Introducing the Cloud

Chapter 1

Episode 1.01

Setting up Your Cloud Accounts

Quick Review

- A credit card is used during cloud account setup to verify identity and process charges if required
- You may be required to use a cell phone and process a security code that is texted to you
- Some cloud providers stop service before charges incur and others do not, be careful

Episode 1.02

Defining the Cloud

Definition

- The NIST definition of cloud computing:

  Cloud computing is a model for enabling ***ubiquitous***, ***convenient***, ***on-demand*** network access to ***shared pools*** of configurable ***computing resources*** that can be ***rapidly provisioned*** and released with ***minimal management effort*** or service provider interaction.

Ubiquitous

- Available anywhere…
  - …there is an Internet connection

Convenient

- No server room required
- No power provisioning required
- No hardware installation required
- Sometimes, no software installation required

On-Demand

- Available when required
- Scheduled availability
- Created in minutes – not hours or days

Shared Pool

- Multi-tenant
  - More than one entity using a shared server
- Automatic prioritization

Computing Resources

- Processing
- Memory
- Storage
- Networking
- Special hardware

Rapidly Provisioned

- Cloud provisioning is simplified through:
  - Single-click launch
  - Template-based launch
  - Solution-based launch

Minimal Management

- Automatic updates/patch management
- Integrated monitoring and reporting
- Automatic scaling
- Scheduled availability

Quick Review

- Cloud computing is ubiquitous and available anywhere Internet access is available
- Computing resources include processing, memory, storage, networking, and specialized hardware
- Cloud solutions provide shared pools of resources with multiple tenants

Episode 1.03

Deployment Models

Cloud Deployment Models

- Deployed by:
  - Others (public/community)
  - Self (private)
  - Both (hybrid)

Selecting a Deployment Model

- Constraints and requirements drive selection
  - Examples:
    - Business policies
    - Functional (tasks)
    - Security (compliance)

Quick Review

- Public and community clouds are hosted online by a service provider
- Private clouds are hosted internally within a company
- Hybrid clouds are some combination of public/community and private cloud

Episode 1.04

Public Cloud Demo

Public Cloud

- Used by public consumers
- Pay-as-you-go/pay-as-you-grow
- Examples
  - Microsoft Azure
  - Amazon AWS
  - Google Cloud Platform (GCP)

DEMO

- Show interfaces for
  - Azure
  - AWS
  - Google Cloud

Quick Review

- Pay-as-you-go/Pay-as-you-grow is provided by public cloud providers
- AWS is the market leader in public cloud
- Azure and GCP are growing rapidly
- All public cloud providers offer a web-based and command line-based management interface

Episode 1.05

Private Cloud Demo

Private Cloud

- Implemented in-house
  - Hardware and software
  - Everything managed in-house
  - Behind a firewall
  - Possibly in a DMZ
- Based on virtualization
  - Virtual machines to rapidly deploy servers as needed
- Need extra physical hardware
  - Rapid deployment
  - Storage space
  - Compute
  - Services

DEMO

- OpenStack
- Cloudify

Quick Review

- Private clouds implement a virtualization layer and a management layer
- Private clouds provide storage, compute and services
- OpenStack and Cloudify are examples or cloud management software

Episode 1.06

Community Cloud Demo

Community Cloud

- Built for specific groups
  - Healthcare
  - Science
  - Education
  - Systems management (partner cloud)
- Sometimes blurred with simple SaaS solutions

DEMO

- coconstruct.com – SaaS/Community Cloud
- Penta.com – SaaS/Community Cloud

Quick Review

- Community cloud and SaaS are similar, but SaaS is a superset of community cloud
- Community clouds exist for healthcare, science, education, government and more
- Partner clouds are in the category of community clouds as well
- CoConstruct.com and Penta.com are examples of community clouds

Episode 1.07

Hybrid Cloud Demo

Hybrid Cloud

- Mixture of public and private cloud solutions
  - Data in the private cloud, processing in the public cloud
  - Processing in the private cloud, data in the public cloud
- Applications/APIs used to integrate the two

Hybrid Cloud

- Driving factors
  - Cost
  - Data
  - Security
- Hybrid cloud solution
  - Using applications in the cloud to create data stored in a private cloud

DEMO

- CloudBolt.io

Quick Review

- A hybrid cloud is a mixture of public and private cloud solutions
- Applications and APIs are available to integrate public and private clouds
- Cost, Data, and Security may be driving factors that assist in selection of a hybrid cloud model

Episode 1.08

Additional Models

Single Server

- A private cloud run by a single powerful server
- A single physical server in the cloud
- A single virtual server in the cloud

Single Cloud

- One cloud provides all services
- Small and medium businesses
- Partial security constraints are not in place

Multi-Cloud

- Multiple clouds
  - Services
  - Departments
  - Divisions/companies
  - Security constraints
- Orchestration platforms help manage multi-cloud deployments

Quick Review

- Single server cloud is a cloud implementation that uses a single, usually powerful, server
- Single cloud indicates that only one cloud account or service is used
- Multi-cloud indicates that multiple accounts or services are used

Episode 1.09

Cloud Architectures

Cloud Architectures

- Software as a Service (SaaS)
- Platform as a Service (PaaS)
- Infrastructure as a Service (IaaS)
- Everything as a Service (XaaS)

