### **Planning for MQ in the Cloud**

**Rob Parker, IBM** 

parrobe@uk.ibm.com

MQ Technical Conference v2.0.1.7

10/2/2017

### **Please note**

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.

Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.

The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver

any material, code or functionality. Information about potential future products may not be incorporated into any contract.

The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

### Agenda

### Introduction

- Generic Cloud considerations
- MQ Specific considerations
- Questions



### **GENERIC CLOUD CONSIDERATIONS**





### **Characteristics of a cloud environment**

- Self-service empowers users to provision resources without requiring human intervention, most likely using a web-based portal or an API.
- Elastic scaling enables scaling up and down on demand, driving the need for high levels of automation.
- Shared resources offers economies of scale through the use of shared infrastructure and software, securely separating the resources at a logical level.
- Metered usage allows pay-as-you-go billing through monitoring, measurement and reporting of usage.



### **Types of cloud environments**

Infrastructure as a Service (laaS)		Containers as a Service (CaaS)	Platform as a Service (PaaS)	Functions as a Service (FaaS)	Software as a Service (SaaS)
<ul> <li>IBM Bluemix Infrastructure</li> <li>Amazon Web Services</li> <li>Microsoft Azure</li> <li>OpenStack</li> </ul>		<ul> <li>IBM Containers on Bluemix</li> <li>Amazon Elastic Container Service</li> <li>Microsoft Azure Container Service</li> <li>Kubernetes</li> <li>Docker Swarm</li> <li>Apache Mesos</li> </ul>	<ul> <li>IBM Bluemix Cloud Foundry Runtimes</li> <li>Amazon Elastic Beanstalk</li> <li>Microsoft Azure App Service</li> <li>Cloud Foundry</li> <li>OpenShift</li> </ul>	<ul> <li>IBM Bluemix OpenWhisk</li> <li>Amazon Lambda</li> <li>Microsoft Azure Functions</li> </ul>	<ul> <li>IBM Cloudant (NoSQL database)</li> <li>Salesforce (CRN platform)</li> </ul>
MQ queue	e r	nanagers			

### MQ client applications

### **Choice of Cloud**

#### When planning for the cloud you should consider provider

- Some providers may only offer tools to use them (Vendor locking)
- What does the provider offer? Public only? Private option?

#### What kind of cloud?

- Public Managed by someone else
- Private The cloud is managed by you
- Mixture.

### **Choice of Environment**

#### Infrastructure

- Virtual machines
- Containers

### • Software as a service?

- Messagehub
- No "IBM MQ" SaaS, but being considered. (at least from IBM anyway)



### Choice of tools & which tool for which job

#### Orchestration

- How do you want to deploy your environments?
- Doing it by hand introduces error

#### Management & monitoring

- Once you have your environment, how are you going to manage it?
- You probably already have this system in place so are you going to:
  - A. Work out how to configure cloud Queue Managers to talk to it?
  - B. Create a new system to manage these environments

#### Vendor locking

- Some tools are provided by a cloud provider to use with their cloud
- But if you want to change providers in the future you will have to use another tool.

#### More on this in Matthew Whiteheads talk. (Deploying MQ to the Cloud)

Wednesday 15:50 - Sagewood



### How will you slot it into your network?

#### Hybrid cloud?

- Private cloud with On Prem
- Public cloud with on prem?
- Public cloud talking to Private?

#### What about communication between clouds/on prem?

- Secure gateway
- VPN provided by cloud?
- ► TLS
- AMS



### **Future Proofing**

#### How will you apply updates/patches?

- You need to apply a patch, how will you apply it to all of your Queue Managers?
- You want to upgrade, how do you upgrade all of your Queue Managers?
- Do you have a way to stage the upgrade/patch?
- Not just limited to patches/updates of installation, what about configuration?

#### How will you add new instances in?

- If you need to scale up how will you connect the new instance with your existing network
- How will you remove one when you don't need it anymore.

### **MQ SPECIFIC CLOUD CONSIDERATIONS**

Running in the cloud *could* be as simple as picking up your current infrastructure and moving it to the cloud

You then just manage a bunch of VMs exactly like your datacentre don't you?...

You could, but that would be missing the point

. . .

Clouds bring an opportunity for simplified operations and management, better scalability and better service and therefore application quality of service



### MQ, ready for the cloud

## MQ's capabilities were ready for the cloud before the cloud was even a *thing*

- Dynamic client connectivity
- Dynamic scaling
- Workload balancing
- High security
- High scale and robustness
- Repeatable and remote administration and monitoring for cattle, not pets deployments
- Perfect for integrating systems across any cloud
- ...

