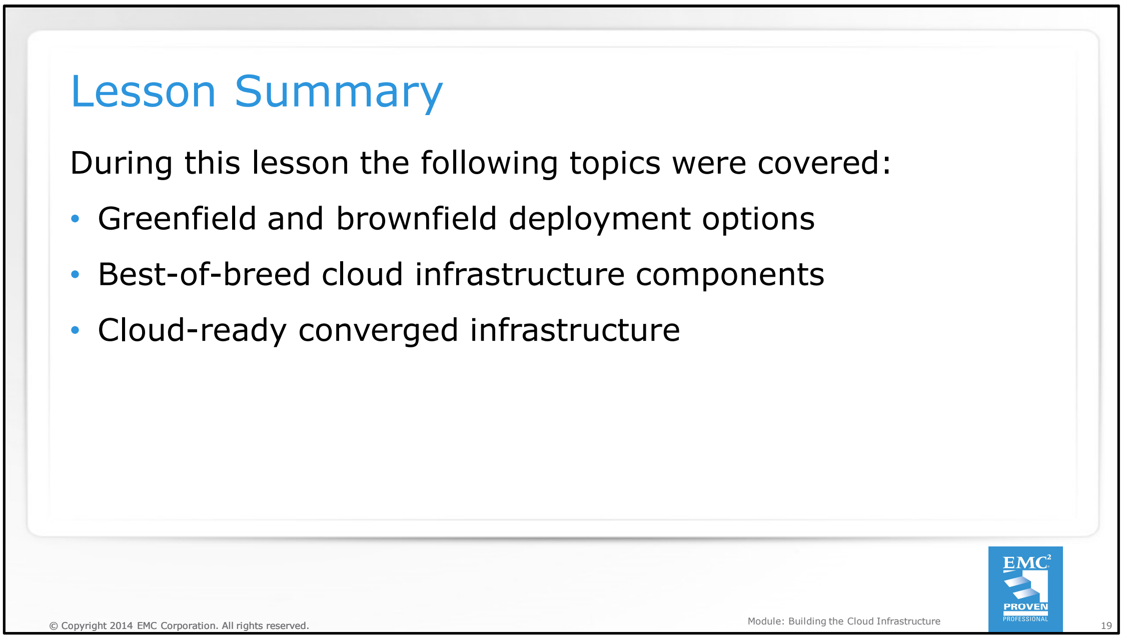

Lesson Summary

During this lesson the following topics were covered:

- Greenfield and brownfield deployment options
- Best-of-breed cloud infrastructure components
- Cloud-ready converged infrastructure

This lesson covered the greenfield and brownfield deployment options for building a cloud infrastructure. The lesson also covered the two technology solutions to build a cloud infrastructure: best-of-breed cloud infrastructure components and cloud-ready converged infrastructure.

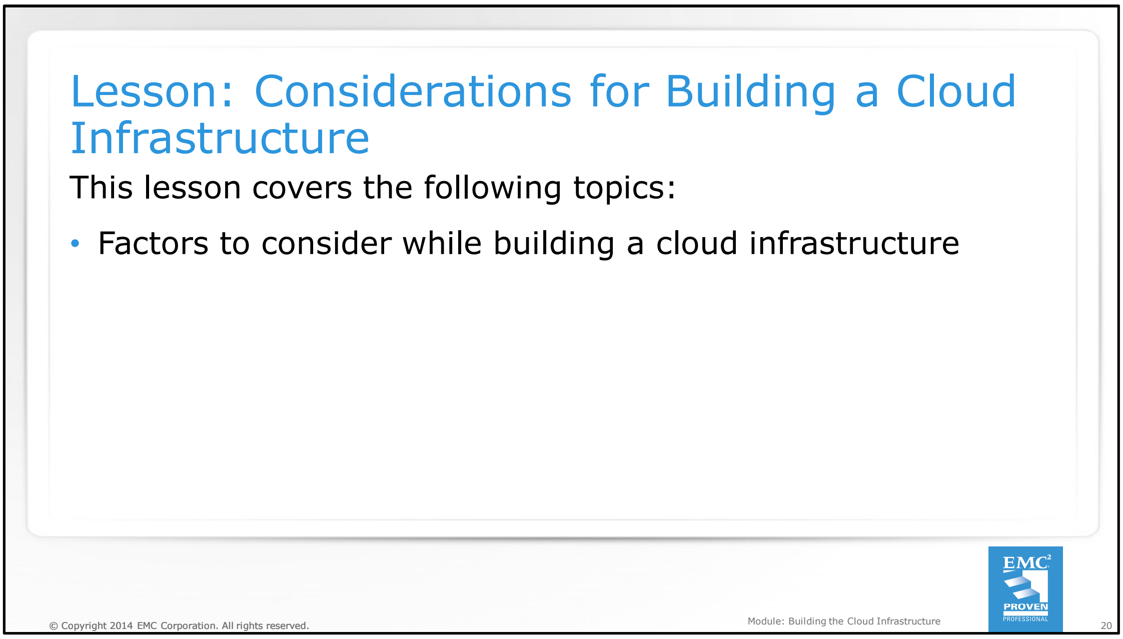

Lesson: Considerations for Building a Cloud Infrastructure

This lesson covers the following topics:

- Factors to consider while building a cloud infrastructure

This lesson covers various factors that should be considered while building a cloud infrastructure.

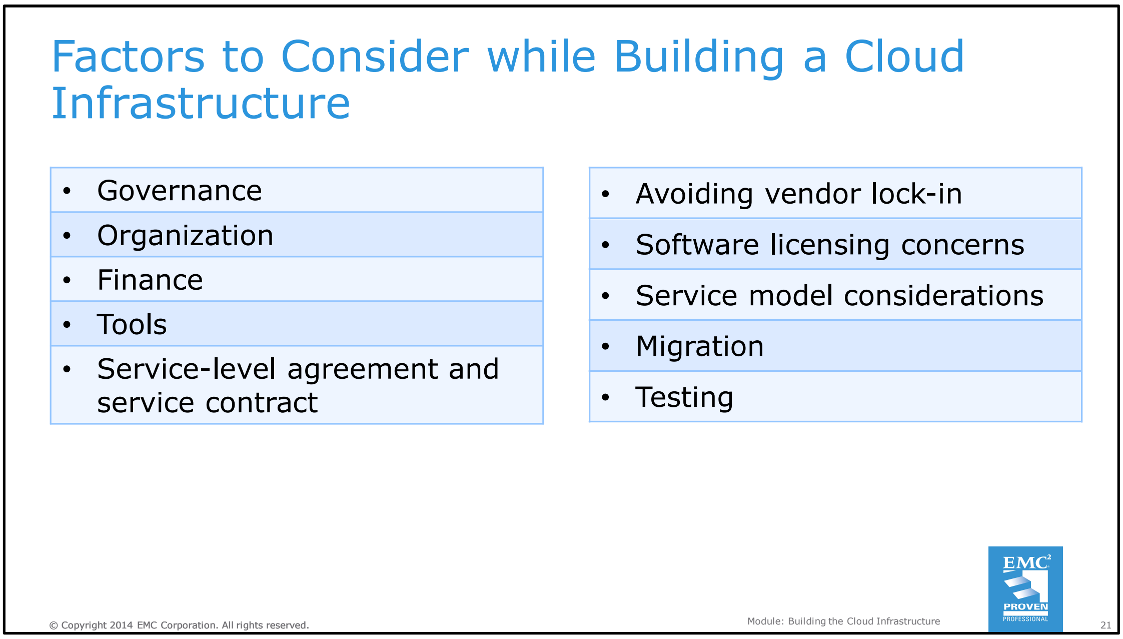

Factors to Consider while Building a Cloud Infrastructure

- Governance
- Organization
- Finance
- Tools
- Service-level agreement and service contract

- Avoiding vendor lock-in
- Software licensing concerns
- Service model considerations
- Migration
- Testing

After deciding on the deployment option and solution to build the cloud infrastructure, a cloud service provider have to consider several factors to deliver cloud services that meet their business objectives and consumer’s expectations. These slides list the key factors a service provider must consider while building a cloud infrastructure.

