

First Steps with IBM Integration Bus: Application Integration for a new world

David Coles

*IBM Integration Bus Level 3 Technical Lead,
IBM Hursley – dcoles@uk.ibm.com*

2nd March 2015

17041



#SHAREorg



**SHARE is an independent volunteer-run information technology association
that provides education, professional networking and industry influence.**



Important Disclaimer

- THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY.
- WHILST EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.
- IN ADDITION, THIS INFORMATION IS BASED ON IBM’S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE.
- IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION.
- NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, OR SHALL HAVE THE EFFECT OF:
 - CREATING ANY WARRANTY OR REPRESENTATION FROM IBM (OR ITS AFFILIATES OR ITS OR THEIR SUPPLIERS AND/OR LICENSORS); OR
 - ALTERING THE TERMS AND CONDITIONS OF THE APPLICABLE LICENSE AGREEMENT GOVERNING THE USE OF IBM SOFTWARE.

Agenda

- Introduction
- Inside IBM Integration Bus
- Development
- Administration
- Developer Edition
- Industry Solutions
- Key Usage Scenarios
- Questions?

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

Introduction

Complete your session evaluations online at www.SHARE.org/Seattle-Eval


What do we mean by Integration?



- Enterprise systems consist of many logical **endpoints**
 - Off-the-shelf applications, services, web apps, devices, appliances, custom built software...
- Endpoints expose a set of inputs and outputs, which comprise
 - Protocols - e.g. MQ, TCP/IP, HTTP, File system, FTP, SMTP, POP3 etc.
 - Message Formats - e.g. Binary (C/COBOL), XML, Industry (SWIFT, EDI, HL7), User-defined
- Integration is about connecting these endpoints together in meaningful ways
 - Route, Transform, Enrich, Filter, Monitor, Distribute, Decompose, Correlate, Fire and Forget, Request/Reply, Publish/Subscribe, Aggregation, Fan-in, Complex Event Processing...

- Three strands are involved in connecting applications together.
- Applications need to talk with each other over a communications protocol. Typical protocols in use today include TCP/IP, and higher level protocols such as FTP, SMTP and HTTP.
- Over the communications protocol applications exchange data, typically in discrete structures known as messages. The format of these messages can be defined from C structures or COBOL copybooks (for example), or simply use a standard format such as XML.
- In order to connect applications together so that their protocols and message formats interoperate, mediation patterns need to be applied to one or both systems you're trying to connect. These mediation patterns can be relatively straightforward, e.g. routing messages from one place to another, or the transformation of one message format into another... to relatively complex patterns such as aggregating multiple outputs from an application into a single message for a target system.

Integration solutions are about reducing cost!

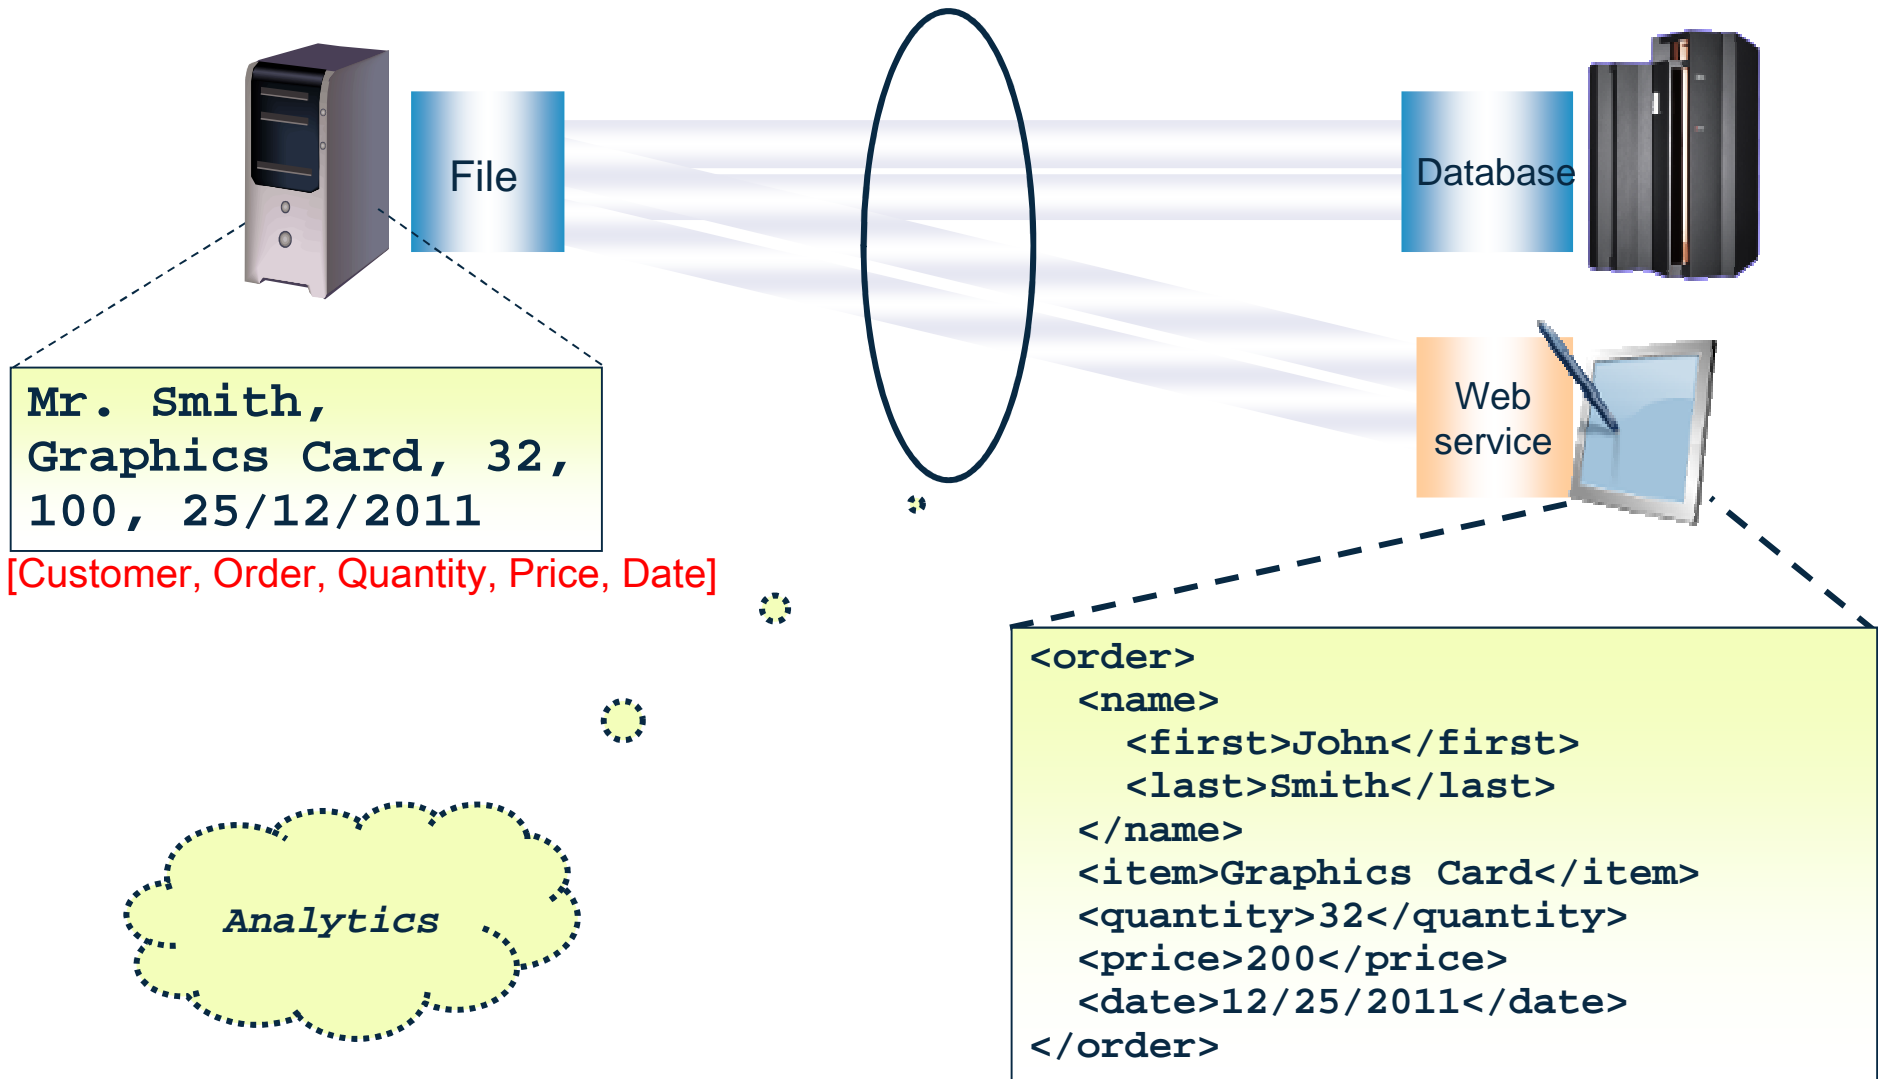


- Integration solutions simplify integration!
 - Avoids rewrites in response to new integration requirements
 - Simplifies maintenance by reducing expensive coupling
 - Flexibility adding anonymity between producers and consumers of data
 - Adds insight into applications and business value they bring

