

MQ on Cloud (AWS)

Suganya Rane

Digital Automation, Integration & Cloud Solutions

Agenda

- ❑ CLOUD Providers
- ❑ Types of CLOUD Environments
- ❑ Cloud Deployments
- ❑ MQ on CLOUD
- ❑ MQ on AWS
- ❑ MQ Monitoring on Cloud

What is Cloud Computing?

Cloud computing is the on-demand delivery of compute power, database storage, applications, and other IT resources through a cloud services platform via the internet with “pay-as-you-go” pricing.

How does Cloud Computing work?

- Cloud computing provides a simple way to access servers, storage, databases and a broad set of application services over the Internet.
- A Cloud services platform maintains the network-connected hardware required for these application services, while you provision and use what you need via a web application.

Advantages / Benefits?

1. Trade capital expense for variable expense
Pay when you consume computing resources and only pay for how much you consume.
2. Benefit from massive economies of scale
By using cloud computing, you can achieve a lower variable cost. This translates into lower “pay as you go” prices.

Advantages / Benefits?

3. Stop guessing capacity

Eliminate guessing on your infrastructure capacity needs.

4. Increase speed and agility

Reduce the time it takes to make those resources available to your developers from weeks to just minutes.

Advantages / Benefits?

5. Reduced \$\$\$ on running & maintaining data centers

Focus on business projects & customers, not the infrastructure.

6. Go global in minutes

Deploy your application in multiple regions around the world with just a few clicks. (low latency & minimal cost).

Types of Cloud Computing

1. Infrastructure as a Service (IaaS)

- Covers networking features, computers (virtual or on dedicated hardware), and data storage space.

2. Platform as a Service (PaaS)

- Covers the underlying infrastructure (usually hardware and operating systems)

3. Software as a Service (SaaS)

- Covers a completed product that is run and managed by a service provider.

Cloud Providers










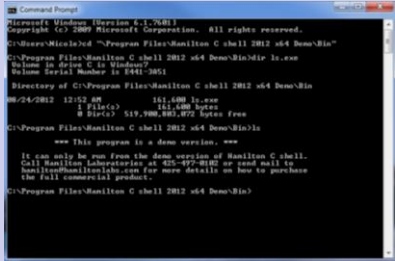
Cloud Computing Deployment Models

CLOUD	HYBRID	ON-PREMISE
A cloud-based application is fully deployed in the cloud and all parts of the application run in the cloud.	A hybrid deployment is a way to connect infrastructure and applications between cloud-based resources and existing resources that are not located in the cloud.	Deploying resources on-premises, using virtualization and resource management tools, is sometimes called “private cloud”.

Deployments Options

ENVIRONMENT PROVISIONING	MQ INSTALLATION	RUNNING MQ QUEUE MANAGERS
<ul style="list-style-type: none">• Create a VM or a container• Create Storage• Create Security• Setup Network	<ul style="list-style-type: none">• Installation• Fixpack Mgmt• Version Migrations• License Management	<ul style="list-style-type: none">• Create QMGRS• Integrating Apps to QMGRS• Failing over

Deployments Options

RESOURCE & ENVIRONMENT MANAGEMENT	MQ INSTALLATION	RUNNING MQ QUEUE MANAGERS
  	  	 

Benefits of Orchestration


- ✓ Forces best practices
- ✓ Simplified optimization
- ✓ Unified automation – Repeatability
- ✓ Improved visibility and control
- ✓ Self-service portal
- ✓ Automated chargeback calculations
- ✓ Long-term cost savings
- ✓ Facilitates business agility
- ✓ Deleting resources when they're no longer needed

IBM MQ Solutions

- ✓ Enterprise features
- ✓ 24x7, 365
- ✓ Multi-platform (z/OS, IBM I, Unix, Windows)
- ✓ Rich feature set (lots of options)
- ✓ Some static components (e.g. full repositories)
- ✓ Hybrid messaging – cloud to/from on-prem
- ✓ Inter-cloud communication
- ✓ Deployable anywhere that can host VMs or containers

Containers

- ✓ Docker
- ✓ IBM Containers
- ✓ Google Container Engine
- ✓ Azure Container Service



- ✓ Resource & Environment Management
 - ✓ Installation of MQ
 - ✓ Starting & Creating QMGRs