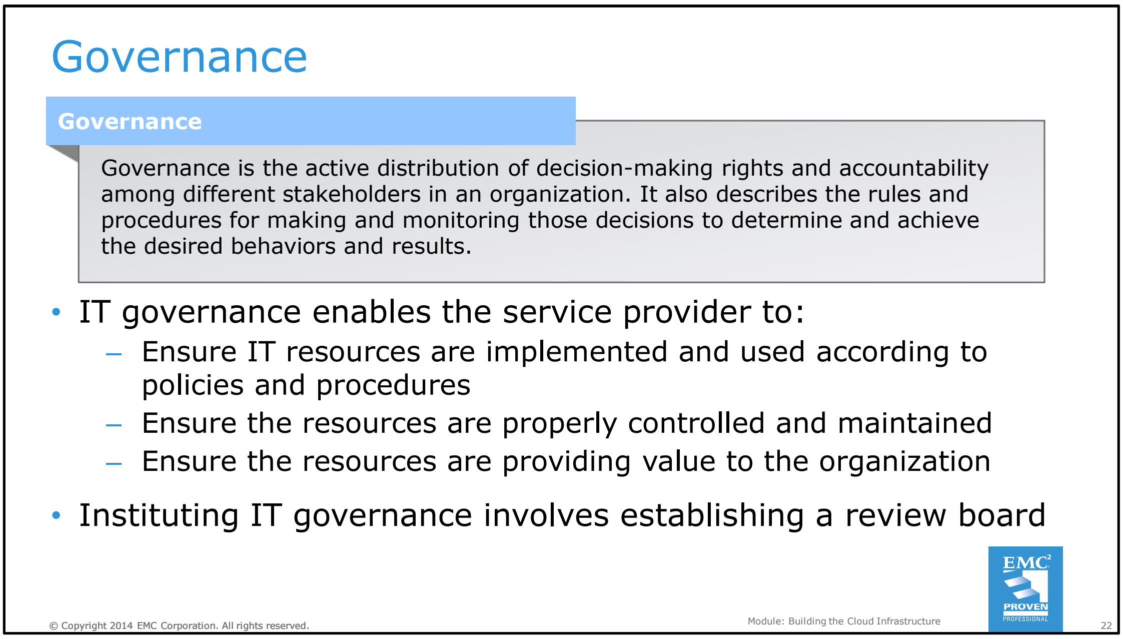

Governance

Governance

Governance is the active distribution of decision-making rights and accountability among different stakeholders in an organization. It also describes the rules and procedures for making and monitoring those decisions to determine and achieve the desired behaviors and results.

- IT governance enables the service provider to:
  - Ensure IT resources are implemented and used according to policies and procedures
  - Ensure the resources are properly controlled and maintained
  - Ensure the resources are providing value to the organization
- Instituting IT governance involves establishing a review board

Governance is the active distribution of decision-making rights and accountability among different stakeholders in an organization. It also describes the rules and procedures for making and monitoring those decisions to determine and achieve the desired behaviors and results. The role of governance in IT is to implement, maintain, and continuously improve the controls on the use of IT resources. IT governance enables a service provider to:

- Ensure that IT resources are implemented and used according to agreed-upon policies and procedures
- Ensure that these resources are properly controlled and maintained
- Ensure that these resources are providing value to the organization

Instituting IT governance usually involves establishing a review board, which is a team of members from across business units including IT. This review board is responsible for creating rules and processes that the organization must follow to ensure that policies are being met.

These rules and processes might include the following:

- Understanding business issues, such as regulatory requirements or funding
- Establishing best practices and monitoring these processes
- Assigning responsibility for things such as standards, design, review, and certifications

If a service provider is using a greenfield deployment option for building a cloud infrastructure, then they must establish governance by choosing appropriate governance model (discussed in the next slide). If a service provider is using a brownfield deployment option, then that service provider must transform their existing governance model to meet the cloud requirement.

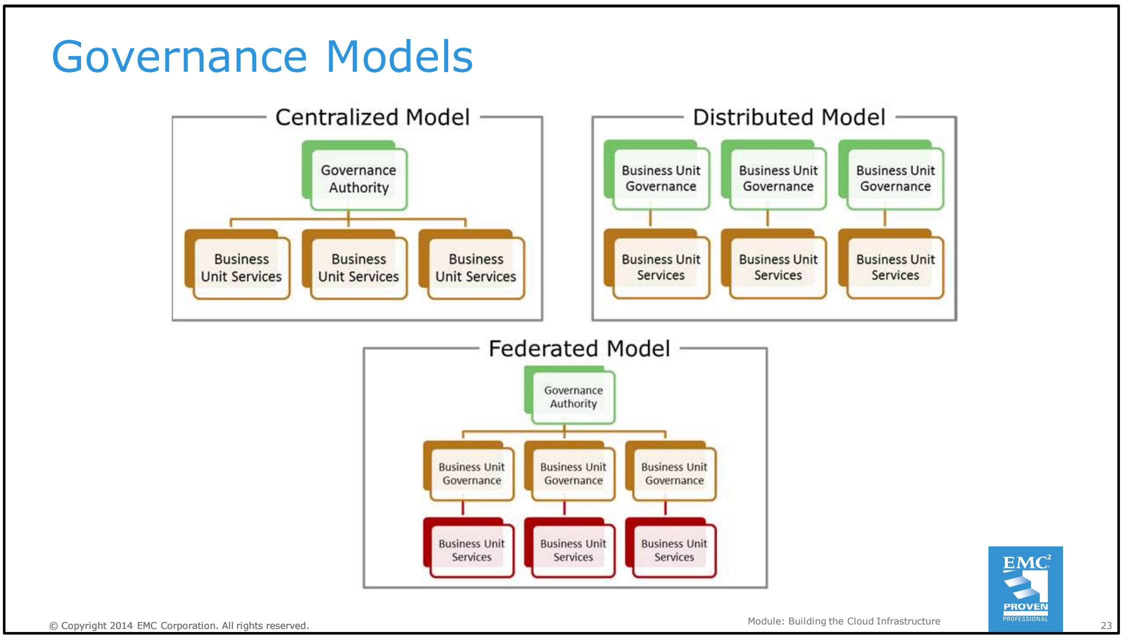

Governance Models

Depending on the size, structure, geographic presence, and culture of an organization, one of these fundamental governance models can be implemented:

- A *centralized model* provides one governance body for the entire organization. This fits best with a smaller or a strongly centralized organization where governance policies are, for the most part, consistent throughout the organization.

- A *federated model* proposes separate governance bodies, one for each business unit. A business unit can be a functional organization, a product group, or a geographic location. Each business unit has its own set of governance policies. Even though the services for a given business unit can be independently standardized, managed and owned, a single, enterprise-wide governance body can still subject all services to a common governance system.

- A *distributed model* proposes separate governance bodies for each business unit. These governance bodies function autonomously and are not controlled by any common governance system.

The organization can choose a governance model that best meets its requirements. After a governance model is chosen, the organization then needs to take steps to establish or transform to the chosen governance model.

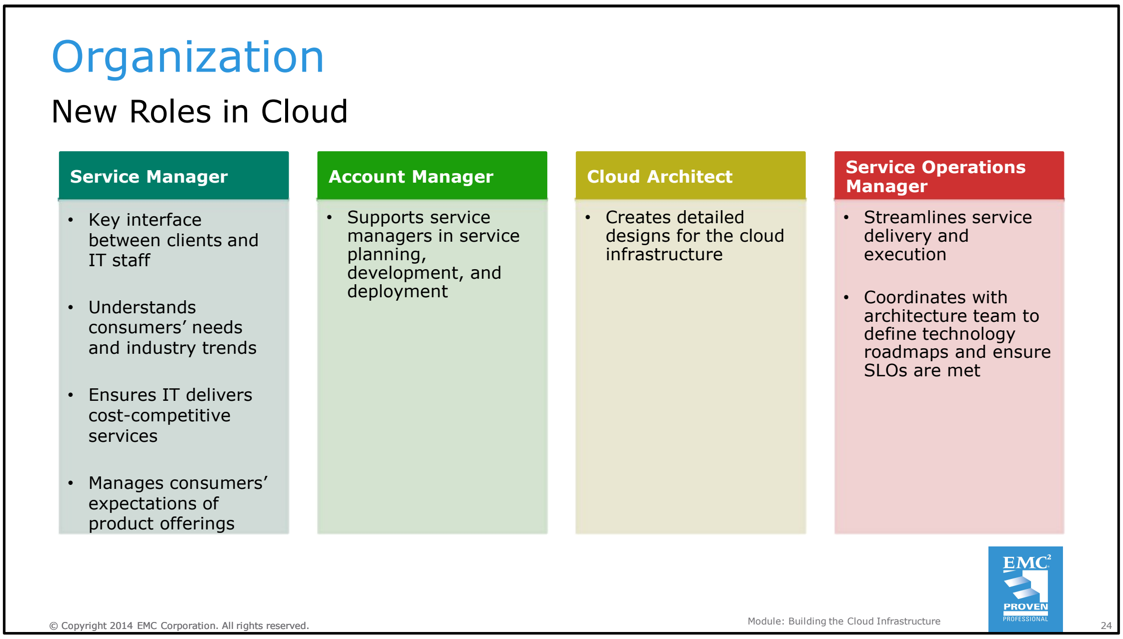

Organization

New Roles in Cloud

Service Manager

- Key interface between clients and IT staff
- Understands consumers’ needs and industry trends
- Ensures IT delivers cost-competitive services
- Manages consumers’ expectations of product offerings

Account Manager

- Supports service managers in service planning, development, and deployment

Cloud Architect

- Creates detailed designs for the cloud infrastructure

Service Operations Manager

- Streamlines service delivery and execution
- Coordinates with architecture team to define technology roadmaps and ensure SLOs are met

A cloud service provider needs to institute or transform the organization to a proactive and services-based model. This requires defining several new roles that perform tasks related to cloud services, such as service definition and creation, service administration and management, service governance and policy formulation, and service consumer management. Some of these tasks can be combined to become the responsibility of an individual or organizational role. A few examples of new roles required to perform tasks within a cloud environment include service manager, account manager, cloud architect, and service operation manager.