### But are you using it that way?



### **Rethink MQ**

### Don't confuse old MQ practices for MQ itself

#### "MQ is too hard to use"

"Our MQ system is too complicated to change"

"MQ isn't cloud, it's too old!"



### How many of these do you have?

- Hand crafted, shared queue managers
- Applications hard coding connection details
- Applications bound to a single IP address
- Edge security at most
- Internal architecture complexity exposed to the applications
- A lengthy change control process
- Manual installation, deployment and configuration

### Rethink how you use MQ!

### **Running MQ in the cloud**

- Persistent storage
- Security
- Scalability
  - Service discovery
  - Load balancing
- Error log management
- Metrics and monitoring
  - Centralized metrics
  - Centralized event messages
- Client applications



### Re-think how you use MQ



**Security** 



- No "armadillo" security
- All data in motion and at rest needs to be secured
  - Configure MQ channels with TLS
  - Use channel authentication for access control
  - Configure disk encryption either at the OS or cloud provider level.
  - IBM MQ Advanced Message Security is also available to provide additional security protections, such as permessage encryption.

### **Persistent storage**



Local storage



Block storage

Bluemix Block Storage Amazon EBS OpenStack Cinder Ceph RBD DRBD



File storage

NFS V4 Bluemix File Storage Amazon EFS OpenStack Manila CephFS

### **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 is moved (e.g. run a container in a different VM), then something needs to re-connect the queue manager to the correct storage.
- Some cloud orchestrators will run identical instances of your image. This could lead to lots of copies of "qm1". Other orchestrators, like Kubernetes, allow you to manage separate identities ("qm1", "qm2", etc.).

### **Scalability**

- In general, scale clients separately from servers
- Scaling up is easy
  - MQ cluster
  - Load-balanced set of identical queue managers

### Scaling down depends more on your applications

Need to remove a queue manager in a controlled manner making sure that all of the messages are safely processed.

#### Message ordering

- Scaling out rather than scaling up, brings concurrent processing
- Messages could be received out of sequence.
- MQ provides features to allow groups of messages to be handled in small ordered batches, or you can manage this yourself in your application.





### **Availability**

#### How do you want to handle Availability?

- Use existing MQ functionality (Active/Passive Queue Managers)
- Use systems built into Kubernetes (Demo in the Container talk)
- Operate across multiple Availability zones

#### Availability zones,

- Pay attention to where you operate.
- Export regulations

You are probably already spreading your MQ infrastructure across multiple regions/Data Centres

Do you use existing MQ features or Cloud provider features?

### **Error log management**



- To manage large numbers of servers, you don't want to SSH into them very often (if ever).
- You will still need to diagnose problems
- Centralized logging is commonly used, where an agent sends MQ error logs and system logs to a centralized location
  - Store
  - Index to make searchable
  - Analyze

#### • For example:

- IBM Monitoring & Analytics
- AWS Cloudwatch
- ElasticSearch



### **Metrics and monitoring**



- MQ V9 makes many statistics available through a pub/sub interface
  - Subscribe to topics under \$SYS/MQ for information on:
    - CPU usage
    - Disk usage
    - Connections and disconnections
    - Opening and closing of queues
    - Pub/sub and put/get
    - Syncpoint calls
    - Changes to MQ objects (MQSET and MQINQ)
- Cloud product insights can be used with IBM MQ.
- Other monitoring tools can be used! ELK stack, Prometheus, etc

### **Client applications**

- Cloud typically emphasizes resilience over robustness
- Reminder: use your MQ client's auto-reconnect feature
  - Cloud servers are likely to get restarted or moved

factory.setClientReconnectOptions(WMQConstants.WMQ\_CLIENT\_RECONNECT);

#### Service discovery will probably be affected

- Dynamic IP addresses
- Scalability may mean multiple equivalent queue managers are available



### **General Principle**

- Create image with IBM MQ.
  - Docker image, EC2 Ami, etc
- Create method to customize the MQ Queue Manager depending on variables
  - Docker example image runs all MQSC files loaded in /etc/mqm
- Use tools to create as many instances as you need
  - Scale up and down on requirement
- Developer works has (and will have)

### examples of different uses of MQ in different scenarios

 Including the automation scripts we used to create and customize our instances



### **MQ as a Service Redbook**

#### Redbooks

### IBM MQ as a Service A Practical Approach



### http://ibm.biz/mqaas\_red

Information about how to build, deploy, and use IBM MQ as a service.

Explains how to apply as a service methodologies to an IBM MQ environment.