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



Example integration



Complete your session evaluations online at www.SHARE.org/Seattle-Eval

[Customer, Order, Quantity, Price, Date]


- An Application Integration Scenario.
- Application A sends some data to application B. At design time, the two applications agreed on the format of the data as the ordered set {Customer, Order, Quantity, Price, Date}. Further, the date is in UK format, the price in UK pounds sterling, and all fields are represented by character strings in codepage 500. Finally, the data is delimited using commas.
- Some time later, Application C is introduced. It needs the same data, but because it is a packaged application from a vendor or may be an application that already existed, it expects data to arrive in a different format. The date is in US format, the price is in dollars and the data is in XML.
- So, we now have an integration choice to make. Either application C must be enhanced to support the data format between A and B, or application A must be enhanced to support application C's data format. (This is an interesting use of the word “enhanced”, but you'll probably want to use it to justify the expenditure!)
- By introducing a solution that can mediate between these applications, you can integrate them without spending time and money modifying and retesting the existing applications. IBM Integration Bus is one such solution.
- In addition, a solution like IBM Integration Bus will also allow you to intercept or record the data as it is processed. Allowing you to satisfy audit requirements or for further analysis for use cases like fraud prevention.

Some examples of integration topologies



Bridges

- Often used for single point-to-point connections
- Usually cheap and quick to configure
- More difficult to scale to larger numbers of endpoints




Enterprise Service Bus (ESB)

- Logical construct that combines messaging and enrichment
- Scales very well; can integrate small and large numbers of endpoints, and can be easily distributed
- Often applied as a backbone for a Service Oriented Architecture (SOA)
- Solutions can usually also be applied to hub and spoke style architectures

Gateways

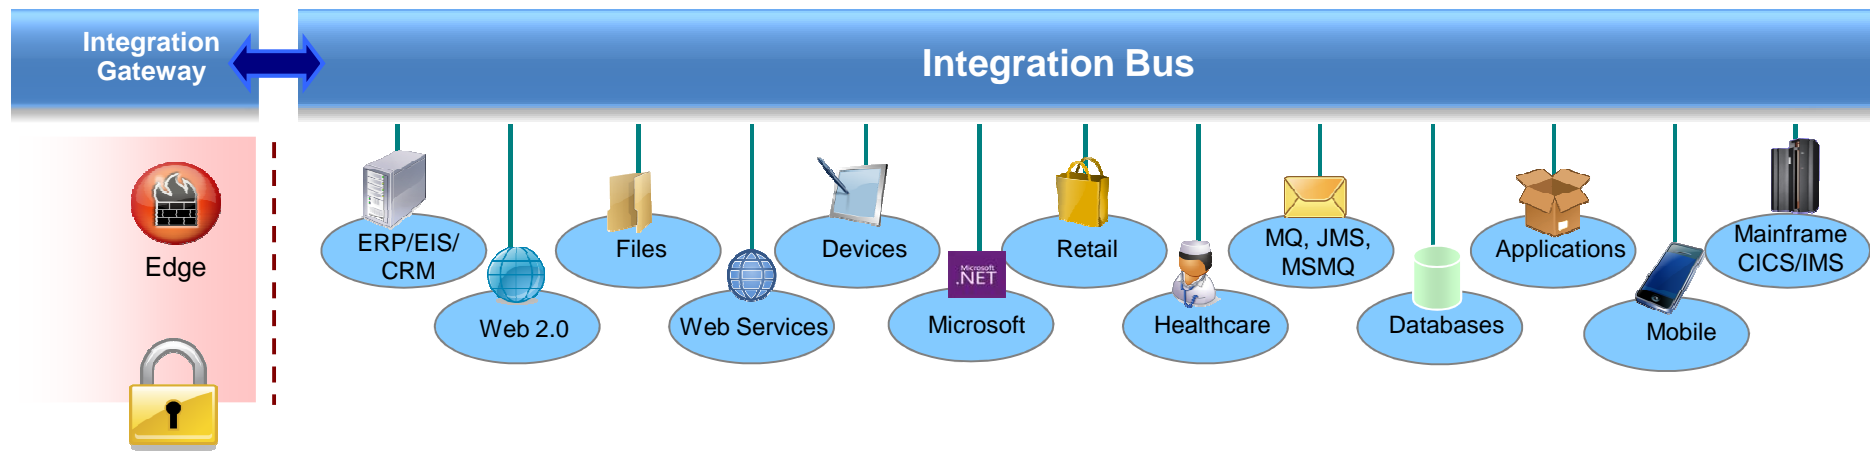
- Provides connectivity to third parties or to a specific class of endpoint
- For example, internet, cloud, security, DMZ, B2B
- Combines simplicity of configuration and a commonly on-ramp to back-end ESB



Introducing IBM Integration Bus



- **IBM's Strategic Integration Technology**
 - Single engineered product for .NET, Java and fully heterogeneous integration scenarios
 - DataPower continues to evolve for integration gateway use-cases



- **IBM Integration Bus is the new name for WebSphere Message Broker**
 - Technology progression over 15 years, installed at 2500+ customers worldwide across all industries
 - Fully supported worldwide by IBM global support network, standard 5 + 3 years support policy
 - Version to version migration is key design consideration
 - Global skills availability - SME's available globally via IBM and partners
 - Close interaction with growing and loyal customer base: beta and lab advocacy programs
 - Also incorporates WebSphere ESB use-cases

IBM Integration Bus



- Provides endpoints and the ability to connect to other endpoints
 - Off-the-shelf applications, services, web apps, devices, appliances, custom built software...
- Protocols and Message Formats
 - Protocols - e.g. MQ, TCP/IP, HTTP, File system, FTP, SMTP, POP3 etc.
 - Message Formats - e.g. Binary (C/COBOL), XML, Industry (SWIFT, EDI, HL7), User-defined
- Mediation Patterns
 - Route, Transform, Enrich, Filter, Monitor, Distribute, Decompose, Correlate, Fire and Forget, Request/Reply, Publish/Subscribe, Aggregation, Fan-in, Complex Event Processing...

- We can now revisit this earlier slide in the context of IBM Integration Bus.
- IBM Integration Bus enables “universal connectivity” by integrating protocols, message formats and mediation patterns.
- IIB provides the ability to be an endpoint and to connect to other endpoints.
- It can do this over a variety of protocols and using a variety of message formats, sometimes with more than one in use at each time.
- IIB also supports a wide range of mediation patterns, helping to support the use of the various message formats and protocols in many ways.
- As we go through the rest of this presentation we will see how IBM Integration Bus supports all of these.

A Broad Range of Supported Platforms and Environments



- Broad range of operating system and hardware platforms supported
 - AIX, Windows, z/OS, HP-UX, Linux on xSeries, pSeries, zSeries, Solaris (x86-64 & SPARC), Ubuntu
 - Optimized 64-bit support on all platforms; 32-bit option available for Windows and x/Linux
 - Support for Windows 8 and Windows Server 2012; .NET CLR V4.5 included on Windows
 - Express, Standard and Advanced editions make IIB applicable for all solutions and budgets
- Virtual images for efficient utilization & simple provisioning
 - Extensive support for virtualized environments, e.g. VMWare, AIX Hypervisor... any!
 - Support for public and private clouds: Softlayer, Pure, non-IBM, RYO etc.
 - Chef scripts for automated building of flexible IIB images (see Github)
 - Pre-built images (Hypervisor editions) available on xLinux and AIX
- Includes access to full range of industry standard databases and ERP systems
 - DB2, Oracle, Sybase, SQL Server, Informix, solidDB
 - Open Driver Manager support enables new ODBC databases to be accessed
 - JDBC Type 4 for popular databases
 - SAP, Siebel, Peoplesoft, JDEdwards
- Technology components and pre-requisites (@v9)
 - Java 7 on all platforms
 - MQ 7.1 prerequisite



