



SOA Expert Series: OSB Internals with Oracle A-Team

Dec 15, 2016

David Shaffer, Managing Partner, Middleworks

Mike Muller, Cloud Solution Architect, Oracle A-Team

Ben Kothari, Solution Architect, Ampliflex

Kathryn Lustenberger, Principal Product Manager, Oracle



Webinar Housekeeping

- Use streaming audio or call into the US or International numbers provided by GoToWebinar
- All attendees are muted for the session – ask questions via Q&A interface
 - We have answered many of the questions submitted with registration. Review the current Q&A at: <http://bit.ly/2gCZnJ6>
- Ask questions throughout the session. We may answer as we go via Q&A and then discuss a few more broadly at the end
- Slides, recording of session and finalized Q&A will be made available within a couple days at www.middleworks.com/soa-expert
- After this session, please fill out the feedback survey:
 - <https://www.surveymonkey.com/r/SOAExpertFeedback>



Agenda

- Dave Shaffer: Intro and Background
- Mike Muller: OSB Internals from SOA Blackbelt training
- Q&A

Please take this unique opportunity to ask any detailed or technical questions of the engineering, partner and A-team experts we have on the webinar!

Thanks to Kiran, former OSB dev director, and Ben Kothari from Ampliflex for answering so many submitted questions



OSB Internals Q&A and Wrapup

- Please fill out the webinar feedback survey:
 - <https://www.surveymonkey.com/r/SOAExpertFeedback>
- Planning next few webinars now – submit the feedback survey to share your preferences
- All slides plus extra content available by Monday at:
 - <http://www.middleworks.com/soa-expert>
- Contact dave@middleworks.com if you want your own on-site version of this level of expertise or OSB or SOA 12c training
- Thanks, happy holidays and Happy New Year all!

ORACLE®

SB Black-belt Training

Oracle Service Bus Internals

Note: this content was created for version 11g, but has been updated (with very minor changes) for 12c. It should be current, as of Oracle Service Bus 12.2.1.





Agenda

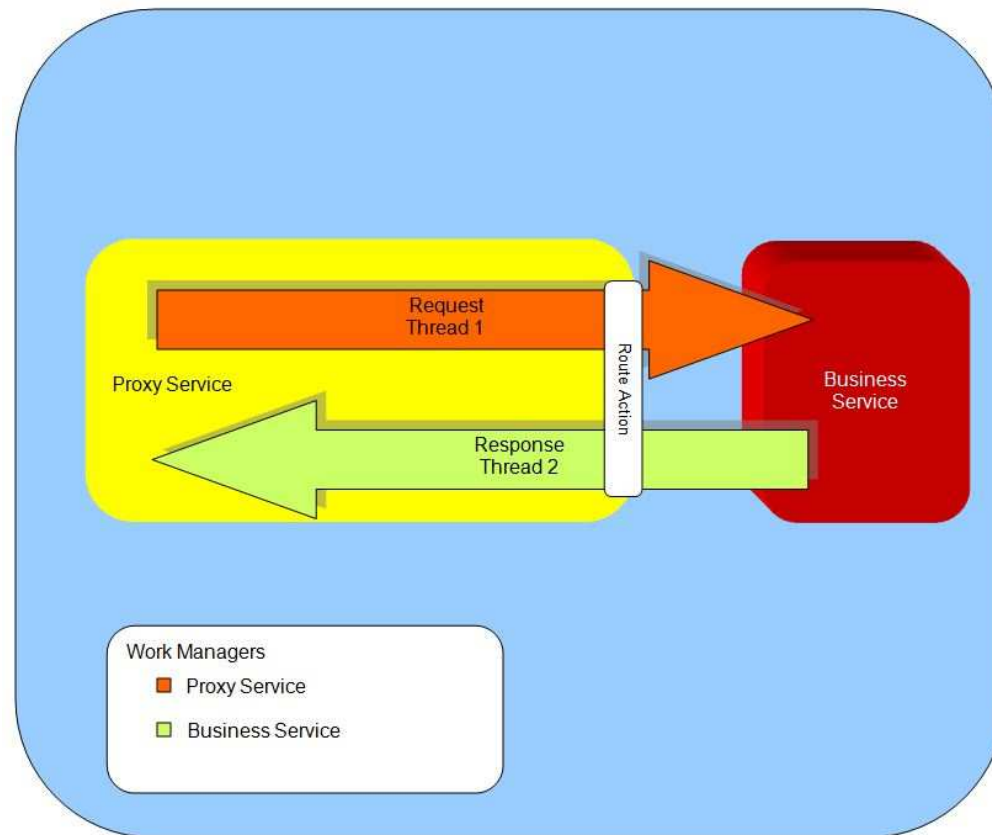
- Service Bus Threading Model
- WebLogic Thread Management
 - Muxer
 - Self-tuning Thread Pool
- WorkManagers
 - Defining
 - Applying to Services
- Throttling
- Transactions