- A *service manager* is responsible for understanding consumers’ needs and industry trends to drive an effective product strategy. The service manager ensures that IT delivers cost-competitive services that have the features that clients need. The service manager is also responsible for managing consumers’ expectations of product offerings and serves as key interface between clients and IT staff.

- An *account manager* supports service managers in service planning, development, and deployment. The account manager maintains day-to-day contact to ensure that consumers’ needs are met.

- A *cloud architect* is responsible for creating detailed designs for the cloud infrastructure.

- The *service operations manager* is responsible to streamline service delivery and execution. Service operations manager is also responsible to provide early warning for service issues, such as emerging capacity constraints, or unexpected increase in cost. The service operations manager also coordinates with the architecture team to define technology roadmaps and ensure that service level objectives are met.

It is not only required to define the roles of IT staff and the skills they need, but it is also essential to identify the skill gaps that should be filled in order to successfully provide cloud services.

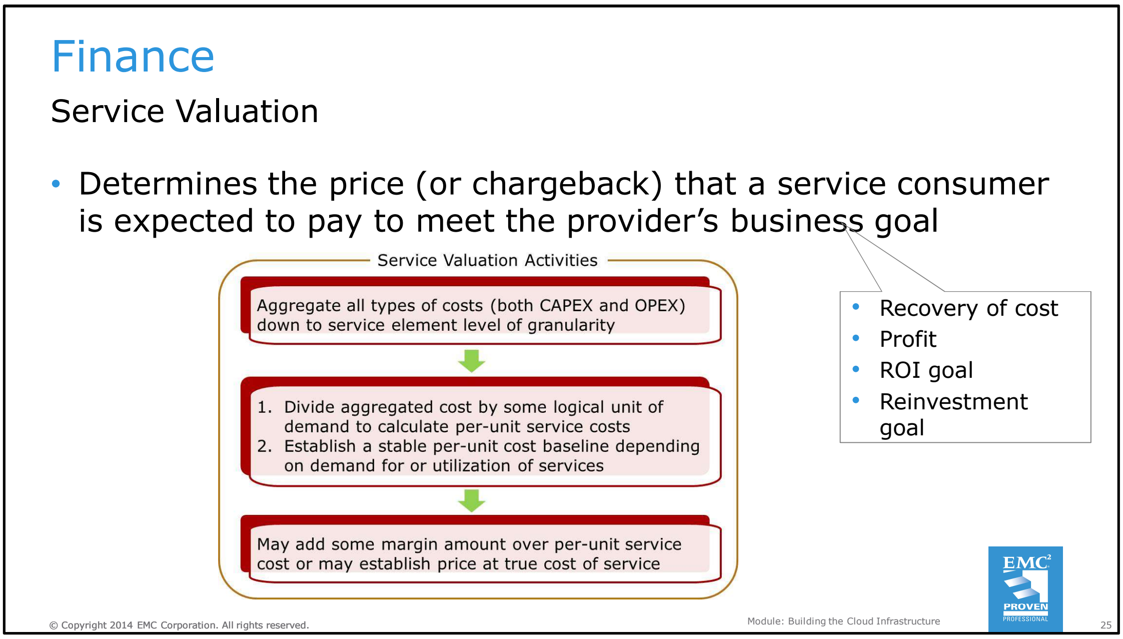

Finance

Service Valuation

- Determines the price (or chargeback) that a service consumer is expected to pay to meet the provider's business goal

Service Valuation Activities

Aggregate all types of costs (both CAPEX and OPEX) down to service element level of granularity

1. Divide aggregated cost by some logical unit of demand to calculate per-unit service costs
2. Establish a stable per-unit cost baseline depending on demand for or utilization of services

May add some margin amount over per-unit service cost or may establish price at true cost of service

- Recovery of cost
- Profit
- ROI goal
- Reinvestment goal

© Copyright 2014 EMC Corporation. All rights reserved.

Module: Building the Cloud Infrastructure

25

A service provider needs to institute or transform the financial/payback/showback/pricing model that will enable them to manage their budgeting, accounting, and chargeback requirements. The model helps the service provider to plan for investments to offer cloud services and determines the IT budget for cloud infrastructure and operations for the lifecycle of services.

The service provider should perform service valuation. Service valuation determines the price (or chargeback) that a consumer is expected to pay for a service, which helps recover the cost of providing the service, ensuring profitability, and meeting the provider's ROI and reinvestment goals. The service provider aggregates all types of costs (both CAPEX and OPEX) down to the service element level of granularity by mapping the elements to the relevant cloud services. Then it calculates the service costs on per-unit basis by dividing the aggregated cost for a service by some logical unit of demand such as GB of storage or an hour of usage for that service.

However, the per-unit service costs may vary over time, depending on the demand for or utilization of the services and service elements. Thus, service provider should track the demand and utilization to establish a stable per-unit cost baseline. Finally, the service provider may add some margin amount over per-unit service cost to define service price, or may establish the price at the true cost of service depending on the provider's business goal. The service provider then defines chargeback or showback model(s) based on the pricing strategy for cloud services.

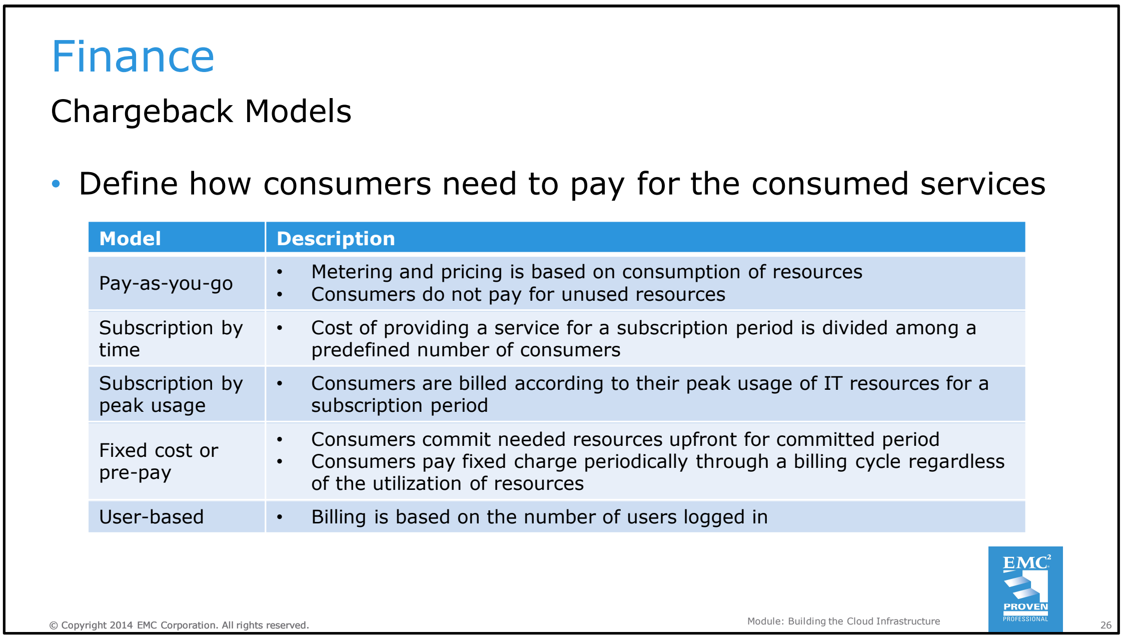

Finance

Chargeback Models

- Define how consumers need to pay for the consumed services

Model	Description
Pay-as-you-go	• Metering and pricing is based on consumption of resources • Consumers do not pay for unused resources
Subscription by time	• Cost of providing a service for a subscription period is divided among a predefined number of consumers
Subscription by peak usage	• Consumers are billed according to their peak usage of IT resources for a subscription period
Fixed cost or pre-pay	• Consumers commit needed resources upfront for committed period • Consumers pay fixed charge periodically through a billing cycle regardless of the utilization of resources
User-based	• Billing is based on the number of users logged in

A chargeback model defines how consumers need to pay for the consumed services. A list of common chargeback models along with their descriptions are provided below.