Traditional OS



IBM Workload Deployer



IBM Pure



Private Cloud



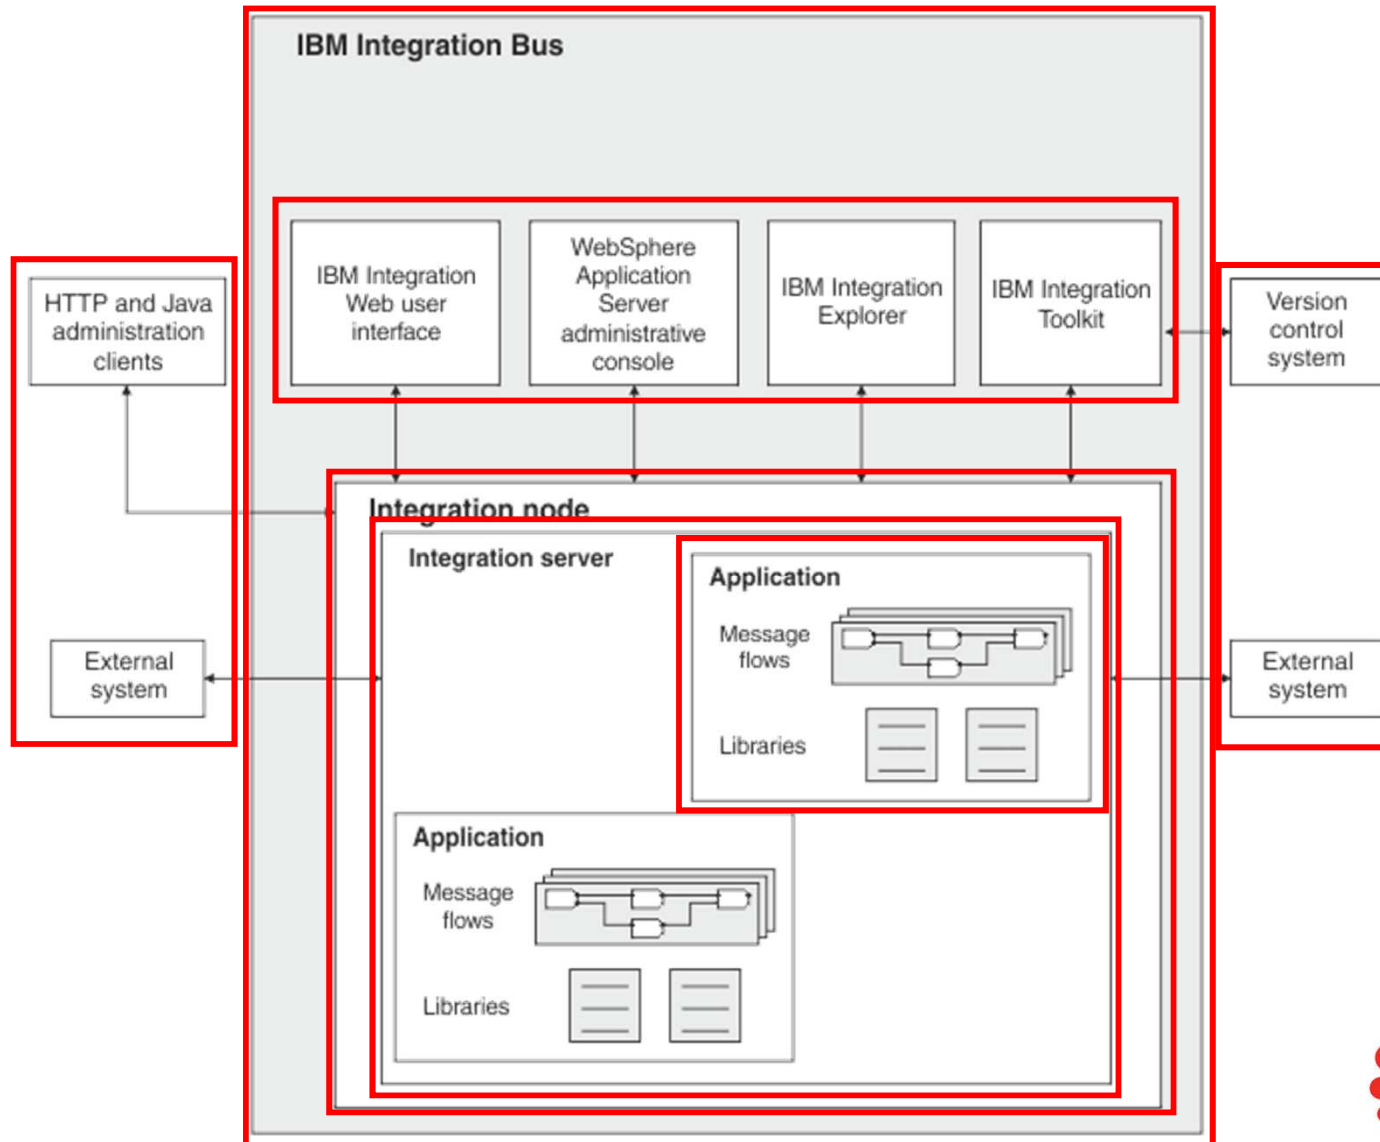


Inside IBM Integration Bus

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

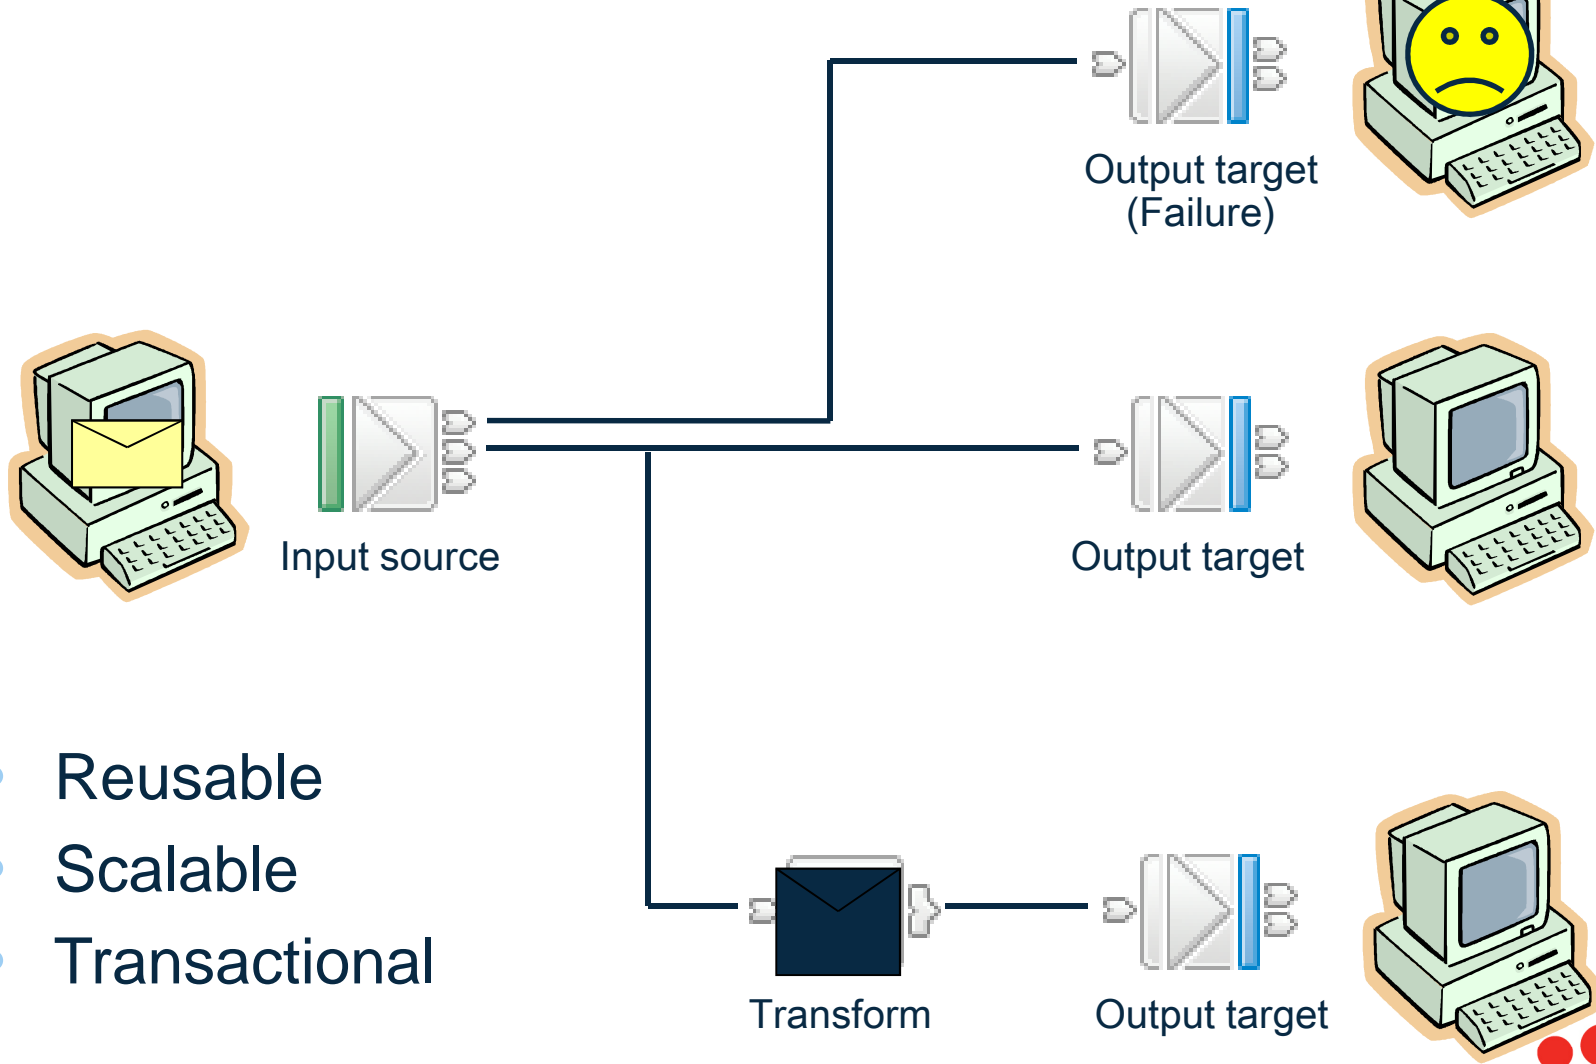


Integration Bus Components



- We'll start off by taking a look at the architectural components of IBM Integration Bus. We'll then see how these elements are used in more detail.
- The Integration Toolkit is the development environment. Based on the Eclipse platform, all the objects required to perform application integration using IIB are developed, deployed and tested here. It provides standard ways to build integration applications, perform version control and provide for the development of custom plug-ins, such as resource editors to allow users to create project resources easily. Examples are custom editors to aid flow creation, ESQL editing and syntax checking, message set modelling, and a raft of other activities. It includes a unit test broker environment.
- The integration node (or broker) is the container that hosts integration servers (or execution groups). Each server is an operating system process that contains a pool of threads responsible for running the integration logic that is deployed to it. The integration servers directly interact with the endpoints that are being integrated.
- There is also a web user interface that provides administration capability, including monitoring of deployed objects and the ability to start, stop, delete, deploy, manage workloads etc. The APIs that the web UI uses can be used by custom administration applications, either through Java, REST or commandline scripts.