Pipeline Threading Model

- At least two threads are required to process a proxy request
 - A minimum of 1 for the request pipeline
 - A minimum of 1 for the response pipeline
- The number of threads will vary based on the actions in the pipeline

Pipeline Threading Model





Blocking vs non-Blocking

- Whenever possible, the request thread does not block waiting for a response
- Whenever possible, calls to remote systems are done in a non-blocking fashion
- Some transports are inherently blocking
 - JEJB
 - SB
- Others are not
 - HTTP
 - JMS
 - MQ



Threading Behavior Modifiers

- Several factors can influence the threading behavior
 - Service invocation method
 - Quality of Service setting
 - Pipeline transactional settings
 - Transport implementation
 - Business Service retries



Service Invocation Methods

- Route
 - Most commonly used
 - Optimized to be the most efficient
 - Can be request only, or request-response
 - Non-blocking in most cases, depends on transport
- Publish
 - One way messaging
 - Often used to implement custom auditing/logging
 - Generally non-blocking, but can block in certain conditions
 - The service is one-way
 - There are retries configured for the service

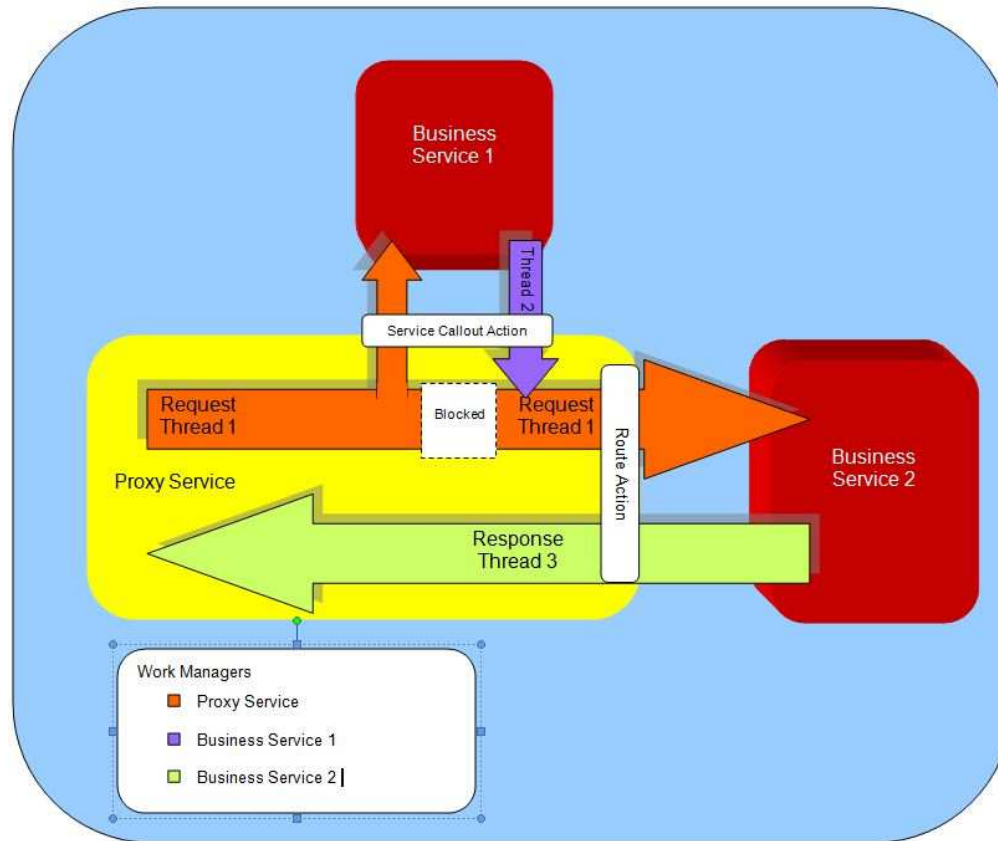


Service Invocation Methods

WSCallout (Service Callout)