- Pay-as-you-go: Metering and pricing is based on the consumption of cloud resources by the consumers. Consumers do not pay for unused resources.

- Subscription by time: Consumers are billed for a subscription period. The cost of providing a cloud service for the subscription period is divided among a predefined number of consumers. For example, in a private cloud, if three business units are subscribing to a service that costs $60,000 a month to provide, then the chargeback per business unit is $20,000 for the month.

- Subscription by peak usage: Consumers are billed according to their peak usage of IT resources for a subscription period. For example, a provider may charge a consumer for their share of peak usage of network bandwidth.

- Fixed cost or pre-pay: Consumers commit up-front on the required cloud resources for the committed period such as one year or three years. They pay fixed charge periodically through a billing cycle for the service they use, regardless of the utilization of resources.

- User-based: Pricing is based on the identity of a user (a person) of cloud service. In this model, the number of users logged in is tracked and billed, based on that number.

Service provider deploys chargeback tools in the cloud infrastructure. These tools enable service provider to define a chargeback model. Based on the model, these tools automatically collect billing data, store billing records in a billing system, and generate the billing report per consumer.

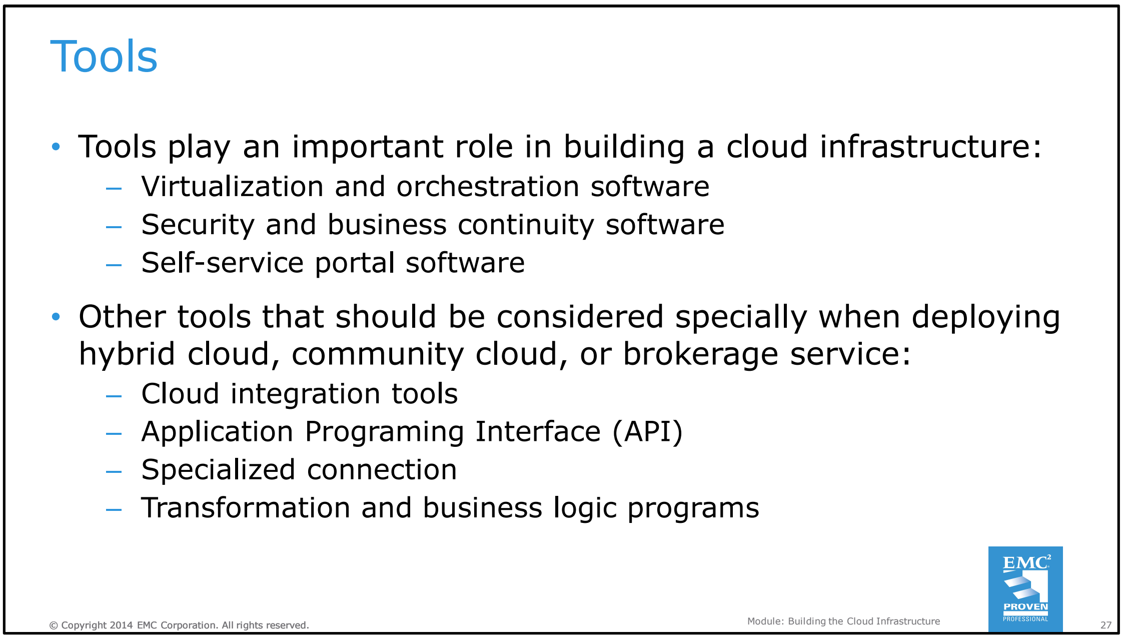

Tools

- Tools play an important role in building a cloud infrastructure:
  - Virtualization and orchestration software
  - Security and business continuity software
  - Self-service portal software
- Other tools that should be considered specially when deploying hybrid cloud, community cloud, or brokerage service:
  - Cloud integration tools
  - Application Programing Interface (API)
  - Specialized connection
  - Transformation and business logic programs

Tools play an important role in building cloud infrastructure; therefore an early step in building the infrastructure is to deploy the necessary technologies using the tools. Examples of the key tools used for building cloud infrastructure include virtualization software, orchestration software, security software, business continuity software, self-service portal software, and so on. These tools enable the service provider to build and offer cloud services to the consumers. Apart from considering the tools that enable the providers to build a cloud infrastructure, providers should also consider tools that will enable them to connect multiple clouds or applications. Examples of such tools include cloud integration tools, APIs, and specialized connection, transformation, and business logic programs. These types of tools are specially useful while deploying hybrid or community cloud. Also, such tools are important to consider when a service provider is providing brokerage services.

Cloud integration tools enable connecting cloud applications with other cloud and non-cloud applications to leverage the capabilities of multiple applications. Cloud integration technology integrates multiple cloud applications using application programming interface (API) support. These APIs enable secure access to the data of integrated applications. However, integration cannot be accomplished only with APIs because they do not perform functions such as transformation of data formats, data mapping, data validation, and error processing. These functions are typically handled by specialized connection, transformation, and business logic programs. These programs gather data with the help of APIs, then transform formats as required, and validate the accuracy of the transformation.

Consumers may avail different cloud services from multiple providers. In such cases, consumers may need assistance in selecting the providers that best meet their requirements. Moreover, using multiple cloud services from different providers may lead to operational complications and integration issues between the various services. Such issues have led to the emergence of cloud consumption assistance services known as *cloud services brokerage*, which are provided by cloud brokers.

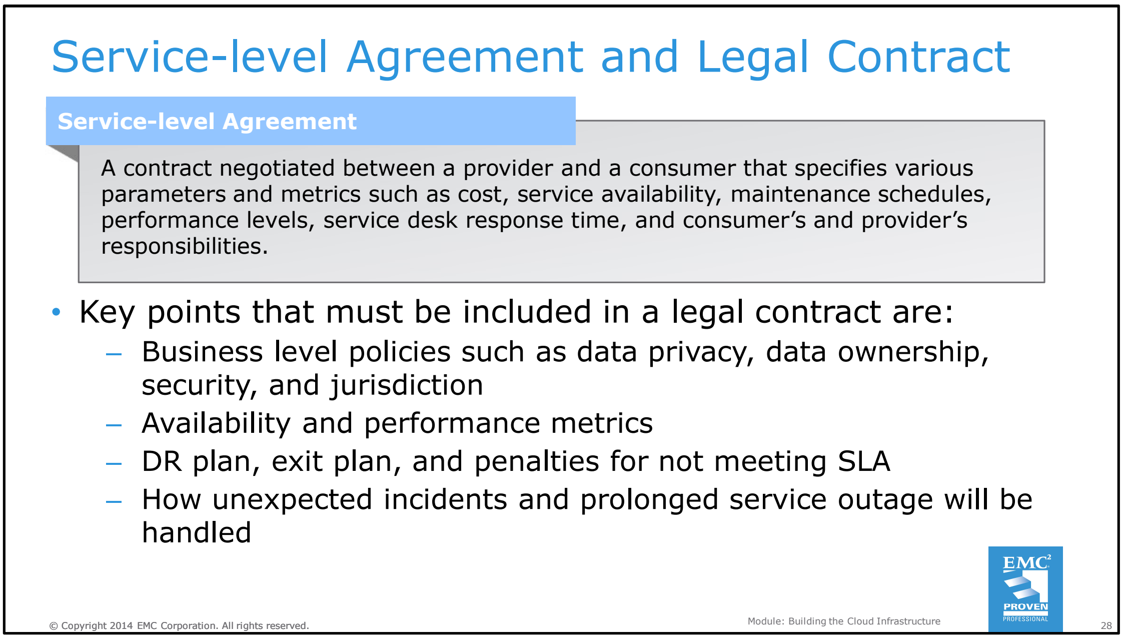

Service-level Agreement and Legal Contract

Service-level Agreement

A contract negotiated between a provider and a consumer that specifies various parameters and metrics such as cost, service availability, maintenance schedules, performance levels, service desk response time, and consumer’s and provider’s responsibilities.

- Key points that must be included in a legal contract are:
  - Business level policies such as data privacy, data ownership, security, and jurisdiction
  - Availability and performance metrics
  - DR plan, exit plan, and penalties for not meeting SLA
  - How unexpected incidents and prolonged service outage will be handled

A service-level agreement (SLA) is a contract negotiated between a provider and a consumer that specifies various parameters and metrics such as cost, service availability, maintenance schedules, performance levels, service desk response time, and consumer’s and provider’s responsibilities.