Integration Bus Message Flows



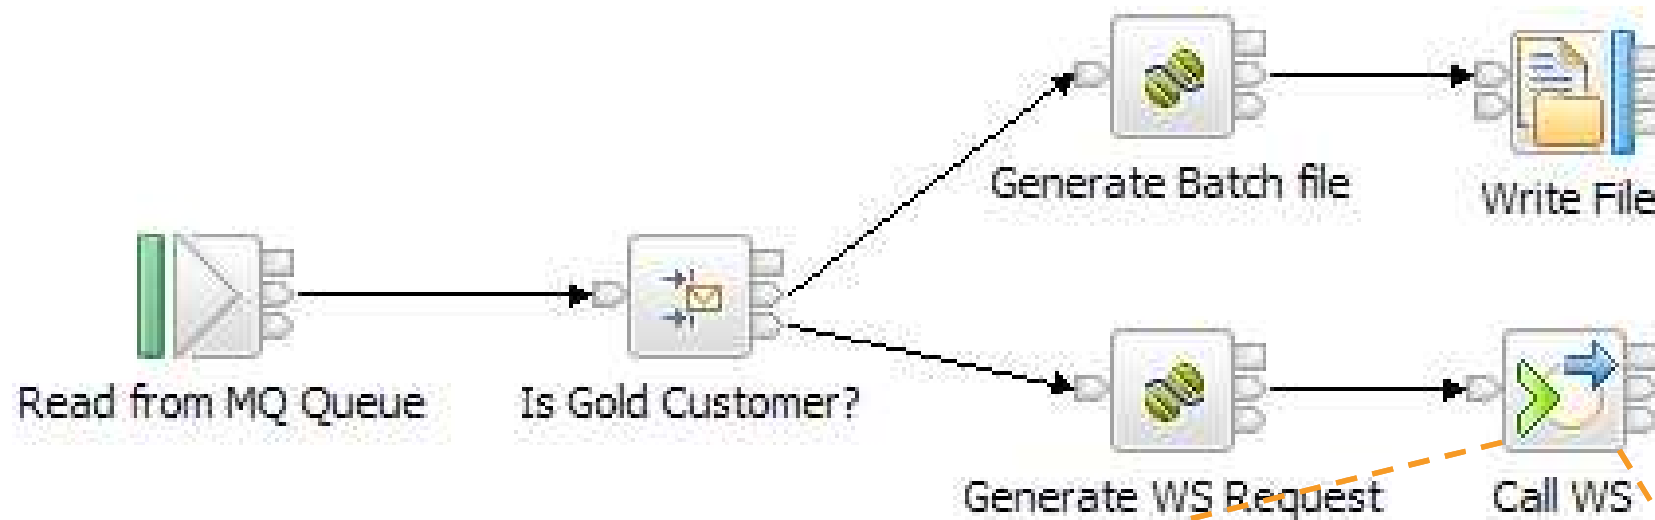
- Reusable
- Scalable
- Transactional

- Message flows provide the processing sequence required to connect applications together.
- A message flow contains the set of operations required to take a message from an originating application and deliver copies of it, some possibly transformed, to any number of connected applications for processing.
- As a message passes through a message flow, it is transformed and routed according to the nodes it encounters, and the processing decisions made within those nodes. Later we'll see how the nodes can modify the values within, or transform the structure of, a message to provide the data transformations necessary to drive backend server applications.
- For a given application scenario, the message flow describes all possible outcomes when processing a message. For example, if the message has a high monetary value, a copy of it might have to be routed to an audit application. Or if the message is not well-formed (maybe it's not encrypted in the right format), it might be routed to a security application to raise an alert.
- Equally important is the visualization of the application integration within the organization. Very often, for any particular application scenario, the application connectivity requirements (*business*) is held within the heads of domain experts. Being able to view the integration structure brings benefits in scenario understanding, reuse potential, and application architecture/standards conformance.
- After a message has been processed by a message flow, the flow does not maintain any state. It is possible to maintain such state in an external database, or within the message by using an extensible header such as the MQRFH2 or message properties.

- Message flows are general purpose, reusable integration applications.
- If you were designing a general purpose integration application, linking client and server applications, the logic would comprise separate routines each performing a well-defined function. The input routine would wait for a message, and after receiving it and checking the its integrity, (well formed etc.), it would transfer to the next routines to continue processing.
- After performing their processing, (e.g. enriching/reformatting/routing), these routines would pass control on through to the lowest functional levels, where output processing would occur. Here, messages would be written to devices, subsequently read by connected applications. At any level of processing an exception could be raised for subsequent processing.
- After the last output routine had completed, control would return back up through the levels to the input routine. Once here, all the changes would be committed and the input routine would wait for more input.

- Message flows are transactional.
 - Message flows provide vital processing and data manipulation and are therefore fully transactional. A message flow either completes all or none of its processing successfully.
 - However, if required, individual nodes can elect to perform operations outside of the message flow transaction. (e.g. audit)
- Message flows are multithreaded.
 - A given message passing through a series of nodes will execute on a single thread. To allow increased message throughput, message flows can be defined with many additional threads assigned to them. Peak workloads use additional threads, which are pooled during inactivity. We'll see more implementation details later. This means application scaling can be an operational rather than design time decision.
- Message flow nesting and chaining allow construction of enhanced capabilities.
 - Sophisticated flows can be rapidly constructed by linking individual flows together as well as nesting flows within each other.
- References:
 - Message Flow overview at http://www-01.ibm.com/support/knowledgecenter/SSMKHH_9.0.0/com.ibm.etools.mft.doc/ac00310_.htm?lang=en

Message flow example



HTTP Request Node Properties - HTTP Request

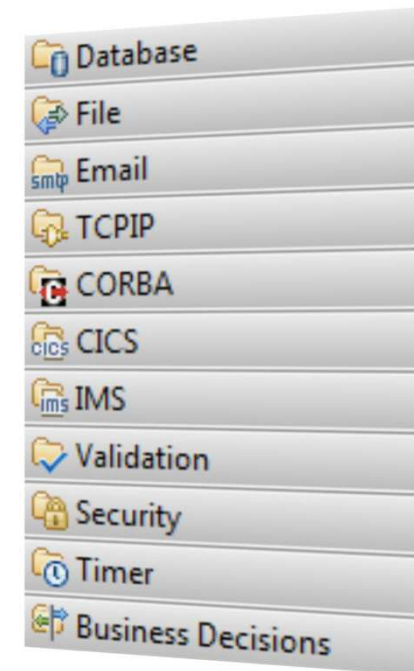
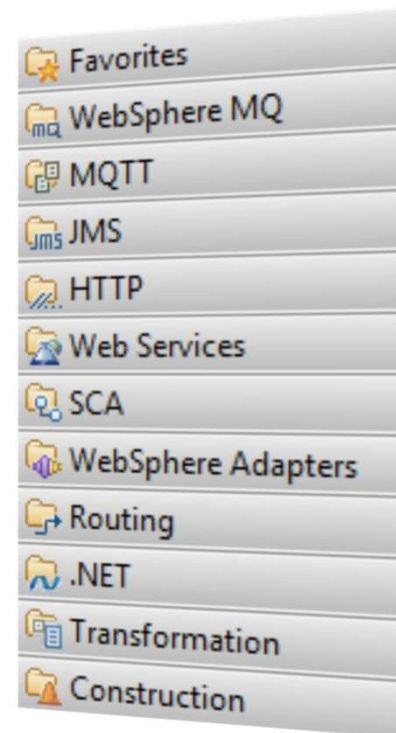
Description: Settings for working with the HTTPRequest node. [More...](#)

Basic	Web service URL*	<input type="text" value="http://example.org/MyWebService"/>
HTTP Settings		<i>e.g. http://server/path/to/service</i>
SSL	Request timeout (sec)*	<input type="text" value="120"/>
Response Message Parsing		
Parser Options		

- Here is an example of a message flow.
 - The ‘Read from MQ Queue’ node tells IIB to take messages from an MQ queue (the name of which is embedded as a property of the node, or overridden by an administrator at deployment time).
 - The message is passed onto the ‘Is Gold Customer?’ node, where a routing decision is made based on a field described in the incoming message, again which is a property on the node itself. We’ll see exactly how this condition is specified later on.
 - If the described condition holds, the message is routed to the ‘Generate WS Request’ node where the message is transformed – presumably into an SOAP message that is recognisable by the web service which is invoked by the subsequent ‘Call WS’ node.
 - If the described condition does not hold, the message is routed to the ‘Generate batch file’ node, which formats the message for subsequent output to a file in the ‘Write file’ node.
- This flow may not tell the complete integration story between the calling application and the target Web Service/File applications. For example, there is no communication back to the calling application to say that the message has been processed (or even received). Nor is there any logic in the message flow to cope with failures – for example, if the web service is not available. This is logic that could be incorporated into the flow, but not visualised here for clarity.