- Intended for message enrichment
- **ALWAYS** a blocking call
- Outbound calls still use asynchronous request response pattern
- Requires an extra thread to receive the response and notify the blocked thread
- Potential server hangs (e.g. if all active threads are used for service callout and blocked waiting for response)

Service Callout





Business Service Retries

- The retry interval causes a delay once the set of endpoints has been exhausted (not between each retry)
- The retry interval is implemented as a `Thread.sleep(nn)`;
- The thread is not released back to the pool
- Long retry intervals can result in STUCK threads



Co-Located Service Optimization

- Proxy to Proxy invocations can be optimized – depends on the transport
- The transport layer is bypassed
- The same thread is used to execute the called Proxy service pipeline
- Supported by the following transports:

- HTTP
- JCA
- Email
- FTP/FILE
- SFTP
- DSP
- SB
- Tuxedo



Transport Specifics

Each transport can vary in its utilization of threads

HTTP – Exactly-once forces request thread to block

JMS – Number of destination consumers defaults to 16

Local – Executed in same thread as calling proxy

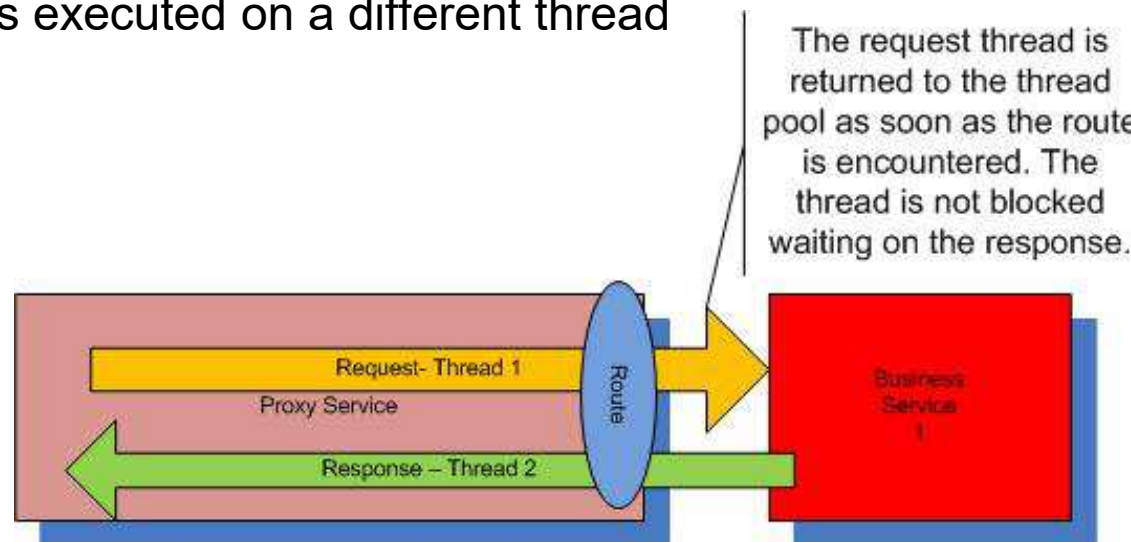
MQ – Polling and Worker threads

JEJB – Synchronous, forces request thread to block until response received

OSB Thread Usage Example

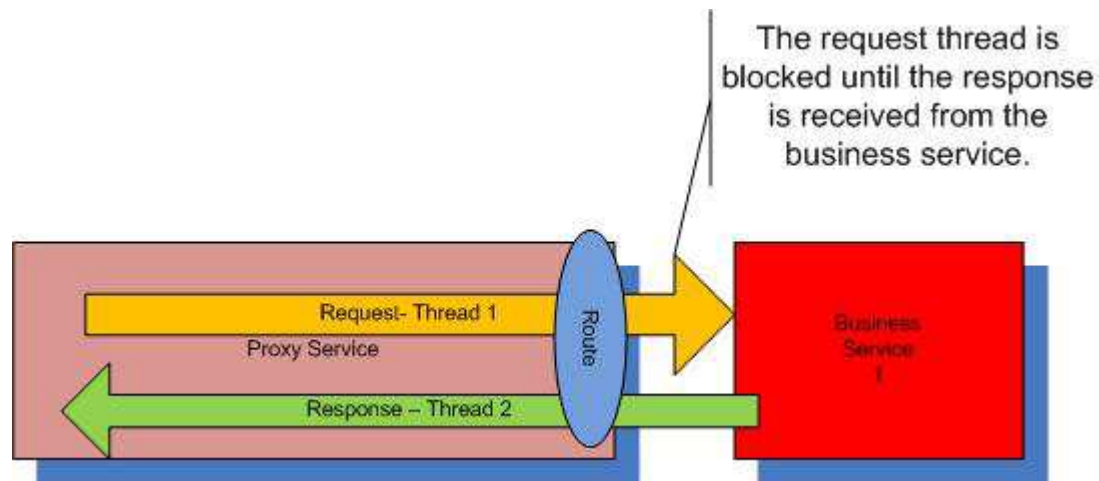
HTTP Transport default processing

- Incoming requests are dispatched to an available execute thread
- After routing to the business service, the request thread is released to the pool
- The response is executed on a different thread