SLAs must be carefully written before offering to a consumer. SLAs are part of a service contract: an agreement between the cloud service provider and the cloud service consumer, stating the terms of service usage. A legal contract must be established with the consumer before a service can be used. When writing a legal contract, the key considerations include business level policies such as data privacy, data ownership, data retention, secure deletion, security, confidentiality, auditing, regulatory requirements, redundancy, jurisdiction, disruption resolution, compensation for data loss and misuse, excess usage, availability and performance metrics, payments and penalty methods, contracted services, a list of services not covered, licensed software, and service termination.

Finally, a disaster recovery plan, penalties, and an exit clause should be included. An SLA should include an indication of how an unexpected incident will be handled and what actions will be taken in case of a prolonged service outage. It should cover penalties for not meeting the SLA. The SLA should also include clauses related to the termination of the service by both the consumer and the provider.

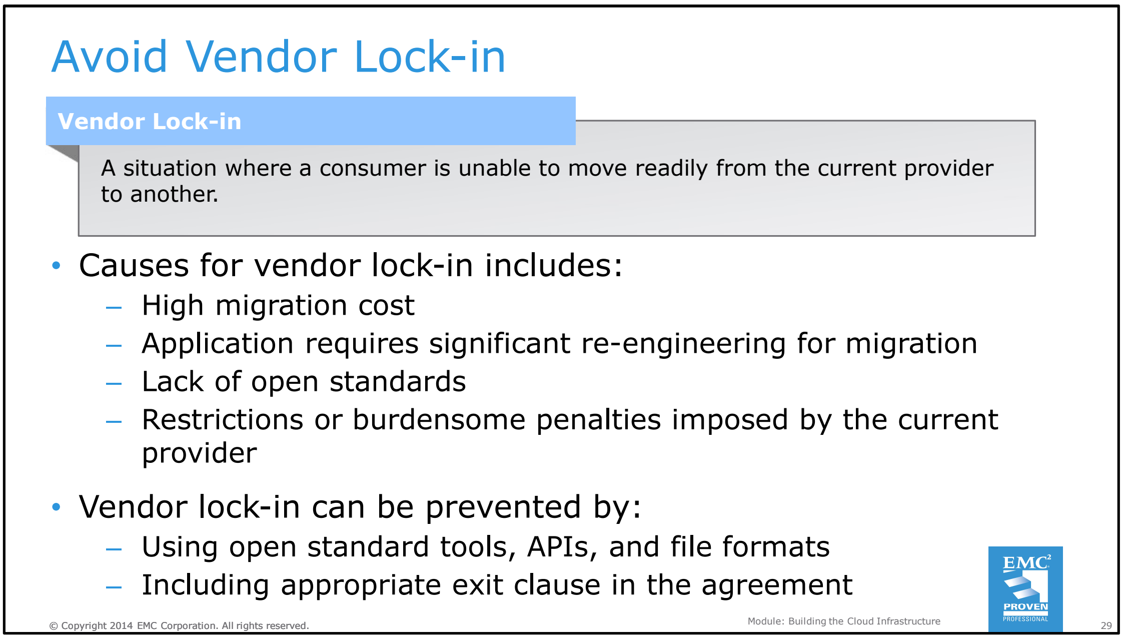

Avoid Vendor Lock-in

Vendor Lock-in

A situation where a consumer is unable to move readily from the current provider to another.

- Causes for vendor lock-in includes:
  - High migration cost
  - Application requires significant re-engineering for migration
  - Lack of open standards
  - Restrictions or burdensome penalties imposed by the current provider
- Vendor lock-in can be prevented by:
  - Using open standard tools, APIs, and file formats
  - Including appropriate exit clause in the agreement

Cloud vendor lock-in refers to a situation where a consumer is unable to move readily from the current provider to another. This condition may result from various causes such as high cost of migration, significant re-engineering effort requirement for an application migration, lack of open standards, or restrictions imposed by the current provider.

When building a cloud infrastructure, providers must avoid using proprietary tools, APIs, or file formats, which may cause vendor lock-in. The use of widely accepted open standard tools, APIs, and file formats not only prevent vendor lock-in, but also make services offered using open tools more acceptable to the consumers. The use of open standards provides interoperability and portability among providers, which consumers typically prefer. For example, the provider may use APIs based on the open standards that enable an application's data to migrate to another provider with minimal or no change to its format. Likewise, if the provider supports the use of Open Virtual Machine Format (OVF), which is an open standard for virtual machine format, then a virtual machine created in one of the provider's environment can be migrated to another provider with minimal or no changes.

Sometimes providers may impose restrictions or burdensome penalties for migrating to another provider, causing lock-in. Including an appropriate exit clause in the SLA can prevent vendor lock-in due to restrictions and penalties.

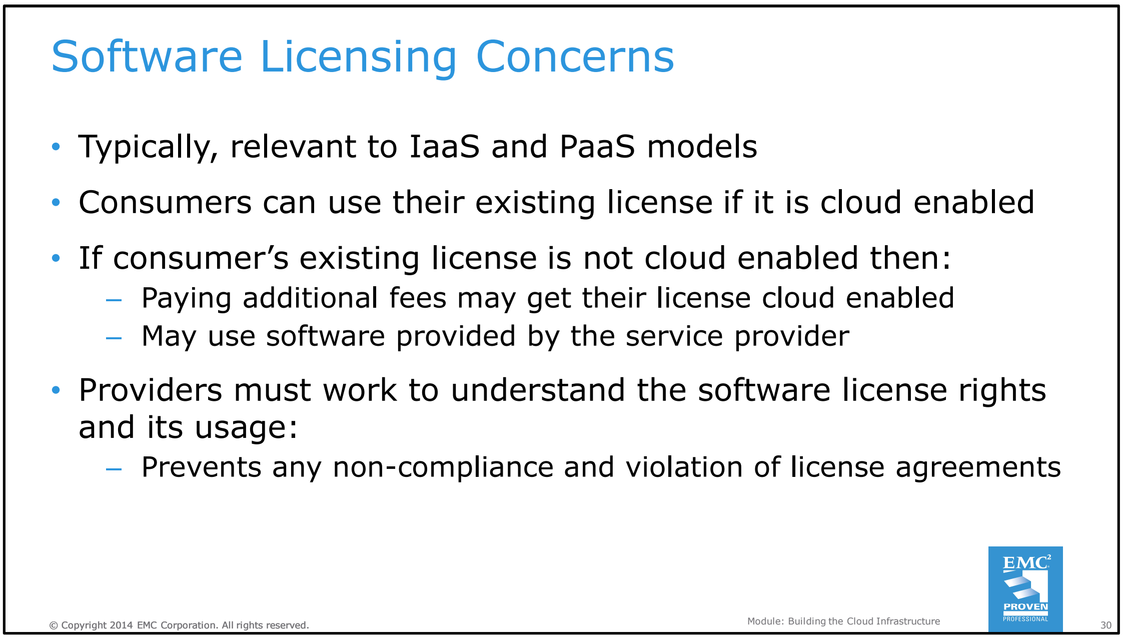

Software Licensing Concerns

- Typically, relevant to IaaS and PaaS models
- Consumers can use their existing license if it is cloud enabled
- If consumer’s existing license is not cloud enabled then:
  - Paying additional fees may get their license cloud enabled
  - May use software provided by the service provider
- Providers must work to understand the software license rights and its usage:
  - Prevents any non-compliance and violation of license agreements

While building a cloud infrastructure, providers must consider challenges associated with software (application and operating system) licenses. It is important to asses these challenges at an early stage. Software licensing challenges are relevant to infrastructure as a service (IaaS) and platform as a service (PaaS) models.

Consumers can use their existing software license in the cloud only if it is cloud enabled. Therefore, providers must identify whether the consumer’s existing software license is cloud enabled. If not, then the consumers can pay additional fees to get their license cloud enabled. Alternatively, consumers can use the software provided by the service provider and pay a fee for the software usage.

Further, the service provider in collaboration with the software vendors and consumers must work to understand the software license rights and its usage. This is important because the cloud service provider may have to create redundant systems by replication to combat against unplanned outage or disasters. Understanding the license rights and its usage will enable the service providers preventing any non-compliance and violation of the license agreement.

