AWS Solution Architecture Patterns

Objectives
Key objectives of this chapter ■ AWS reference architecture catalog ■ Overview of some AWS solution architecture patterns

1.1 AWS Architecture Center

- The AWS Architecture Center portal [http://aws.amazon.com/architecture/] offers a catalog of application architecture blueprints for solutions deployed in the AWS cloud
- The catalog comes with Architecture Datasheets and best practices that you can use as guides to implementing your solutions
- Based on your application's needs, you can select the most suitable reference architecture

1.2 List of Reference Architectures

- Some of the reference architecture blueprints offered by the architecture catalog are as follows:
  - Fault tolerance and High Availability (HA)
  - Web / Mobile-web Application Hosting
  - Log Analysis
  - Financial Services Grid Computing
  - Time Series Processing

1.3 High Availability Solution Architecture Blueprint

- AWS uses the Elastic Load Balancing service to achieve fault tolerance and higher levels of application availability
- For best results, Elastic Load Balancing should be used to balance user traffic across instances started in multiple availability zones (AZ's); in example below *AZ A* & *AZ B* are used

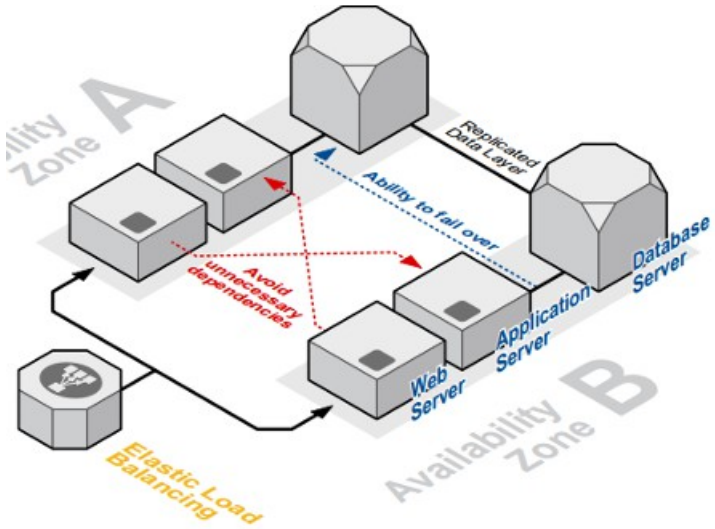

Notes:

Availability zones (AZ's) in the AWS cloud can be seen as logically different data centers; they are hooked to separate power grids that help minimize effect of power disruption in one AZ on the overall availability of your applications.

Your applications should be designed to minimize application state sharing between components running in different AZs.

Canada

821A Bloor Street West
Toronto, Ontario, M6G 1M1
1 866 206 4644
getinfo@webagesolutions.com

United States

436 YorkRoad, Suite 1
Jenkintown, PA, 19046
1 877 517 6540
getinfousa@webagesolutions.com

1.4 Log Analysis Solution Architecture Blueprint Summary

- For processing large volume of log files, AWS offers Elastic MapReduce service (EMR) which is backed up by a hosted Hadoop framework
- The EMR service is integrated with the S3 service for data input / output
- Integration with RDS for storing data processing results is available
- Optionally, clients can opt for using Spot instances (EC2 instances at a reduced cost) that are made available when Amazon EC2 has some underutilized computing capacity

1.5 Scalable Web App Solution Architecture Blueprint Summary

- This pattern help with hosting scalable web applications with HA quality of service
- This blueprint offers a reliable, scalable solution architecture that is also cost efficient under variable web traffic
  - Efficiency of the solution is achieved by using the Auto-scaling service that automatically adjusts the processing capacity up or down in correlation with the incoming traffic
- Fault tolerance is enabled by using Elastic Load Balancing that distributes incoming traffic for a cluster of EC2 instances deployed across two AZ's
- The DNS services are provided by Route53
- Application data is stored in RDS deployed in Master-Slave mode using cross- AZ data replication

Canada

821A Bloor Street West
Toronto, Ontario, M6G 1M1
1 866 206 4644
getinfo@webagesolutions.com