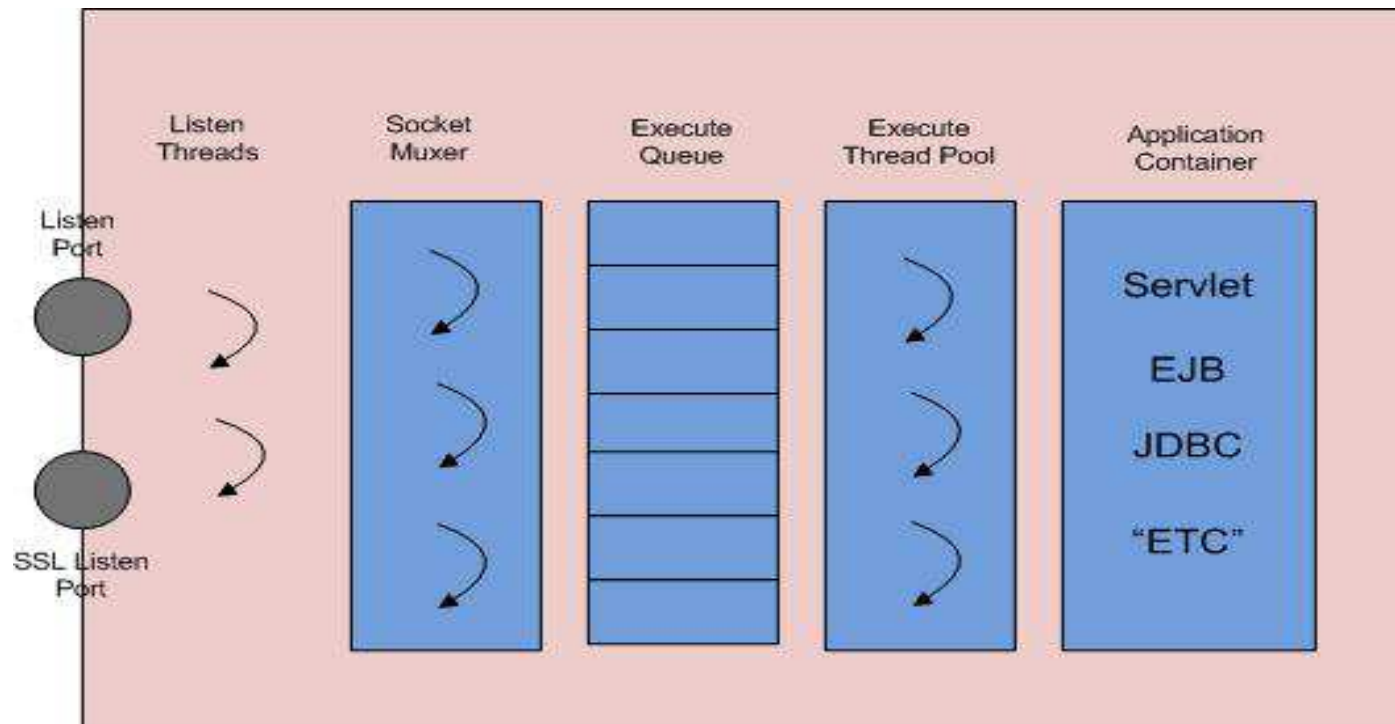


Effect of Quality of Service

- The Quality of Service specified for a service invocation can modify the threading behavior
- For HTTP, Exactly-Once forces the request thread to wait for a response
 - Allows errors to be handled by the request thread



Threads and Work Managers Server Architecture





Threads And Work Managers

Listen Ports and Listen Threads

- Waits for connection requests, accepts the request, hands the socket off to the muxer, and goes back to wait for the next request
- By default, WLS listens on two ports
- Listen Port – All non-SSL based protocols (T3, HTTP, IIOP)
- SSL Listen Port – All SSL-based protocols (T3S, HTTPS, IIOPS)
- May also configure additional network channels to allow additional ports to be defined and limit the protocols accepted



Threads And Work Managers

Socket Muxer Overview

- Processes incoming requests on established connections
- Handles requests for any supported protocol
- Inspects first few bytes of input stream to:
 - Determine the protocol
 - Determine the request target
 - Determines the work manager
- Packages the request up and puts it on the execute queue



Self Tuning Thread Pool

- Single user thread pool that automatically grows and shrinks
- 400 max for “active” state
- Can be tuned in some cases
 - specify the min thread pool size in –D...
- Don't change max

Thread State

- Active (ready to accept or executing requests)
- Idle
- Hogging
- Stuck
- Standby (not ready to accept requests but can be promoted if needed but also able to handle request in some situations)

Console-> Server: Monitoring -> Threads