Containers

- Containers provide a similar environment to a VM but lighter in weight
- Linux containers
 - ✓ Containers all share the same OS kernel
 - ✓ Images are constructed from layered filesystems
 - ✓ Containers isolate applications from each other and the underlying infrastructure

MQ Docker Container

- MQ 8.0.0.4+ supported to run inside a Docker image
 - Details: <https://ibm.biz/mqdocker>
- Brings the benefits of Docker to MQ
 - Lightweight containers for running MQ
 - Predictable and standardized units for deploying MQ
 - Process, resource and dependency isolation
- IBM samples for customizing and building your own Docker images
 - Runs an MQ queue manager inside a container, isolated from the rest of your system

Persistent Storage

Reliability of storage

- Replicated across failure domains / availability zones? Are disk writes cached?
- What's the failure rate of disks?

Connecting to the right persistent storage


- When a queue manager's compute resource is moved (ex: run a container in a different VM), then something needs to connect the queue manager to the correct storage.
- For Example, the correct block storage volume, or directory on networked file storage.

Identifying the right persistent storage


- A very basic cloud orchestration setup could result in multiple instances of "QM1"

Persistent Storage Considerations

Region: US West 1
us-west-1 (N. Calif)




Region: US West 2
us-west-2 (Oregon)




Persistent Storage Considerations – Local SSD's

Region: US West 1
us-west-1 (N. Calif)




Region: US West 2
us-west-2 (Oregon)




Persistent Storage Considerations – Local SSD's

Region: US West 1
us-west-1 (N. Calif)




Region: US West 2
us-west-2 (Oregon)




Persistent Storage Considerations – Elastic Block Storage (EBS)

Region: US West 1
us-west-1 (N. Calif)




Region: US West 2
us-west-2 (Oregon)




Persistent Storage Considerations – Elastic block Storage

Region: US West 1
us-west-1 (N. Calif)



Region: US West 2
us-west-2 (Oregon)




Persistent Storage Considerations – Elastic block Storage

Region: US West 1
us-west-1 (N. Calif)




Region: US West 2
us-west-2 (Oregon)




Persistent Storage Considerations – Elastic File Storage

Region: US West 1
us-west-1 (N. Calif)




Region: US West 2
us-west-2 (Oregon)




Persistent Storage Considerations – Elastic File Storage

Region: US West 1
us-west-1 (N. Calif)



Region: US West 2
us-west-2 (Oregon)