### Where can I get more information?

IBM Messaging developerWorks developer.ibm.com/messaging ~

IBM Messaging Youtube https://www.youtube.com/IBMmessagingMedia

LinkedIn Ibm.biz/ibmmessaging

Twitter @IBMMessaging

30/2/2017

IBM MQ Facebook Facebook.com/IBM-MQ-8304628654/



Blog posts tagged with

"cloud"

### Would you like to take part in IBM MQ Design Research?

The IBM MQ team is currently conducting some long term research with our MQ customer base.

#### With this survey we would like to understand:

- Who is interreacting with MQ and what are their responsibilities?
- Which customers are interested in moving IBM MQ into the cloud?
- Which customers would like to take part in future research?
- We estimate the survey should take 4 minutes to complete.

Please note: This survey is for distributed users only.

If you're interested, go to <u>ibm.biz/MQ-Customer-Survey</u>



### Other Cloud sessions from the IBM MQ team.

#### MQ in Containers – Rob Parker

- Monday 9:50 Leopardwood Room
- Wednesday 14:30 Leopardwood Room

#### MQ Automation: Config Managenment using Amazon S3 – T.Rob Wyatt

- Monday 15:50 Aloeswoood Room
- Wednesday 8:30 Aloeswood Room

#### MQ Hybrid Cloud Architectures – Matt Whitehead

- Tuesday 8:30 Sagewood Room
- Wednesday 9:50 Sagewood Room

#### MQ Automation: Config Management using Baselines, Patterns and Apps – T.Rob Wyatt

- Monday 9:50 Aloeswood
- Tuesday 13:00 Aloeswood Room

#### What's up DOCker – Rob Sordillo

- Monday 11:15 Zebrawood Room
- Wednesday 11:15 Sagewood Room

### Other Cloud sessions from the IBM MQ team.

#### Introduction to Kafka (and why you care) – Richard Nikula

- Monday 14:40 Zebrawood Room
- Wednesday 14:30 Aloeswood Room

#### MQ Console & REST API – Matt Leming

Wednesday 15:50 – Rosewood Room

#### Deploying MQ to the Cloud – Matt Whitehead

- Monday 9:50 Sagewood
- Wednesday 15:50 Sagewood

#### Meet the experts! – Various

Tuesday 15:50 – Zebrawood Room

### **Questions & Answers**



### **Notices and disclaimers**

Copyright © 2017 by International Business Machines Corporation (IBM). No part of this document may be reproduced or transmitted in any form without written permission from IBM.

U.S. Government Users Restricted Rights — use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.

Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. This document is distributed "as is" without any warranty, either express or implied. In no event shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity. IBM products and services are warranted according to the terms and conditions of the agreements under which they are provided.

IBM products are manufactured from new parts or new and used parts.

In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply."

Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.

Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented

IBM products ano

the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.

Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the

views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.

It is the customer's responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer's business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer is in compliance with any law.

# Notices and disclaimers continued

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. IBM expressly disclaims all warranties, expressed or implied, including but not limited to, the implied warranties of merchantability and fitness for a particular, purpose.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.

IBM, the IBM logo, ibm.com, Aspera®, Bluemix, Blueworks Live,

CICS, Clearcase, Cognos<sup>®</sup>, DOORS<sup>®</sup>, Emptoris<sup>®</sup>, Enterprise Document Management System<sup>™</sup>, FASP<sup>®</sup>, FileNet<sup>®</sup>, **Global Business Services**<sup>®</sup>, Global Technology Services<sup>®</sup>, IBM ExperienceOne<sup>™</sup>, IBM SmartCloud<sup>®</sup>, IBM Social Business<sup>®</sup>, Information on Demand, ILOG, Maximo<sup>®</sup>, MQIntegrator<sup>®</sup>, MQSeries<sup>®</sup>, Netcool<sup>®</sup>, OMEGAMON, OpenPower, PureAnalytics<sup>™</sup>, PureApplication<sup>®</sup>, pureCluster<sup>™</sup>, PureCoverage<sup>®</sup>, PureData<sup>®</sup>, PureExperience<sup>®</sup>, PureFlex<sup>®</sup>, pureQuery<sup>®</sup>, pureScale<sup>®</sup>, PureSystems<sup>®</sup>, QRadar<sup>®</sup>, Rational<sup>®</sup>, Rhapsody<sup>®</sup>, Smarter Commerce<sup>®</sup>, SoDA, SPSS, Sterling Commerce<sup>®</sup>, StoredIQ, Tealeaf<sup>®</sup>, Tivoli<sup>®</sup> Trusteer<sup>®</sup>, Unica<sup>®</sup>, urban{code}<sup>®</sup>, Watson, WebSphere<sup>®</sup>, Worklight<sup>®</sup>, X-Force<sup>®</sup> and System z<sup>®</sup> Z/OS, are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: www.ibm.com/legal/copytrade.shtml.