Self-Tuning Thread Pool (Filtered - More Columns Exist)

Active Execute Threads	Execute Thread Total Count	Execute Thread Idle Count	Queue Length	Pending User Request Count	Completed Request Count	Hogging Thread Count	Standby Thread Count	Throughput	Health
401	406	0	8293	8293	282801	400	5	0.9995002498750625	Warning



Work Manager Fundamentals

- A WorkManager's purpose is to define a set of execution characteristics
- These characteristics provide hints to the WorkScheduler as to how requests should be serviced by the Self Tuning thread pool
- A work manager definition consists of a request class and optional constraints



Constraints

Constraints allow you to set limits on how a work manager's requests are processed.

- Maximum Threads Constraint – limit concurrent requests
- Minimum Threads Constraint – ensure threads are always available
- Request Class Constraint
 - Fair Share – relative value of thread usage time (default 50)
 - Response Time – response time goal (ratio, not absolute value)
 - Context – users or groups context of the request
- Capacity Constraint – limiting requests in queue



Work Managers in Service Bus

- Configured on service transport tab as “Dispatch Policy”
- Can be configured on both proxy and business services
- Control how many MDB (JMS Proxy) instances (max threads constraint)



Work Managers in Service Bus

- Constraints can be shared by multiple work managers
- The constraint is applied to all work managers, not per work manager
- For example
 - a Max Thread Constraint of 10 shared by 2 work managers
 - The maximum total number of threads processing for both work managers at any given time is 10, not 20.