SaaS

- Software as a Service
  - Provides the software in the cloud
    - Google apps
    - Adobe cloud
    - Microsoft Office 365
    - Salesforce automation
    - Inventory management
    - Project management
  - May be API as a Service (AaaS)

PaaS

- Platform as a Service
  - Simple implementation of platforms
  - Runtimes, modules, components
  - Easy application deployment
  - Often used by developers and programmers

IaaS

- Infrastructure as a Service
  - Complete solution from hardware up
  - VMs and operating systems
  - Network configuration
  - Network services
  - Ex: AWS, Azure, Google Cloud Platform

XaaS

- Everything as a Service
  - Networking
  - Analytics
  - Artificial intelligence
  - Device management
  - Data extract, transform, load (ETL) operations
  - Ex: DBaaS (Database as a Service), WSaaS (Web Server as a Service)

Quick Review

- Examples of SaaS include Office 365 and Salesforce.com
- PaaS includes LAMP and other application development environments
- IaaS includes networking and security filtering features

Episode 1.10

Advanced Cloud Services

Internet of Things

- Connecting “things” to the network
  - May connect to the Internet
- Cloud Service Providers (CSPs) offerings:
  - IoT communication protocols
  - Application processing
  - Database storage
  - Aggregate IoT solutions

Artificial Intelligence

- **Artificial Intelligence (AI):** *ability* of a *machine* to *perform tasks* that normally require human-like *perception*, *cognition*, *decision-making*, and *action* in a *smart*, *efficient*, and *effective* way
- **Machine Learning (ML):** *subset* of AI that involves the development of *algorithms* and *statistical models* that enable a system to *learn* from *data* and *improve* its performance on a task *over time*

AI Cloud Offerings

- Image recognition
- Text recognition and natural language processing
- Data analysis and prediction
- Generative language models
- Automated decision support

Episode 1.11

Capacity, Elasticity, and Support Agreements

Capacity

- The workload capability of a system
  - Storage
    - Amount
    - Speed of read/write for I/O operation concurrency
      - IOPS – input/output operations per second
  - Networking
    - Speed of transfer/number of users
  - Processing
    - Speed of workload processing

Elasticity

- The ability to expand and contract as required
  - CPU resources
  - Storage
  - Servers
  - Threads/requests

Service Model Maintenance

- Metering mused to measure cloud resource consumption
- Chargeback can be made to departments other than IT
  - Each department is charged for their IT resources
- Pay-as-you-grow
  - Allows for low cost of entry

Service Model Maintenance

- Demand-driven service
  - Service to be provisioned based on current demand

Service Model Maintenance

- Responsibility
  - You manage what you put in the cloud
  - Cloud provider manages the cloud itself
  - Support agreements
    - SLAs
    - Response time
    - Service contact methods

Quick Review

- Capacity is the workload capability of a system and applies to storage, networking, and processing
- Processing may include CPU and memory
- Elasticity is the ability to expand and contract as required

Episode 1.12

Advanced Capacity Planning Concepts

Capacity Planning

- Define intentions
- Define requirements
- Create a design
- Test the design
- Implement the design
- Validate the design

Requirements Engineering

- Requirement types
  - Business
  - Stakeholder/User
  - Technical
    - Functional
    - Non-Functional
- Business > Stakeholder > Technical (System)
- BRS > StRS > SyRS/SRS

Trend Analysis

- System use changes over time
- Trend analysis explores the trajectory of change
- Allows for prediction of future capacity needs
  - Long-term predictions
  - Short-term predictions

Standard Templates

- Cloud providers offer prebuilt templates
- Cloud engineers can create templates
- Templates may be deployed in different ways
  - Manually
  - Scripted
  - Automated with scaling features

Licensing Requirements

- Know the license model
  - Per user, per machine, per cpu
- Know the license duration
- Ensure scalability

Episode 1.13

Resource Balancing

Resource Balancing

- Provides for the resources required at optimal cost
- Several techniques can be used
  - On-demand provisioning
  - Auto-scaling
  - Hybrid clouds
  - Serverless processing

Ideal Resources for the Cloud

- Public resources
  - Ex: your website
- Private scale resources
- Resources used for distributed work
  - Local access to resources across regions
- Anything requiring central processing

Poor Resources for the Cloud

- Anything requiring offline access
  - Possible solution: synchronize with the cloud
- Some security-related resources
  - Cloud provider might not be able to adhere to strict security requirements
- Low-latency demand resources
  - Edge processing may be better

Quick Review

- Capacity is the workload capability of a system and applies to storage, networking, and processing
- Processing may include CPU and memory
- Elasticity is the ability to expand and contract as required

Episode 1.14

Change Management

Change Management Components

- Advisory board
  - Usually an expert approves/rejects changes
- Approval process
  - When should something go through official approval process?
- Documentation
  - Change Management Database (CMDB)
  - Spreadsheet

Change Approval Process

- Submit a change request
  - Description
  - Positive reason for the change
  - Possible negative consequences
- Change reviewed by advisory board
- Approval/rejection
  - Implement or alter

Quick Review

- Advisory boards assist in determining if a proposed change should be implemented
- All changes should be documented for future reference
- Change approval processes include descriptions, change justification, potential problems, and review