Nodes

- The building blocks of message flows
- Each node type performs a different (input, output or processing) action
- Many different node types
 - Grouped into logical categories in the editor
 - Nearly 100 nodes available out-of-the-box



- Nodes can be grouped in several ways; for example, by where in the flow they are used:
 - Input nodes do not have input terminals; processing of the message flow starts when a message is retrieved from an input device, for example WebSphere MQ.
 - Output nodes do not have output terminals (or at least, they are not wired to any other node). The final stage of output processing is after a message is put using one or more output nodes, and processing control returns to the input node which commits or backs out the transaction. Recalling that a message flow is analogous to a functional decomposition, it makes sense that the top most level (i.e. the input node) controls the overall transaction.
 - Processing nodes are nodes that are neither input nor output nodes. They will be connected to nodes both upstream (i.e. towards the input nodes) and downstream (i.e. towards the output nodes).
- They can also be grouped by the function that they perform.
 - Protocol-specific nodes give the broker the ability to interact with particular systems, such as MQ and Web Services.
 - Transformation nodes will take a message in one format on the input terminal and output a converted message on the output terminal.
 - Logical constructs give the message flow designer the vocabulary required to solve complex integration scenarios, for example, the ability to aggregate messages from multiple places or the ability to filter messages based on their content.
- References
 - More on nodes can be found here:
http://www-01.ibm.com/support/knowledgecenter/SSMKHH_9.0.0/com.ibm.etools.mft.doc/ac04550_.htm?lang=en

Message Flow Nodes



WebSphere MQ	Web Services	WebSphere Adapters	.NET	File	Business Decisions
MQInput	SOAPInput	PeopleSoftInput	.NETInput	FileInput	DecisionService
MQOutput	SOAPReply	PeopleSoftRequest	Transformation	FileOutput	CICS
MQReply	SOAPRequest	SAPInput	.NETCompute	FileRead	CICSRequest
MQGet	SOAPAsyncRequest	SAPRequest	Mapping	FTEInput	IMS
MQHeader	SOAPAsyncResponse	SAPReply	XSLTransform	FTEOutput	IMSRequest
JMS	SOAPEnvelope	SiebelInput	Compute	CDInput	Validation
JMSInput	SOAPExtract	SiebelRequest	JavaCompute	CDOOutput	Validate
JMSOutput	RegistryLookup	JDEJDEdwardsInput	PHPCompute	Email	Check (Deprecated)
JMSReply	EndpointLookup	JDEJDEdwardsRequest	Construction	EmailInput	Security
JMSReceive	SCA	TwineBallInput	Input	EmailOutput	SecurityPEP
JMSHeader	SCAInput	TwineBallRequest	Output	TCPIP	Timer
JMSMQTransform	SCAReply	Routing	Throw	TCPIPClientInput	TimeoutControl
MQJMSTransform	SCARequest	Filter	Trace	TCPIPClientOutput	TimeoutNotification
HTTP	SCAAsyncRequest	Label	TryCatch	TCPIPClientReceive	
HTTPInput	SCAAsyncResponse	Publication	FlowOrder	TCPIPServerInput	
HTTPReply		RouteToLabel	Passthrough	TCPIPServerOutput	
HTTPRequest		Route	ResetContentDescriptor	TCPIPServerReceive	
HTTPHeader		AggregateControl	Database	CORBA	
HTTPAsyncRequest		AggregateReply	DatabaseInput	CORBARequest	
HTTPAsyncResponse		AggregateRequest	Database		
		Collector	DatabaseRetrieve		
		Resequence	DatabaseRoute		

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



- Here's a list of the protocol specific nodes built in to IBM Integration Bus V9. For example:
 - The WebSphere MQ nodes allows Integration Broker to interact with queues on MQ Queue Managers. For example, MQInput is an input node that triggers a flow when a message arrives on a queue; MQOutput puts a message to a queue.
 - The WebSphere Adapters nodes provides native support in Integration Bus for inbound and outbound communication with Enterprise Information Systems.
 - Web Services nodes provide a rich environment for running as a Web Services requestor, provider and intermediary. Support for WS-Security, WS-Addressing, import and export of WSDL and validation against the WS-I Basic profile. The RegistryLookup and EndpointLookup nodes provide support for WebSphere Registry and Repository (WSRR).
 - The File nodes are very sophisticated and include support for the FTP and SFTP protocols, as well as advanced processing scenarios such as record detection. These nodes are complemented by additional nodes which provide support for managed file transfer systems (IBM's MQ File Transfer Edition and Sterling Connect:Direct).
 - HTTP nodes complement the Web Services capability. Support is provided for HTTP 1.0, 1.1 and HTTPS.
 - JMS nodes work with *any* JMS 1.1 compliant provider.
 - The EmailOutput node is a highly configurable node that allows e-mail messages to be sent over the SMTP protocol. EmailInput allows e-mails to be received from POP3 or IMAP servers.
 - TCP/IP nodes allow the Integration Bus to communicate with any client or server talking the ubiquitous TCP/IP protocol.
 - CORBA, IMS and CICS request nodes for integrating with CORBA, IMS and CICS applications respectively.

- Database nodes allows message flows to interact with many different data sources, including DB2, Oracle and Sybase.
- Timer nodes provide support for triggering message flows and certain times or intervals.
- The Routing category allows messages to easily flow around a network, and also allow multiple messages to be aggregated or propagated in the correct sequence.
- The Transformation category provides Integration Bus with the capability to transform messages from one format into another. Six ways of doing this are available out-of-the-box. More on these later.
- Construction nodes
 - Nodes have error handling as part of their design. If an error is detected within a primitive node (e.g. database error), the message is transferred to the failure output terminal. If the failure terminal is not connected, an exception is generated and propagated back towards the input node. There is also a specialized Throw node which allows a flow designer to generate an exception in a controlled way. Nodes can have transaction scope inside or outside of the flow.
 - A TryCatch node is used to process any such exceptions. Its 'try' terminal is used for normal processing, but if an exception occurs along this path, the TryCatch node regains control and the original message is propagated through the 'catch' terminal.
 - If the message reaches the input node, it is subject to "back out" processing. In this case, it will be propagated down its catch or failure terminal, returned to the input queue, put to a back out or dead letter queue, or discarded, as appropriate.

IBM and third-party extensions



OT4i: Open Technologies for Integration

[View on GitHub](#)

OT4i

Open Technologies for Integration is a collection of software components that can be used for integrating different technologies.

Systems, devices and applications need access to each other's services, events and documents. Open Technologies for Integration provides tools for connectivity and transformation of data as well as operational monitoring and control of the interactions.

- Patterns**
Re-usable and customizable designs of integration solutions that embody best practices.
- Connectors**
Discover and connect to various IT resources in a heterogeneous landscape.
- Data models**
Data format definitions for standards based or custom data.
- DevOps**
Automated provisioning and configuration on virtualized or physical systems.
- Conversion Tools**
Assisted conversion of code and data between different integration platforms.