Persistent Storage Considerations – Elastic File Storage

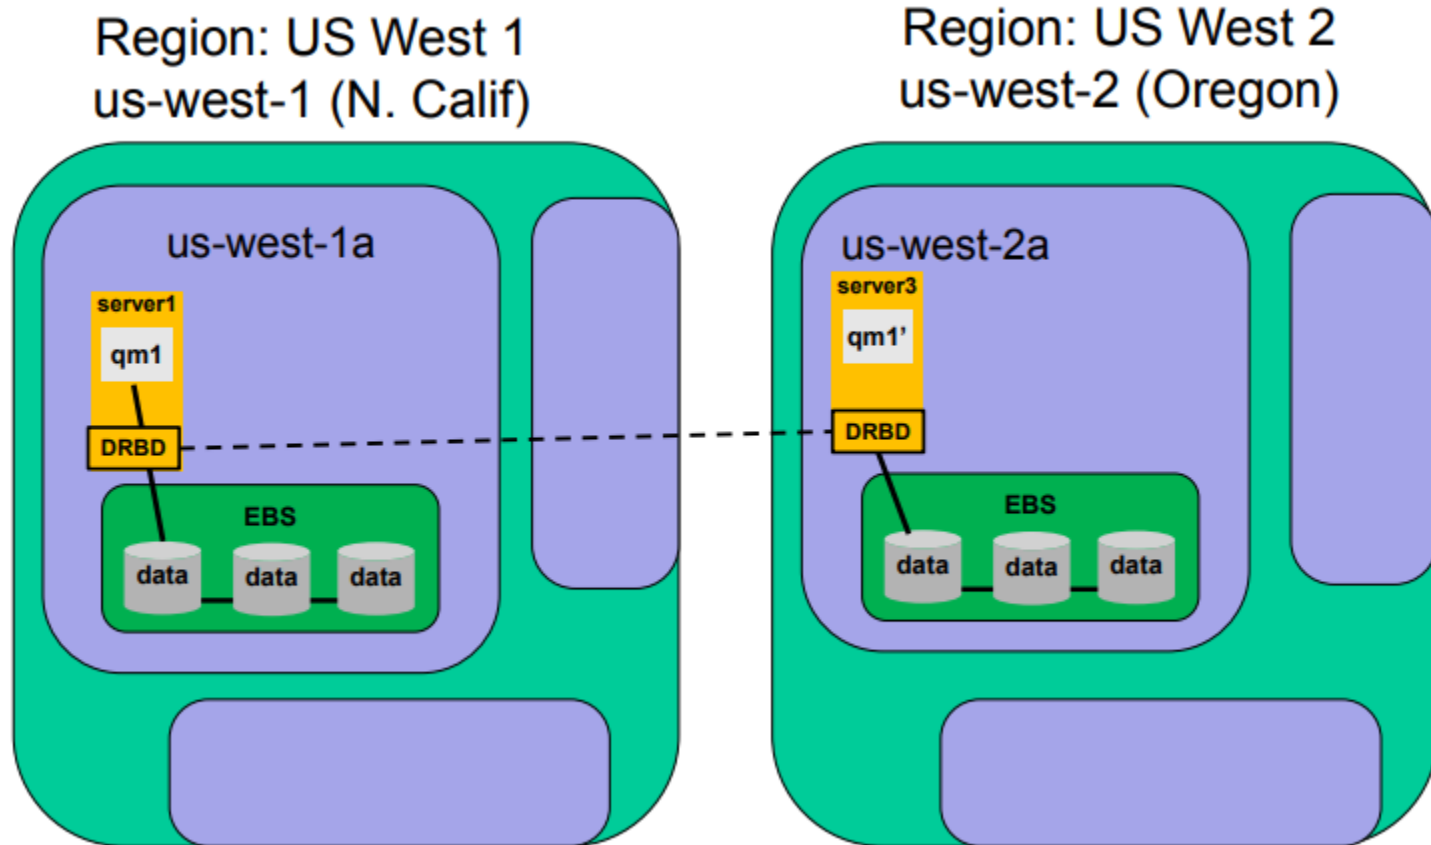
Region: US West 1