Work Manager Best Practices

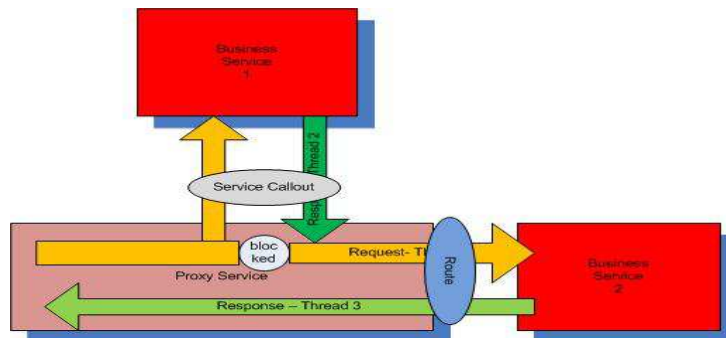
- Assign a minimum constraint work manager to services invoked by a service callout (a value of 1 or 2 is typically sufficient)
- STUCK Threads - This is one of the biggest issues seen in the field

```
"[STUCK] ExecuteThread: '166' for queue: 'weblogic.kernel.Default (self-tuning)'" daemon prio=10 tid=0x00007fec4c18a000 nid=0x4b54 in  
Object.wait() [0x00007fec1b270000]
```

```
java.lang.Thread.State: WAITING (on object monitor)  
    at java.lang.Object.wait(Native Method)  
    - waiting on <0x000000074ede2858> (a java.lang.Object)  
    at java.lang.Object.wait(Object.java:503)  
    at com.bea.wli.sb.pipeline.PipelineContextImpl$SynchronousCallback.waitForResponse(PipelineContextImpl.java:1391)  
    - locked <0x000000074ede2858> (a java.lang.Object)  
    at com.bea.wli.sb.pipeline.PipelineContextImpl.dispatchSync(PipelineContextImpl.java:460)  
    at stages.transform.runtime.WsCalloutRuntimeStep$WsCalloutDispatcher.dispatch(WsCalloutRuntimeStep.java:1380)  
    at stages.transform.runtime.WsCalloutRuntimeStep.processMessage(WsCalloutRuntimeStep.java:247)  
  
....
```

Work Manager Best Practices

- Starting in 11.1.1.7
 - Default Response Work Manage created by default
 - Automatically assigned to all Business Services
 - Doesn't avoid issues with WSCallout in Response pipeline





Work Manager Best Practices

- Advisable not to use the same work manager for both the proxy service and business service
- A work manager assigned to a business service comes into play on the response pipeline
- Work managers are per managed server
- Start simple with an objective in mind, get more complicated as required



Work Manager Monitoring

Console->-Server->Monitoring->Workload
You will only see constraints usage

Settings for AdminServer

Configuration Protocols Logging Debug **Monitoring** Control Deployments Services Security Notes


General Health Channels Performance Threads Timers **Workload** Security Default Store JMS SAF JDBC JTA

Use this page to view statistics for the Work Managers, constraints, and request classes that are configured for this server.

▶ [Customize this table](#)

Work Managers

Showing 1 to 10 of 32 [Previous](#) | [Next](#)


Name 	Pending Requests	Completed Requests
CdsAsyncRegistration	0	1
DataRetirementWorkManager	0	0
direct	0	0
ImageWorkManager	0	0
JmsAsyncQueue	0	0
JmsDispatcher	0	0
JTACoordinatorWM	0	0
OneWayJTACoordinatorWM	0	0
UserLockout	0	0
WatchManagerEvents	0	14

Showing 1 to 10 of 32 [Previous](#) | [Next](#)

▶ [Customize this table](#)

Min Threads Constraint

Showing 1 to 1 of 1 [Previous](#) | [Next](#)

Name 	Completed Requests	Pending Requests	Executing Requests	Out Of Order Execution Count	Must Run Count	Max Wait Time	Current Wait Time
MinThreadsConstraint-0	0	0	0	0	0	0	0

Showing 1 to 1 of 1 [Previous](#) | [Next](#)



Throttling - Overview

- Provides a means of protecting back end systems from overload
- Allows the user to define the maximum concurrent calls
- Once Maximum Concurrency has been reached, additional requests must wait
- Applies only to Business Services
- Throttling Group added in 12c

Throttling - Configuration

- Can be enabled/disabled per business service
- Maximum concurrency is **per cluster, not per managed server**
- Throttling Queue – How many messages can wait
- Additional messages are rejected immediately
- Message Expiration – How long will a given message wait in queue

Throttling	
Throttling State	<input type="checkbox"/> Enabled
Maximum Concurrency	<input type="text" value="0"/>
Throttling Queue	<input type="text" value="0"/> messages
Message Expiration	<input type="text" value="0"/> msecs



Throttling – Blocking vs Non-Blocking

- When a request is throttled it may or may not block a thread while waiting
- Throttling will inherit the blocking characteristics from the transport and invocation method
- If the call would have blocked, the throttling will block
- One exception, JEJB proxy services.