Display a menu

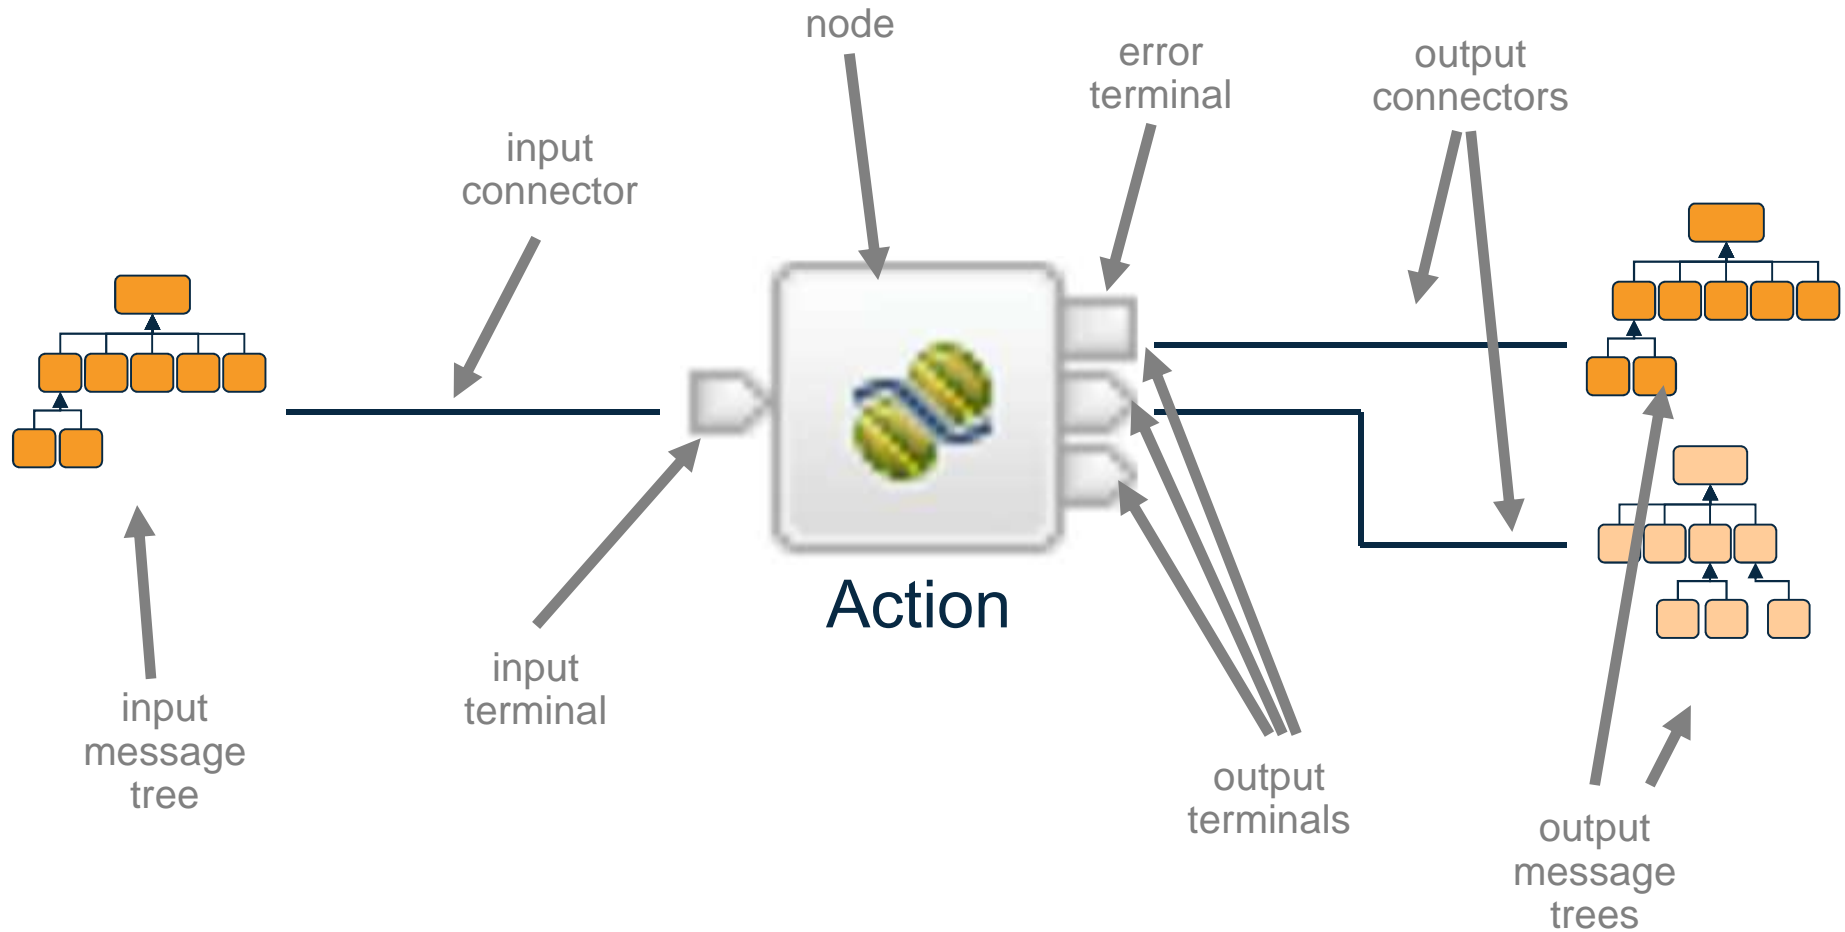
V3.0.0.1
Aug 2014

V1.0.0.0
Jun 2014

V1.0.0.0
Dec 2013

- Many other nodes and features available through product extensions
 - WebSphere TX, Tibco RV, VSAM, QSAM
- Write your own User-Defined Nodes
 - Native node framework available in C and Java
 - OT4i connector framework provides means to implement full lifecycle, including endpoint discovery