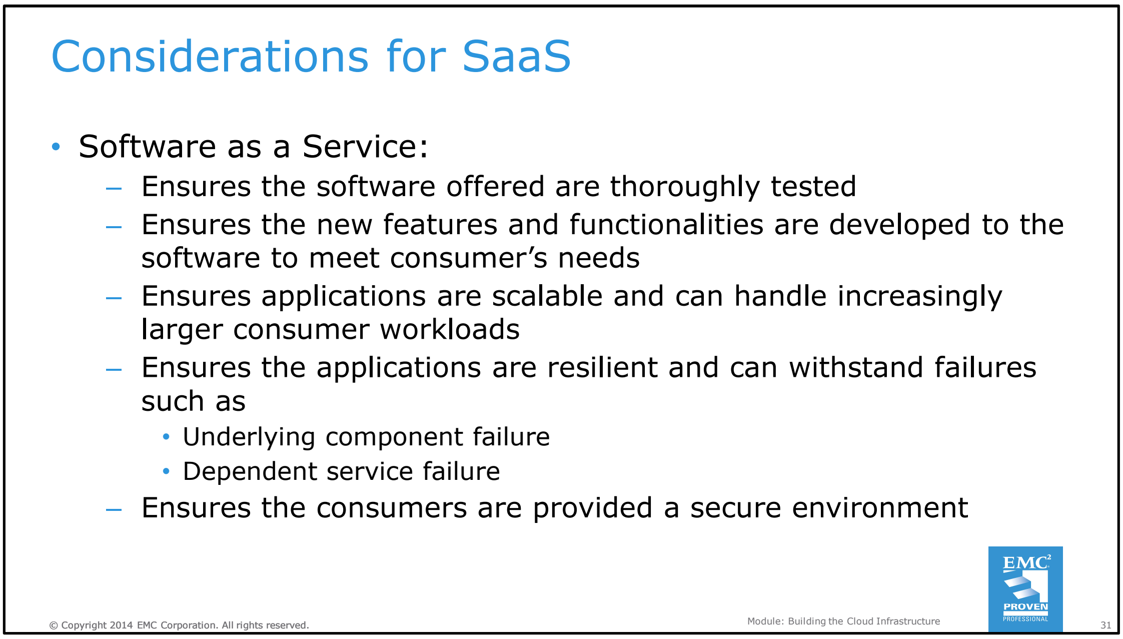

Considerations for SaaS

- Software as a Service:
  - Ensures the software offered are thoroughly tested
  - Ensures the new features and functionalities are developed to the software to meet consumer’s needs
  - Ensures applications are scalable and can handle increasingly larger consumer workloads
  - Ensures the applications are resilient and can withstand failures such as
    - Underlying component failure
    - Dependent service failure
  - Ensures the consumers are provided a secure environment

The slide lists the key factors that must be considered while deploying SaaS.

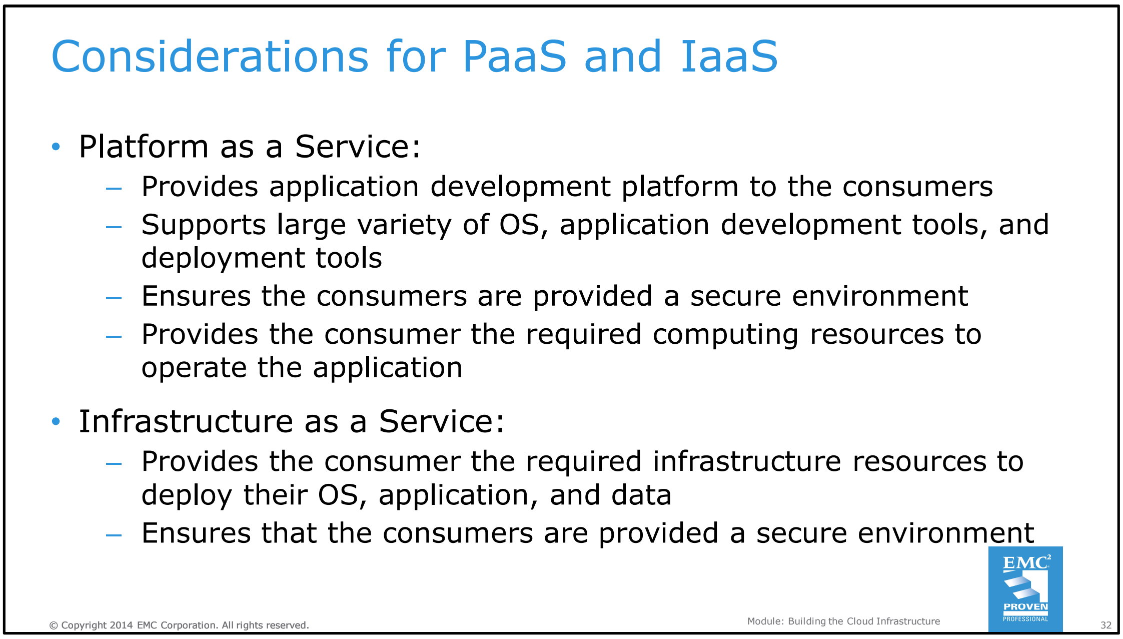

Considerations for PaaS and IaaS

- Platform as a Service:
  - Provides application development platform to the consumers
  - Supports large variety of OS, application development tools, and deployment tools
  - Ensures the consumers are provided a secure environment
  - Provides the consumer the required computing resources to operate the application
- Infrastructure as a Service:
  - Provides the consumer the required infrastructure resources to deploy their OS, application, and data
  - Ensures that the consumers are provided a secure environment

The slide lists the key factors that must be considered while deploying PaaS and IaaS.

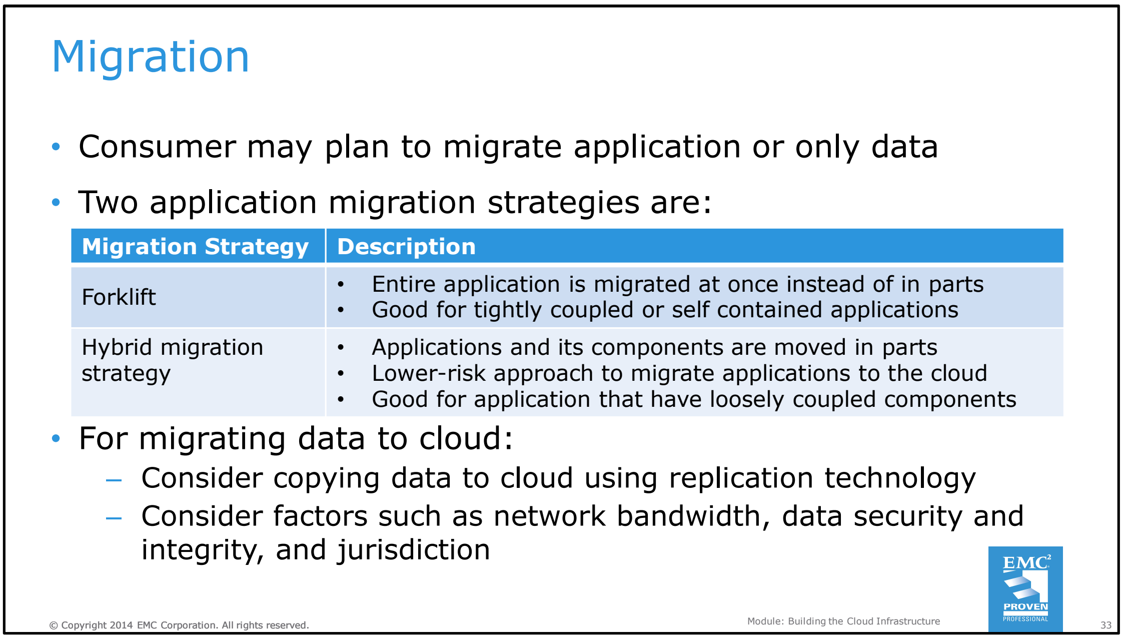

Migration

- Consumer may plan to migrate application or only data
- Two application migration strategies are:

Migration Strategy	Description
Forklift	• Entire application is migrated at once instead of in parts • Good for tightly coupled or self contained applications
Hybrid migration strategy	• Applications and its components are moved in parts • Lower-risk approach to migrate applications to the cloud • Good for application that have loosely coupled components

- For migrating data to cloud:
  - Consider copying data to cloud using replication technology
  - Consider factors such as network bandwidth, data security and integrity, and jurisdiction

Migration strategy and considerations depend on whether a consumer plans to migrate their application (in case of IaaS) or only their data (in case of SaaS).

For application migration, service providers must work with consumers to develop a migration strategy for their application. Also, they must identify the various dependencies of the application. For example, if an application depends on an authentication service that is on-premise, then appropriate configuration changes are required in order to make the application work after migrating to a cloud. Based on dependencies, a consumer may choose one of the two migration strategies. The strategies are forklift migration and hybrid migration.