us-west-1 (N. Calif)



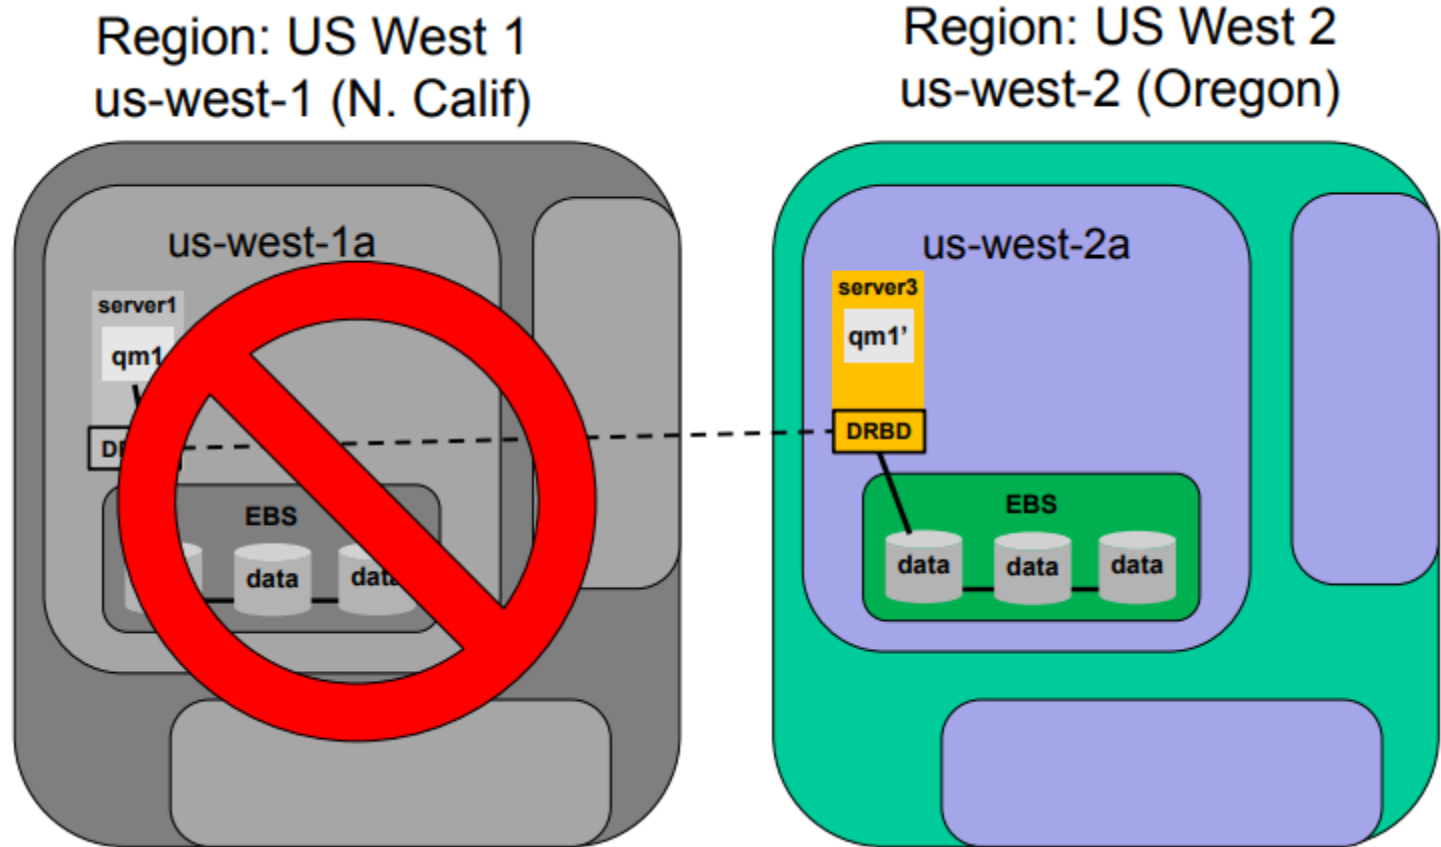
Region: US West 2
us-west-2 (Oregon)