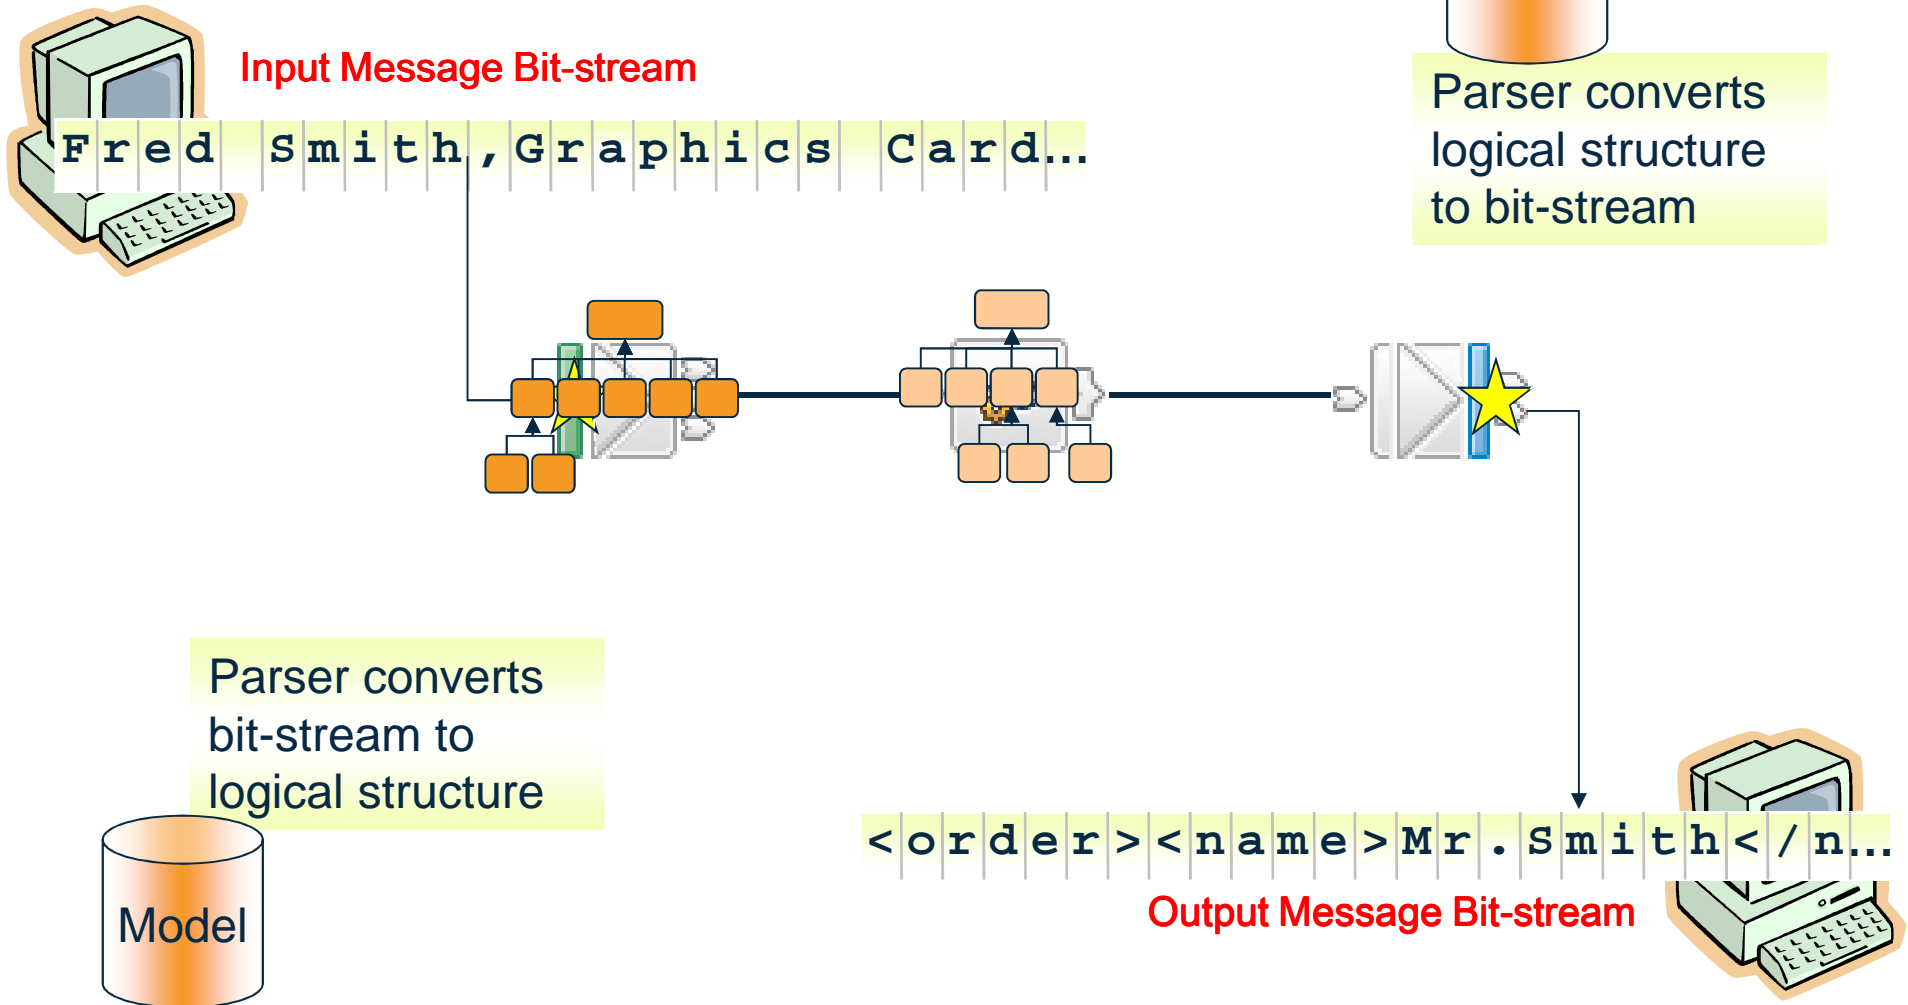
Node Terminology




- Message flow nodes provide the individual processing elements that make up a message flow.
- We've seen that a message flow is the combination of operations required to achieve application integration. We build a message flow from small units called nodes; these nodes represent the base elements required to connect messaging applications together.
- Looking at a message flow, you can see several objects identifiable with this processing.
 - Nodes represent functional routines encapsulating integration logic
 - Terminals represent the various outcomes possible from node processing
 - Connectors join the various nodes through their terminals
- A message processing node defines a single logical operation on a message.
- A message processing node is a stand alone procedure that receives a message, performs a specific action against it, and outputs zero or more messages as a result of the action it has taken.
- The action represented by a message processing node encapsulates a useful and reusable piece of integration logic. Nodes can be thought of as reusable components in an integration library.

- A node is joined to its neighbours in the data flow through connectors attached to its data terminals.
- Every node has a fixed number of connection points known as "input" terminals and "output" terminals. These allow it to be connected to its neighbours. Each node normally has one input terminal upon which it receives messages, and multiple output terminals for different processing results within the node. Different types of node have different numbers of terminals.
- A connector joins an output terminal of one node to an input terminal of the next node in the message flow. You can leave an output terminal unconnected, or you can connect a single output terminal to more than one target node.
- After a node has finished processing a message, the connectors defined from the node's output terminals determine which nodes process the message next. If a node has more than one output terminal connected to a target node, it is the node (not you) that determines the order in which the different execution paths are executed. If a single output terminal has more than one connector to a target node, it is the broker (again, not you) which determines this execution order.
- A node does not always produce an output message for every output terminal: often it produces one output for a specific terminal depending on the message received. E.g. a filter node will typically send a message on either the true or false terminal, but not both.
- When the processing determined by one connector has been completed, the node reissues the message to the next connector, until all possible paths are completed. Updates to a message are never propagated to previously executed nodes, only to following nodes.
- The message flow can only start processing the next message when all paths through the message flow (that is, all connected nodes from all output terminals, as appropriate) have been completed.