- In the *forklift migration strategy*, the application and all of its related components are migrated to the cloud at once. This strategy is typically used for tightly coupled applications or self-contained applications. Tightly coupled applications are multiple applications that are dependent on each other and cannot be separated. Self-contained applications are applications that can be treated as a single entity.

- In a *hybrid migration strategy*, an application and its components are moved to the cloud in parts. This strategy is a lower-risk approach to migrate applications to the cloud. This is because parts of an application can be moved to the cloud and optimized before moving other parts. This reduces the risk of unexpected behavior of the application when it is moved to the cloud. This strategy is typically good for applications with many loosely coupled components.

In some cases, consumers may only require migration of data. The data can be migrated to the cloud by deploying replication technology to copy the data from the consumer's data center to the cloud. While migrating data to the cloud, the provider must consider the factors such as network bandwidth, data security, data integrity, data consistency, jurisdiction, and so on.

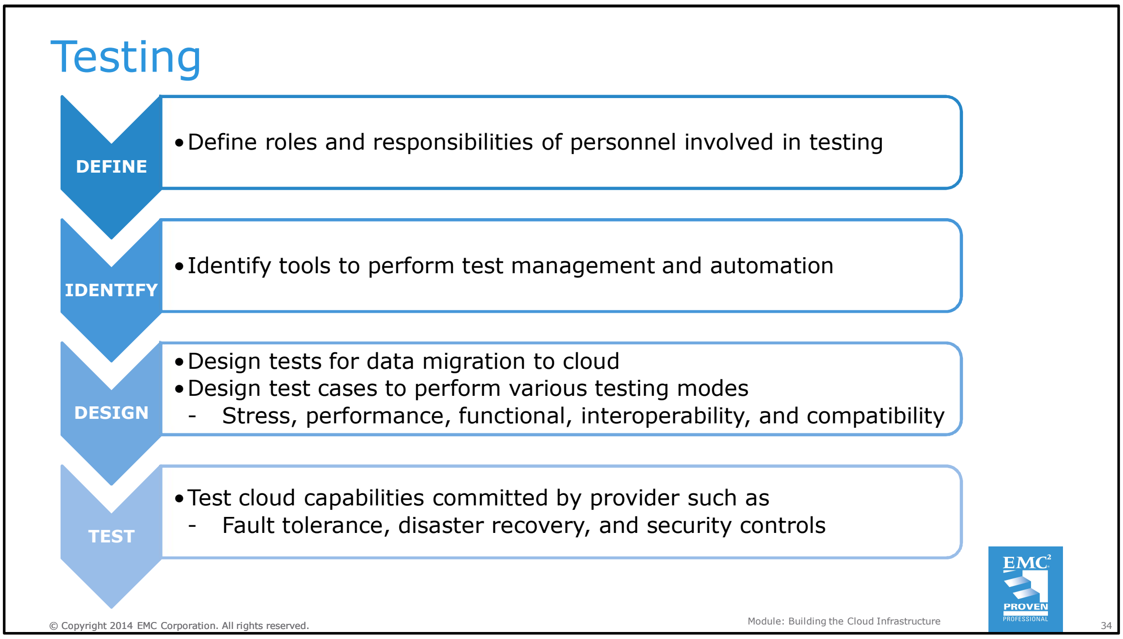

Testing

DEFINE
- Define roles and responsibilities of personnel involved in testing

IDENTIFY
- Identify tools to perform test management and automation

DESIGN
- Design tests for data migration to cloud
- Design test cases to perform various testing modes
  - Stress, performance, functional, interoperability, and compatibility

TEST
- Test cloud capabilities committed by provider such as
  - Fault tolerance, disaster recovery, and security controls

After the application or data is migrated to the cloud, the provider must work with the consumer to test their application to ensure that it is working as expected. The degree of testing may vary depending on the scope and magnitude of the consumer's requirements. While developing a test strategy, the providers in collaboration with the consumers must consider the following:

- Define roles and responsibilities of the personnel involved in test and quality assurance (QA) process
- Identify the tools required to perform test management and automation
- Design tests for data migration to the cloud
- Design test cases to perform various testing modes such as stress, performance, functional, interoperability, and compatibility

Apart from testing the application, the provider must also test other cloud capabilities such as fault tolerance, disaster recovery, security controls, and any other capabilities to ensure that the migrated application has successfully been configured with the capabilities that are committed by the provider.

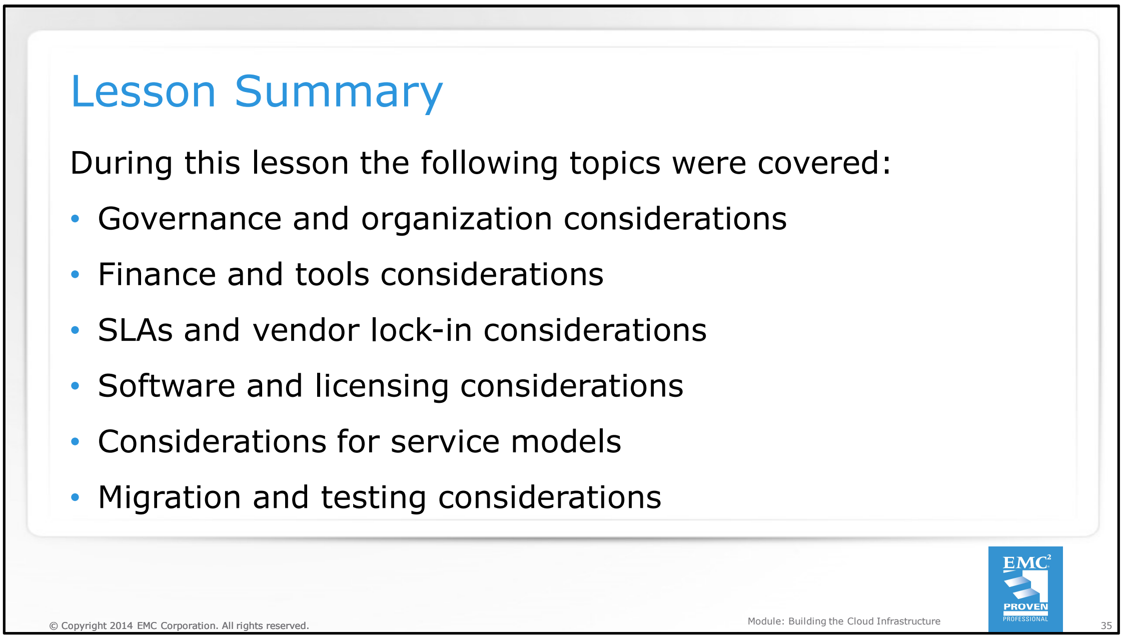

Lesson Summary

During this lesson the following topics were covered:

- Governance and organization considerations
- Finance and tools considerations
- SLAs and vendor lock-in considerations
- Software and licensing considerations
- Considerations for service models
- Migration and testing considerations

This lesson covered several factors that must be considered while building a cloud infrastructure.

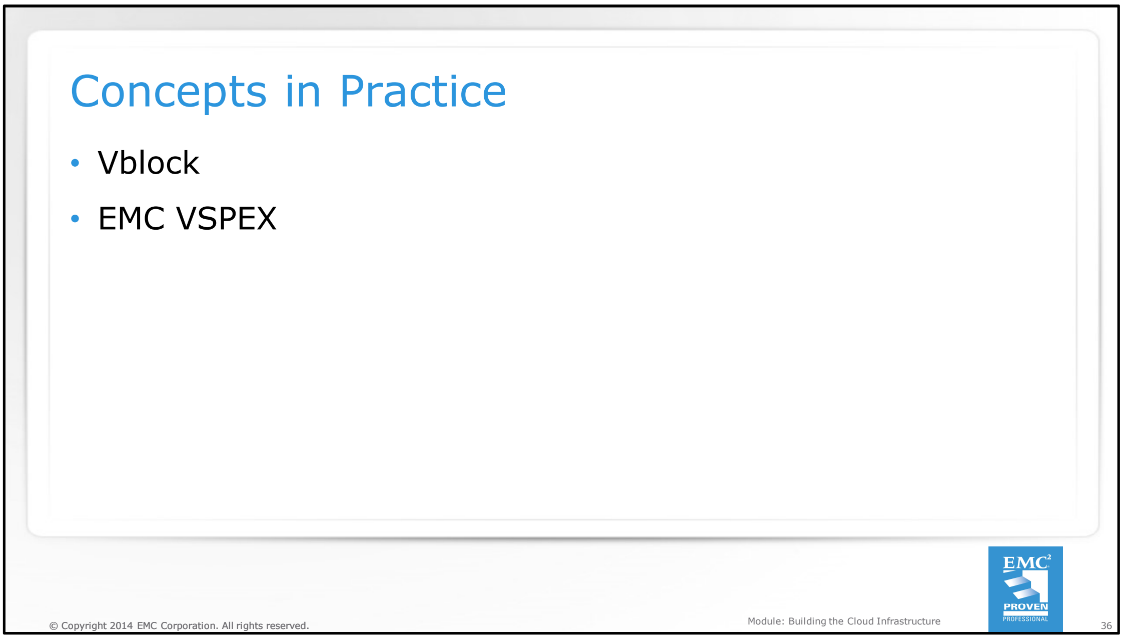

Concepts in Practice

- Vblock
- EMC VSPEX

The Concepts in Practice section covers two product examples Vblock and EMC VSPEX.