Persistent Storage Considerations – Elastic Block Storage



Persistent Storage Considerations – Elastic Block Storage



QM Availability Decisions

- Single instance with failure detection and automatic restart
- Active-Passive, warm instance waiting in another availability zone to take over in event of failure
- Active-Active, workload balancing between instances in multiple availability zones

Service Delivery

- Queue manager IP addresses more likely to change in a dynamic/cloud environment
- Most cloud providers offer manual management of IP addresses
 - Amazon Elastic IP
 - Azure Reserved IP
 - Google Cloud Engine Reserved IP
- DNS can provide basic service discovery for clients
 - Amazon Route 53
 - Azure DNS
 - Google Cloud DNS

Service Delivery


- MQ channels need to know where they're connecting to?
- MQ clusters instances tell each other about their own IP addresses
- Client configuration – how do clients find the right queue manager

Data Ingress Vs Egress

- What about data entering/leaving the Cloud?
- What level of availability do you need between Cloud and on-premise?
- How is connectivity secured? (Vpn, Secure Tunnel, Tls/Tcp)
- Performance
- How is data workload balanced between Cloud & on-prem?

Cloud Ingress/Egress – MQ Clusters

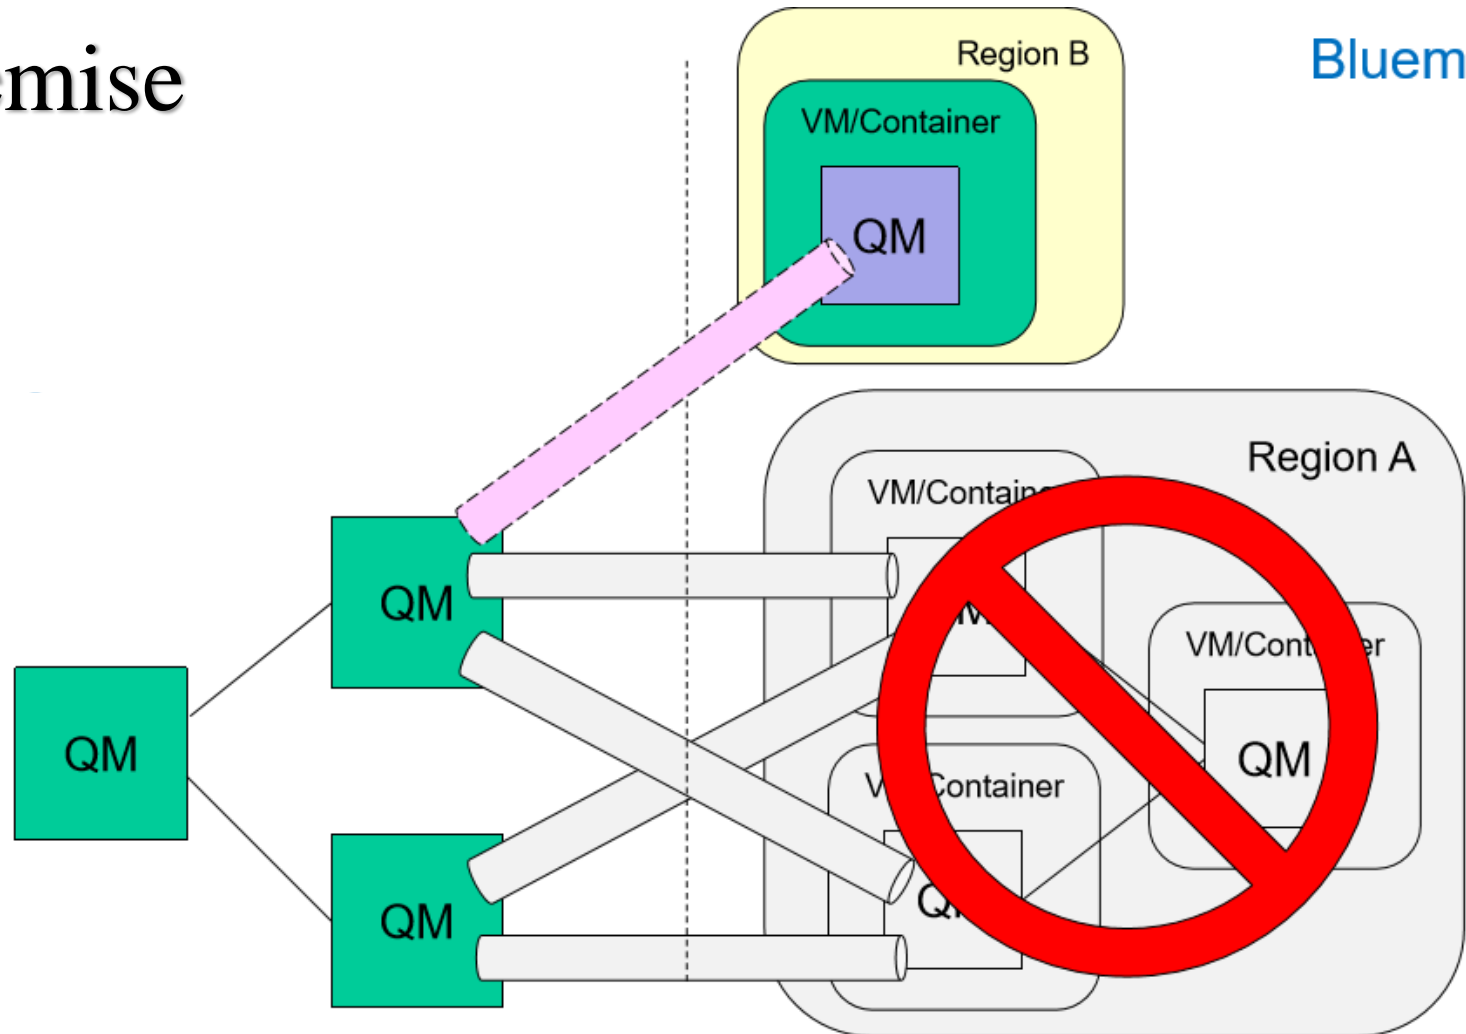
On-Premise



Cloud Ingress/Egress – MQ Clusters

On-Premise

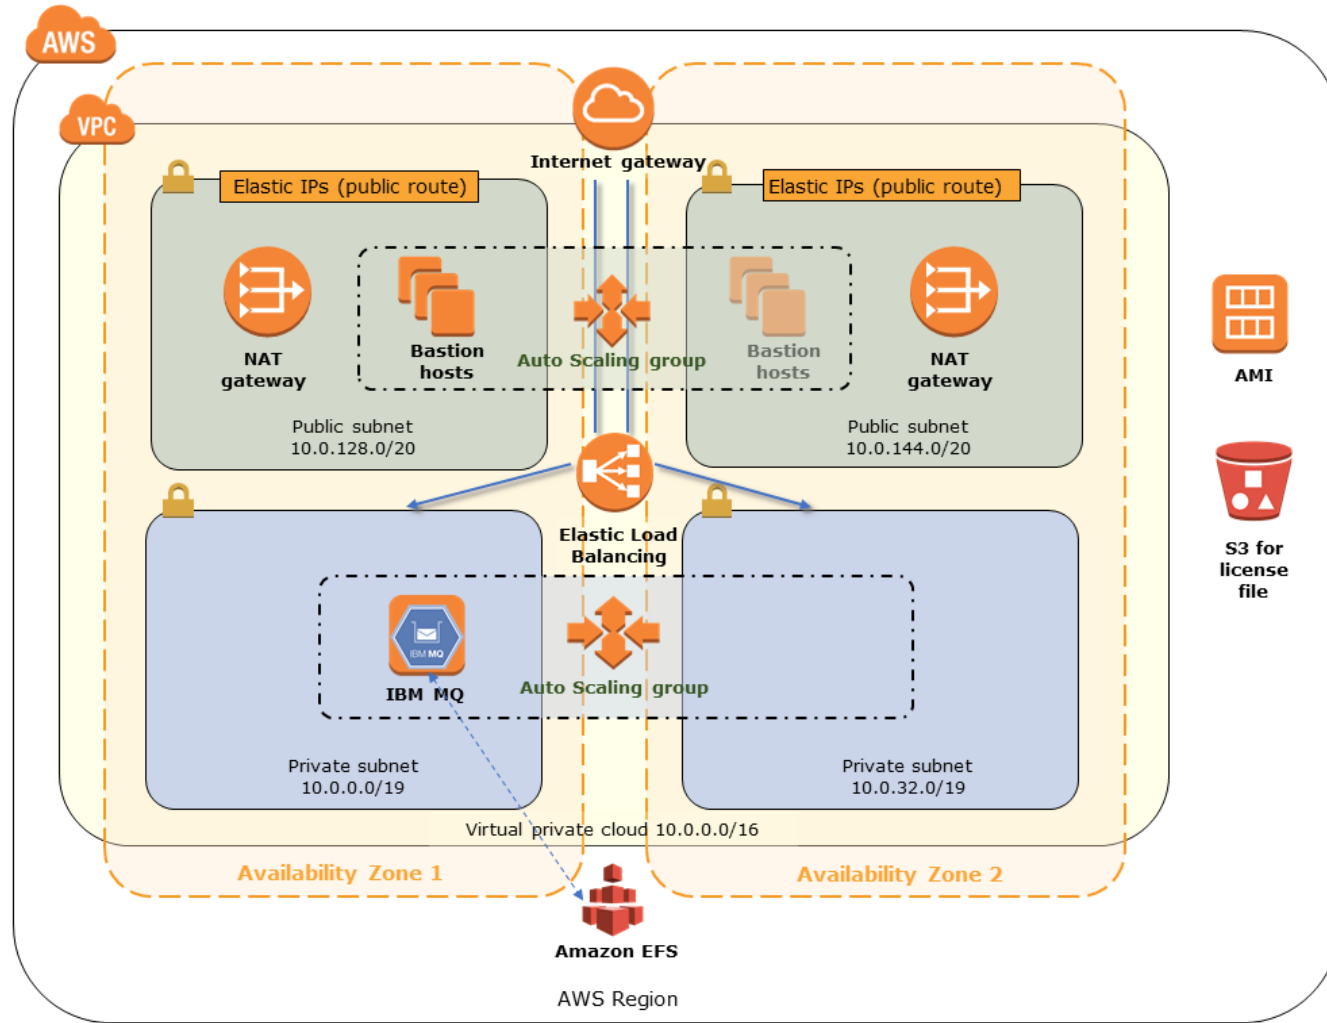
Bluemix



MQ on AWS Options

- Quick Start Guide for AWS (MQV9)
- Partner World Licensing (non-prod environments)
- BYOL from your current license stack
- Purchase of MQ from AWS Marketplace (MQV8)
- MQ Clients on AWS (connecting back to your on-prem qmgrs) – no software license costs!

MQ on AWS (the quick start with trials)



MQ from the AWS MarketPlace

MQ from the AWS Marketplace

This AMI has WebSphere MQ v 8.0 installed on Red Hat Enterprise Linux 7.5 base.