Transactions

Know the transaction demarcation

<http://atheek.wordpress.com/2011/04/21/transaction-handling-within-osb/>

- Verify that the quality of service is correctly set and understand what impact it will have on the service from a global and local transaction perspective
- When using JMS protocol and referencing a remote connection factory make sure to enable “Is XA Required” in the advanced settings.

This will ensure that the generated MDB will have the transaction attribute set properly (Required)

- HTTP transport does not support XA and cannot be part of the global transaction



Transactions

Pipeline Options

- Use XA Transaction – Creates a new XA transaction before processing the request pipeline
 - Only if one doesn't already exist
 - Intended for non-transactional transports (HTTP)
- Same Tx For Response – Process the response pipeline within the same XA transaction as the request pipeline
 - Not possible for Request-Response XA scenarios (JMS)
 - Request message won't be sent until transaction commits



Transactions and Quality of Service

- An inbound transactional transport, such as JMS, proxy service will change the default QoS to Exactly-Once
- A JMS destination not participating in the global transaction may have unintended consequences of duplicate messages being delivered if there is a rollback and retries
- As an example, setting the QoS to Best-Effort in the routing options will force the business service to commit in a local transaction and therefore will not be part of the global transaction



Transactions and Quality of Service

Message Loss

Scenario

- Configured a non-XA connection factory
- Message was posted to Queue and was consumed by an inbound JMS proxy service
- Error is encountered during the processing of the request pipeline

Result

- The message will be “lost” unless an exception handler is configured to repost the message



Transactions and Quality of Service

Message Loss - Corrected

Scenario

- Configured a XA connection factory for Proxy Service and Business Service
- Message was posted to Queue and then consumed by an inbound JMS proxy service
- Error is encountered processing the request pipeline

Result

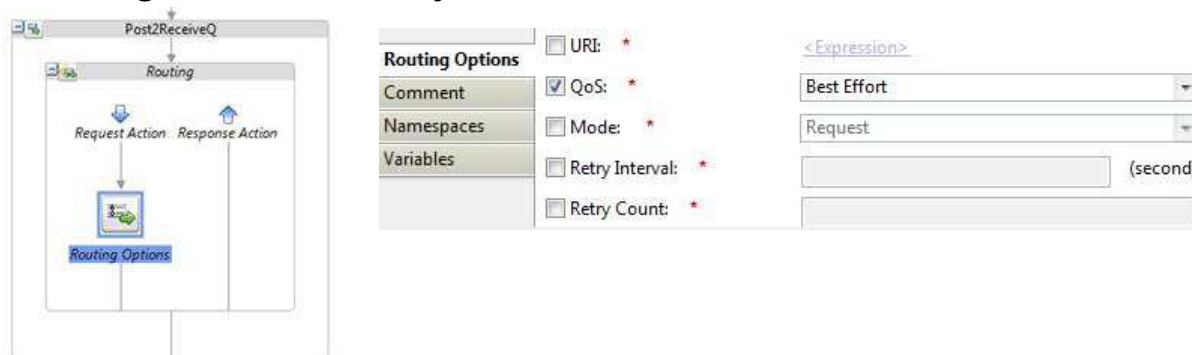
- The message will be rolled back to the original queue. No messages are lost.

Transactions and Quality of Service

Errors Result In Duplicate Messages

Scenario

- Configured an XA connection factory for Proxy Service and Business Service
- Business Service invoked via Post action in request pipeline
- Changed the Quality of Service of the business service to Best-Effort





Transactions and Quality of Service

Errors Result in Duplicate Messages

- Message was posted to Queue and then consumed by an inbound JMS proxy service
- Error was encountered on the request pipeline of the message flow

Result

- No messages will be lost, but there will be “number of retries + 1” messages sitting on the business service queue
- The Best-Effort forced the business service to a local transaction; therefore, its commit was outside of the global transaction

ORACLE®