Note:

For the latest information on Vblock, visit www.vce.com.

For the latest information on EMC products, visit www.emc.com.

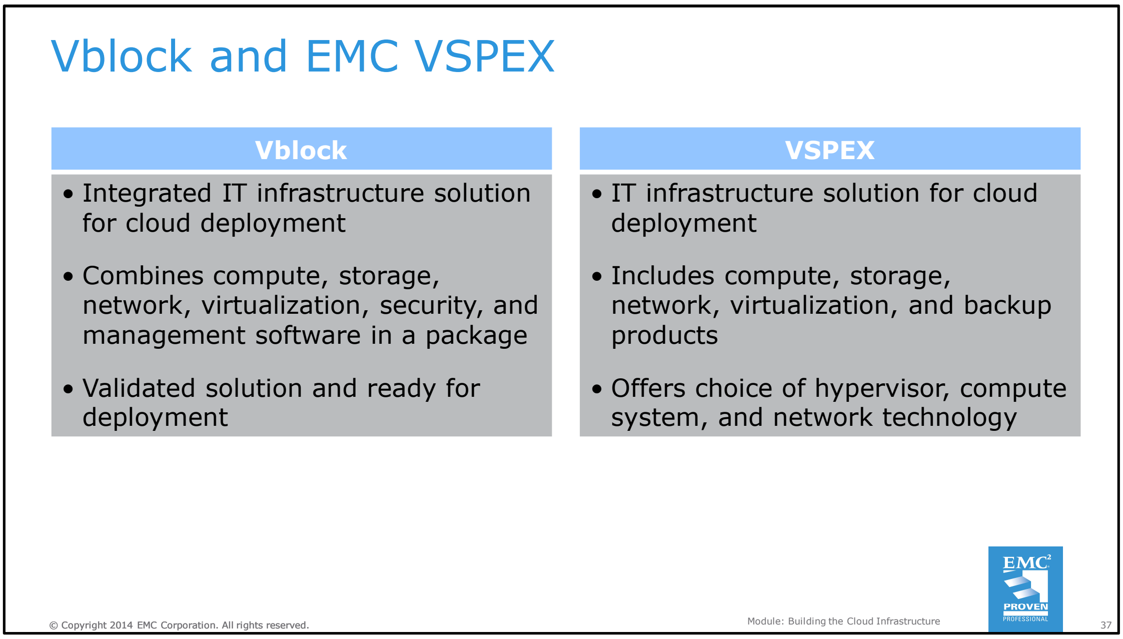

Vblock and EMC VSPEX

Vblock	VSPEX
• Integrated IT infrastructure solution for cloud deployment	• IT infrastructure solution for cloud deployment
• Combines compute, storage, network, virtualization, security, and management software in a package	• Includes compute, storage, network, virtualization, and backup products
• Validated solution and ready for deployment	• Offers choice of hypervisor, compute system, and network technology

Vblock is a completely integrated cloud infrastructure offering that includes compute, storage, network, and virtualization products. These products are provided by EMC (storage solution provider), VMware (virtualization solution provider), and Cisco (networking and compute solution provider), who have formed a coalition to deliver Vblocks.

Vblock is an integrated IT infrastructure solution that combines compute, storage, network, virtualization, security, and management software into a single package. This solution is a self-contained unit that accelerates deployment of a cloud infrastructure. Vblocks are pre-architected, preconfigured, pretested and have defined performance and availability attributes. Rather than the customers buying and assembling the individual IT infrastructure components, Vblock provides a validated solution and is factory-ready for deployment and production. This saves significant cost and deployment time associated with building a cloud infrastructure.

EMC VSPEX is an end-to-end virtualized infrastructure solution for cloud deployment, which includes compute, storage, network, virtualization, and backup products. The product vendors include EMC, Brocade, Cisco, Citrix, Intel, Microsoft, and VMware. VSPEX offers choice to the customers in terms of the hypervisor, compute systems, and networking components. Therefore, customers have the flexibility to choose the infrastructure components that fit their existing IT infrastructures.

EMC VSPEX is a complete virtualization solution that accelerates the deployment of cloud infrastructures. It provides the customers the flexibility to choose the hypervisor, compute system, and network technology they prefer along with EMC’s VNX and VNXe unified storage, EMC Data Domain, EMC Avamar, and EMC NetWorker backup and recovery solutions. Regardless of customer’s choice of hypervisor, compute system, and network technologies, validation of VSPEX by EMC ensures fast and low-risk deployment. VSPEX significantly reduces the planning, sizing, and configuration burdens that typically come with designing, integrating, and deploying a best-of-breed solution.

VSPEX, unlike Vblock, does not offer unified management. It comes with element management tools such as Microsoft System Center, VMware vCenter Operations Management Suite, and EMC Unisphere. But, it offers customers the choice of service elements that make up the solution. The tradeoff made for this freedom of choice is less integrated management.

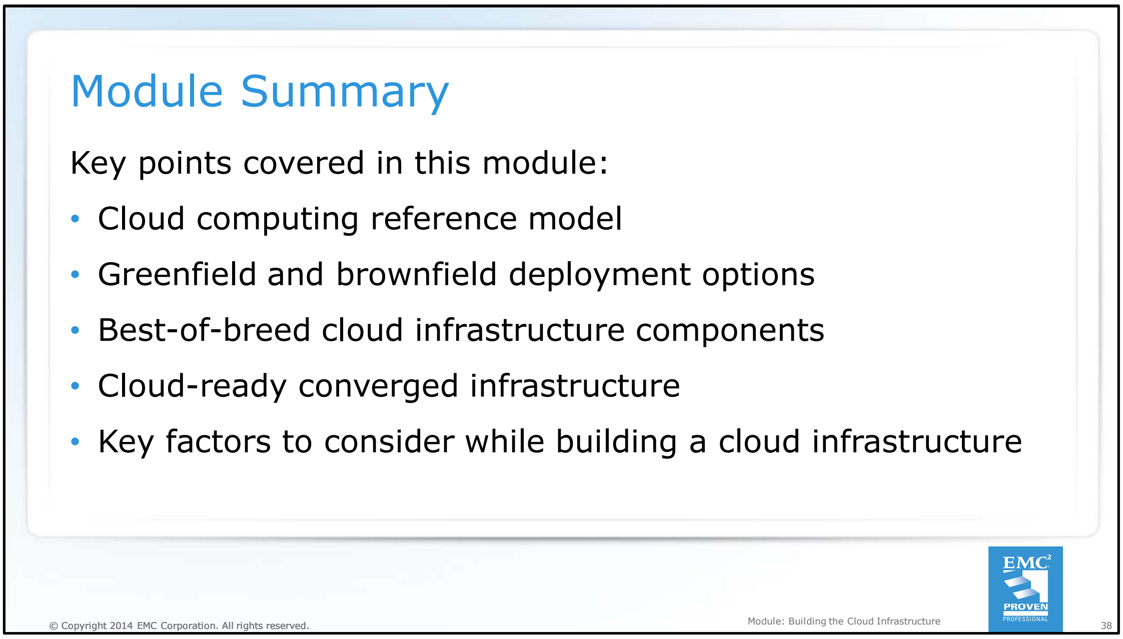

Module Summary

Key points covered in this module:

- Cloud computing reference model
- Greenfield and brownfield deployment options
- Best-of-breed cloud infrastructure components
- Cloud-ready converged infrastructure
- Key factors to consider while building a cloud infrastructure

This module covered the cloud computing reference model. It also covered the greenfield and brownfield deployment options. Further, it covered the two technology solutions—best-of-breed cloud infrastructure components and cloud-ready converged infrastructure—that can be used to build the cloud infrastructure. Finally, it covered the various factors to consider while building a cloud infrastructure.