Highlights:

- Install WebSphere MQ in just a few minutes
- Ideal for Development, Test and peak production workloads
- If the version you require is not listed, you can request it.

COSTS

Hourly Cost:

EC2 Instance type	Software/hr	EC2/hr	Total/hr
<input checked="" type="radio"/> m3.medium ★ Vendor Recommended	\$0.8	\$0.127	\$0.927

Annual Cost:

EC2 Instance type	Software/yr	EC2/hr	Percent Savings (%)
<input checked="" type="radio"/> m3.medium ★ Vendor Recommended	\$5,957	\$0.127	15%

AWS Free Tier

12 months free and always free products

AWS Free Tier includes offers that expire 12 months following sign up and others that never expire.

Amazon EC2

750 hours per month of Linux, RHEL, or SLES t2.micro instance usage

750 hours per month of Windows t2.micro instance usage

Amazon QuickSight

1 GB of SPICE capacity 1 user perpetual free tier

10GB of SPICE capacity the first 2 months for free for a total of 4 users

Amazon CloudWatch

10 Custom Metrics and 10 Alarms

1,000,000 API Requests

5GB of Log Data Ingestion and 5GB of Log Data Archive

3 Dashboards with up to 50 Metrics Each per Month

Amazon RDS

750 Hours per month of db.t2.micro database usage (applicable DB engines)

20 GB of General Purpose (SSD) database storage

20 GB of storage for database backups and DB Snapshots

Amazon S3

5 GB of Standard Storage

20,000 Get Requests

2,000 Put Requests

AWS Lambda

1,000,000 free requests per month

Up to 3.2 million seconds of compute time per month

AWS ElasticSearch

The ElasticSearch stack is a useful way to store, search and visualize data.

The ELK stack comprises 3 components:

- ElasticSearch - A distributed search engine
- Logstash - a log collection & processing tool
- Kibana - a browser based visualization tool for displaying and searching log data

Kibana

- You can add and remove different components to the dashboard
- You can save dashboards once you're happy with their layout
- Create various charts and graphs of your data.

CloudWatch


- CloudWatch allows you can search individual log streams (per instance), and set up filters and alarms for certain events.
- The search facilities offered aren't anywhere near as advanced as ElasticSearch, but you can choose to forward log events on to ElasticSearch if you want to.
- Importantly, you can also trigger AWS Lambda functions based on certain messages being logged, which will allow you to take programmatic action.
- You can use either solutions and also add email functionality to automatically indicate when specific errors occurs.

Prometheus

- An open-source monitoring and alerting solution, whose particular strengths:
 - the collection of time series data, with the ability to easily query that data.
 - it provides libraries in several languages to enable products to export data to it.
 - it works by pulling information from exporters such as this MQ program at configured intervals over an HTTP connection.

Grafana Dashboards

As an example, this dashboard is looking at several items including the same queues as above, and CPU and logging information:



Cloud Monitoring Summary

- These slides and presented material show how the statistics generated by MQ can easily be used in some of the monitoring packages that are commonly used with various cloud and container-based systems.
- The MQ data can be integrated with other metrics to give a complete view of your environment.

Useful Links

- MQ on Cloud – IBM Messaging

<https://developer.ibm.com/messaging/mq-on-cloud/>

- GitHub Repositories

<https://github.com/ibm-messaging/mq-aws>

- IBM MQ - Using AWS CloudWatch to monitor queue managers

https://www.ibm.com/developerworks/community/blogs/messaging/entry/IBM_MQ_Using_CloudWatch_to_monitor_queue_managers?lang=en

- IBM MQ - Using Prometheus & Grafana to monitor qmgrs

https://www.ibm.com/developerworks/community/blogs/messaging/entry/IBM_MQ_Using_Prometheus_and_Grafana_to_monitor_queue_managers?lang=en

Useful Links

- Storing and searching MQ error logs in Elasticsearch
https://www.ibm.com/developerworks/community/blogs/messaging/entry/Storing_and_searching_MQ_logs_in_Elasticsearch?lang=en
- MQ on AWS: Sending MQ error logs to CloudWatch
https://www.ibm.com/developerworks/community/blogs/messaging/entry/mq_aws_cloudwatch_logs?lang=en

Grafana Support For Prometheus

<https://prometheus.io/docs/visualization/grafana/>

Questions & Answers