United States

436 YorkRoad, Suite 1
Jenkintown, PA, 19046
1 877 517 6540
getinfousa@webagesolutions.com

1.6 Simplified Web App Solution Architecture Blueprint

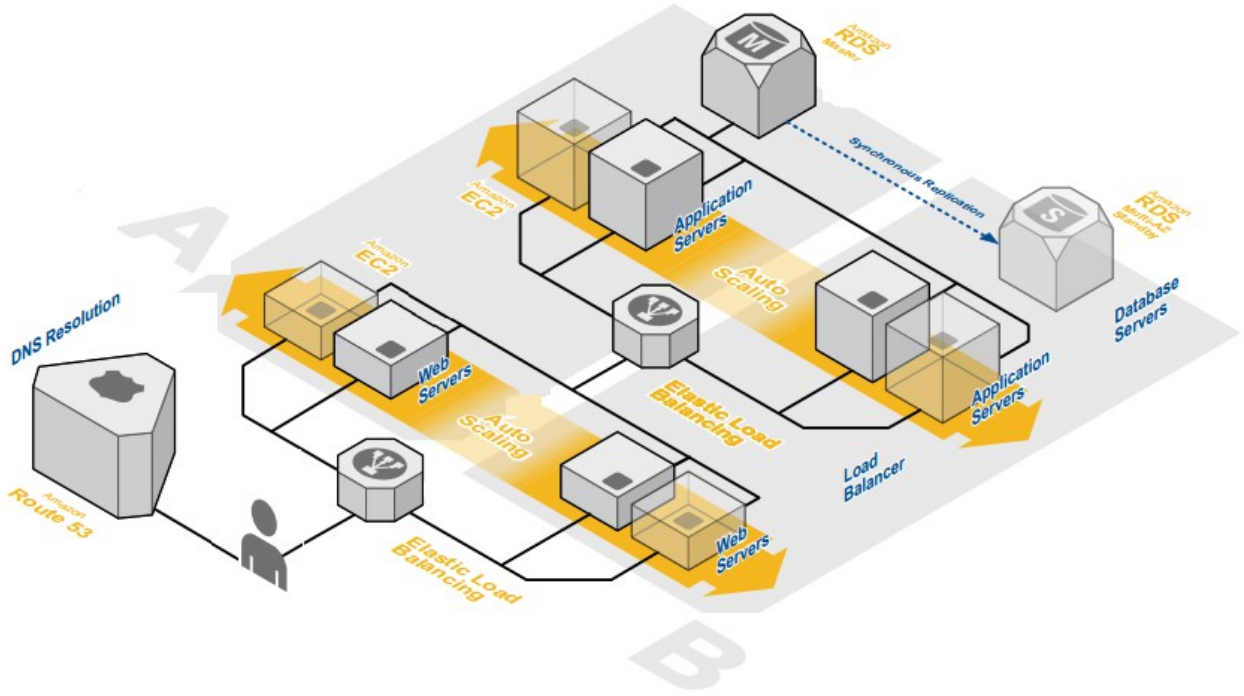

1.7 Architecting for AWS: Design for Failure - Take 1

- To design your solutions for failure, use the following mechanisms:
  - Design workflows and process that are interruption-tolerant and can resume on instance reboot
    - ✔ Make your application's design stateless; state should be outsourced to a centralized persistence store, if needed
  - Have an adequate backup and restore automation strategy in place

Canada

821A Bloor Street West
Toronto, Ontario, M6G 1M1
1 866 206 4644
getinfo@webagesolutions.com

United States

436 YorkRoad, Suite 1
Jenkintown, PA, 19046
1 877 517 6540
getinfousa@webagesolutions.com

1.8 Architecting for AWS: Design for Failure - Take 2

- Leverage AWS's multiple Availability Zones
- Use the Amazon CloudWatch service to monitor the health of your application
  - You can also use a number of available open source monitoring tools
  - Monitor the following system metrics of your applications: CPU, memory, Disk I/O, Network I/O
- Utilize the Auto scaling group to maintain a fixed number of healthy EC2 instances

1.9 Go with SOA and Asynchronous Communication Patterns

- The SOA design principles help you build loosely coupled components of your solutions that are more fault-tolerant and scale better
- One way to achieve service decoupling is to go asynchronous in component interactions
- You implement asynchronous communication patterns using messaging queues
- Messaging queues also help absorb (buffer) load spikes
- Amazon SQS offers a simple yet powerful message queuing system infrastructure

Canada

821A Bloor Street West
Toronto, Ontario, M6G 1M1
1 866 206 4644
getinfo@webagesolutions.com

United States

436 YorkRoad, Suite 1
Jenkintown, PA, 19046
1 877 517 6540
getinfousa@webagesolutions.com

1.10 Secure Your Applications

- Every EC2 instance can be protected by one or more simple firewall-like rules (referred to as security groups) for incoming network traffic

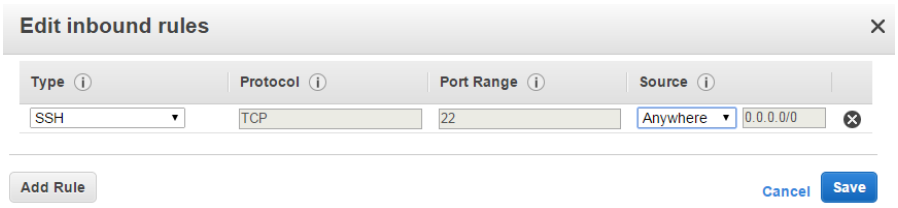

  - The above rule permits TCP traffic for SSH; public access from the Internet is allowed
  - You can narrow down source IP address ranges that can reach your EC2 instance using the CIDR notation
  - If you run Microsoft Windows, you may want to open the Remote Desktop Protocol (RDP) port 3389 for system administration
    - **Note:** In its default configuration, the RDP protocol is vulnerable to a man-in-the-middle attack; administrators must enable TLS to mitigate this risk

1.11 Securing your Web Application Example

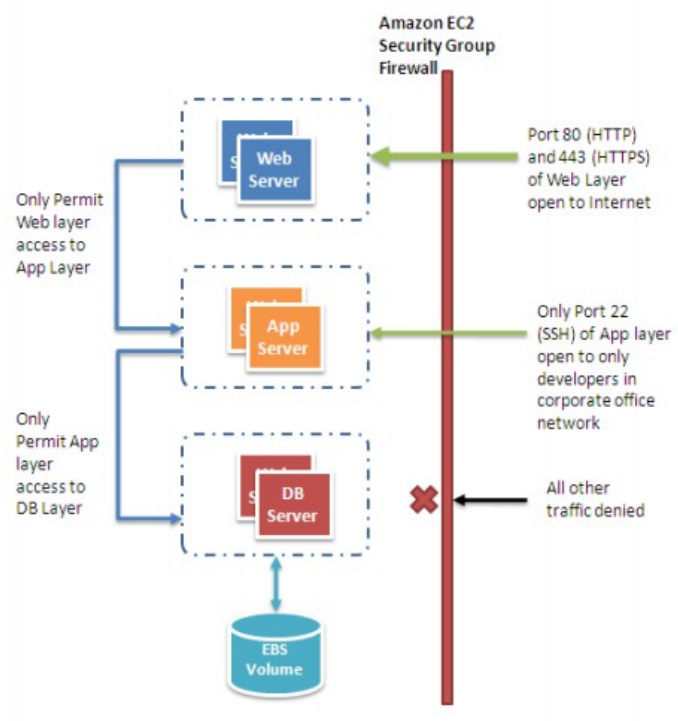

Source: Architecting for the Cloud: Best Practices by jvaria@amazon.com

1.12 Other Security Considerations

- On any IaaS platform, you can restrict inbound traffic by configuring software-based firewalls on your instances
  - Linux can use *netfilter* and *iptables*
  - Windows can use built-in firewall
- Basically, you need to adopt a "defense-in-depth" strategy which is about layering security from the outer layer (perimeter network) down to the inner layer (host firewall protection)

- Review the *Auditing Security Checklist for Use of AWS paper* (*AWS_Auditing_Security_Checklist.pdf*) published by AWS

1.13 Operational Checklists for AWS

- The Operational Checklists for AWS document published by Amazon (AWS_Operational_Checklists.pdf) can help you with regard to:
  - Evaluating your applications against a list of essential and recommended best practices
  - Reviewing operational and architectural aspects of your cloud solutions

Checklist	Intended Usage	Target Customer
Basic Operations Checklist	To help customers assess their application's use of specific services and features before they launch	Developers and system architects
Enterprise Operations Checklist	To assist enterprises identify key items to think about as they build a cloud migration and operational strategy	Enterprise architects
Auditing Security Checklist	To assist customers when they evaluate the security controls required by their specific industry or governing body like the AICPA, NIST, ISO, PCI SSC, etc.	Auditors or risk and compliance professional

1.14 Excerpts from Operational Checklists

- "We use AWS Identity and Access Management (IAM) to provide user-specific, rather than shared credentials for making AWS infrastructure requests."
- "Before sharing our customized Amazon Machine Images with others, we removed all confidential or sensitive information including embedded public/private instance key pairs and reviewed all SSH authorized_keys files"

- "Does the implemented AWS solution meet or exceed the application's high availability and resilience requirements?"
- "Does your organization have a configuration and change management strategy for its AWS resources?"

1.15 Summary

- The AWS Architecture Center portal offers a catalog of application architecture blueprints that you can use when building solutions for the AWS cloud
- You can choose the following reference architecture blueprints:
  - Fault tolerance and High Availability (HA)
  - Web / Mobile-web Application Hosting
  - Log Analysis
  - etc.
- You can secure your solutions in the AWS cloud by using security groups and by software-based firewalls

Canada

821A Bloor Street West
Toronto, Ontario, M6G 1M1
1 866 206 4644
getinfo@webagesolutions.com

United States

436 YorkRoad, Suite 1
Jenkintown, PA, 19046
1 877 517 6540
getinfousa@webagesolutions.com