Parsers

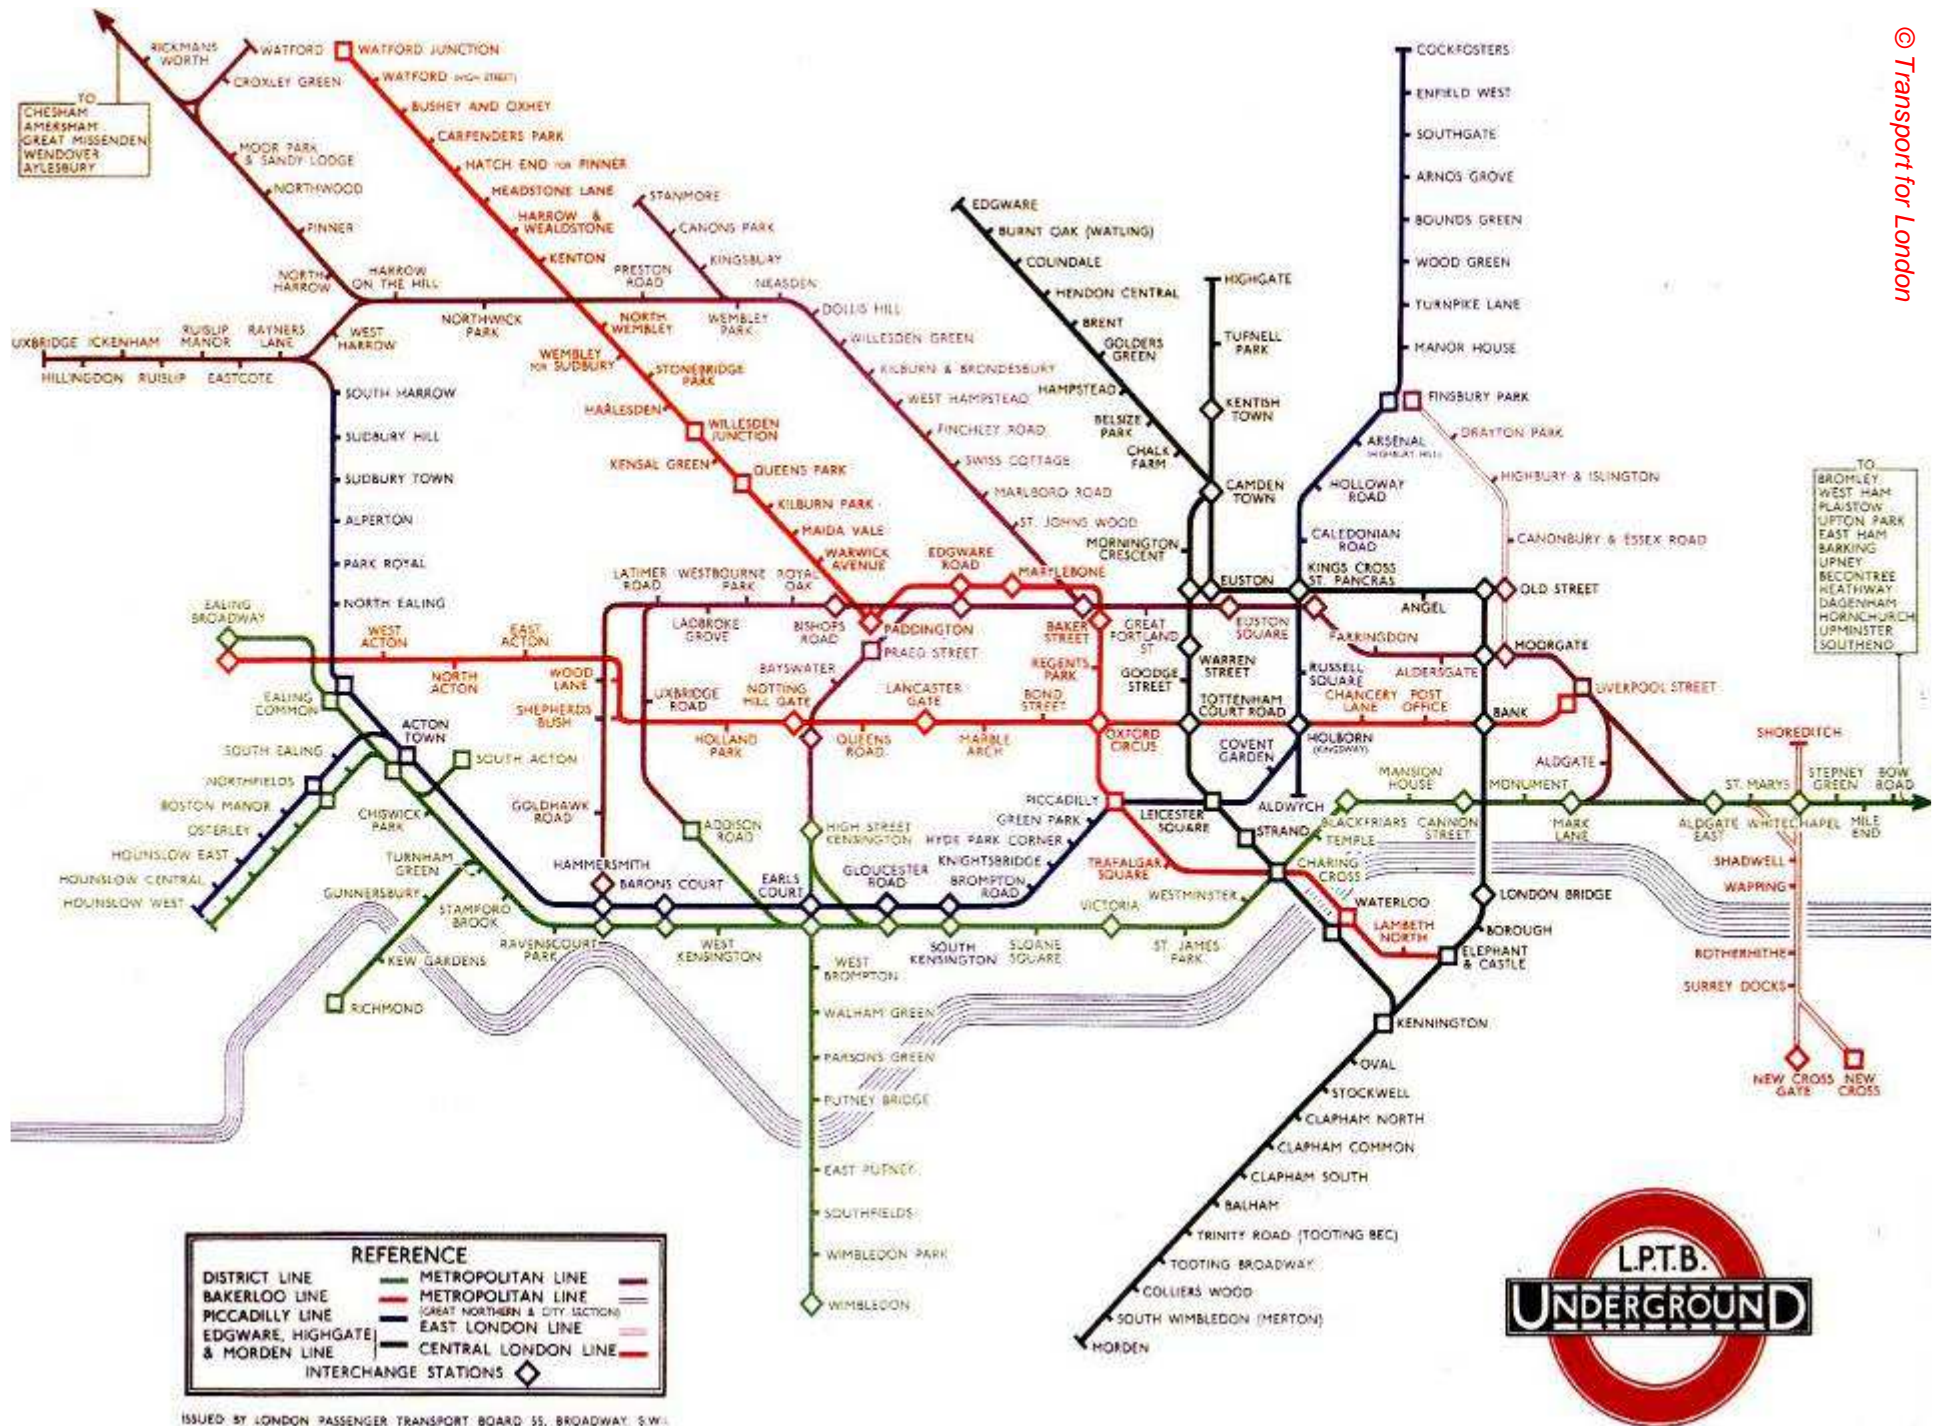


- On the previous slide we saw that objects called “message trees” are sent to a node’s input terminals, and either the same or different message tree is propagated from a node’s output terminals.
- The message tree is a logical definition of a message processed by the broker. It’s described as a tree because messages are typically hierarchical in structure; a good example of this is XML. Other message formats too, are also often derived from complex structures which themselves can be derived from complex structures, and so on, which gives them a tree-like shape with leaf nodes representing simple data types.
- In IBM Integration Bus, parsers have the job of converting between physical messages (bit-streams) and logical trees. When a message arrives at the broker through an input node, the message bit-stream is converted into a tree structure by the parser, which typically uses a model to drive the form of the logical tree. Built-in parsers handle well known headers within the message (MQMD, MQRFH2 etc.). Finally the user data is parsed into the tree using the domain parser as identified in the MQRFH2 (or input node). Message Broker’s built-in parsers support multiple domains (MRM, SOAP, XMLNSC, Data Object, XMLNS, JMSMap, JMSSStream, MIME, IDOC, BLOB and XML) to enable parsing of user and industry standard formats.
- As the logical tree is passed from node to node, the form of the logical tree may change depending on what the node is doing.
- When the message arrives at an output node, the parser converts the logical tree back into a physical bit-stream where it can be output to the external resource, where it can be read by the target (receiving) application. However, note that an output node need not indicate the end of a flow; it is possible to output to multiple destinations within a single invocation of a message flow. In this case, the logical tree can be passed on to other nodes and manipulated further, even after it has been converted back into a physical bit-stream for this particular output node.



© Transport for London

- This is a map of the London Underground railway system from 1921. Notice how the lines relate to the physical layout of the stations rather than the logical layout. It makes it really difficult to work with this map, particularly around the central area.
- Harry Beck (pictured) produced a new style map in 1933. Beck was an Underground employee who realised that, because the railway ran mostly underground, the physical locations of the stations were irrelevant to the traveller wanting to know how to get to one station from another — only the topology of the railway mattered.
- See http://en.wikipedia.org/wiki/Harry_Beck_%28graphic_designer%29 for more information.



REFERENCE	
DISTRICT LINE	METROPOLITAN LINE
BAKERLOO LINE	METROPOLITAN LINE (GREAT NORTHERN & CITY SECTION)
PICCADILLY LINE	EAST LONDON LINE
EDGWARE, HIGHGATE & MORDEN LINE	CENTRAL LONDON LINE
INTERCHANGE STATIONS	

- This is Harry's map from 1933.
- He had devised a vastly simplified map, consisting of stations, straight line segments connecting them, and the Thames; lines ran only vertically, horizontally, or at 45 degrees. The Underground was initially sceptical of his proposal — it was an uncommissioned spare-time project, and it tentatively introduced it to the public in a small pamphlet. It was immediately popular, and ever since the Underground has used topological maps to illustrate the network.
- In Integration Bus terms, Harry was the “parser” – he took the physical representation of the London Underground and converted into a logical structure. This logical structure is very much easier to work with.
- Integration Bus uses logical structures to describe physical data for similar reasons: it makes them much easier to work with, particularly when addressing or converting between data elements. We'll see examples of this later on.
- See http://en.wikipedia.org/wiki/Tube_map for more information.

Message Modeling

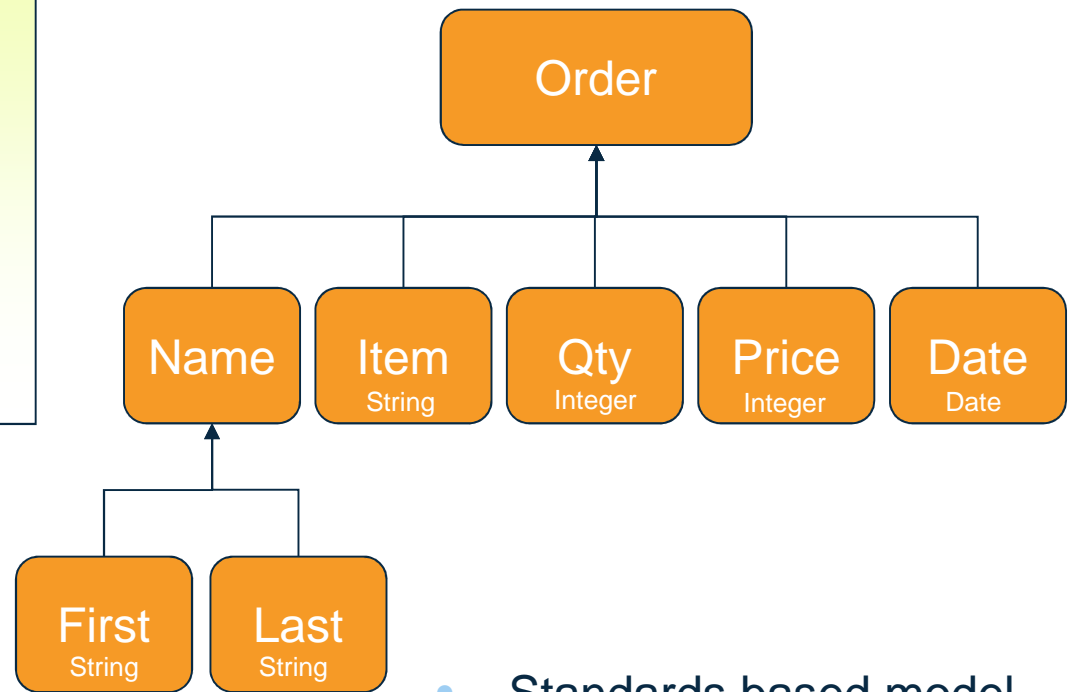


```

<order>
  <name>
    <first>John</first>
    <last>Smith</last>
  </name>
  <item>Graphics Card</item>
  <quantity>32</quantity>
  <price>200</price>
  <date>07/11/09</date>
</order>
  
```

John,Smith,Graphics Card,
32,200,07/11/09

John Smith.....
Graphics Card.....
3220020071109.....

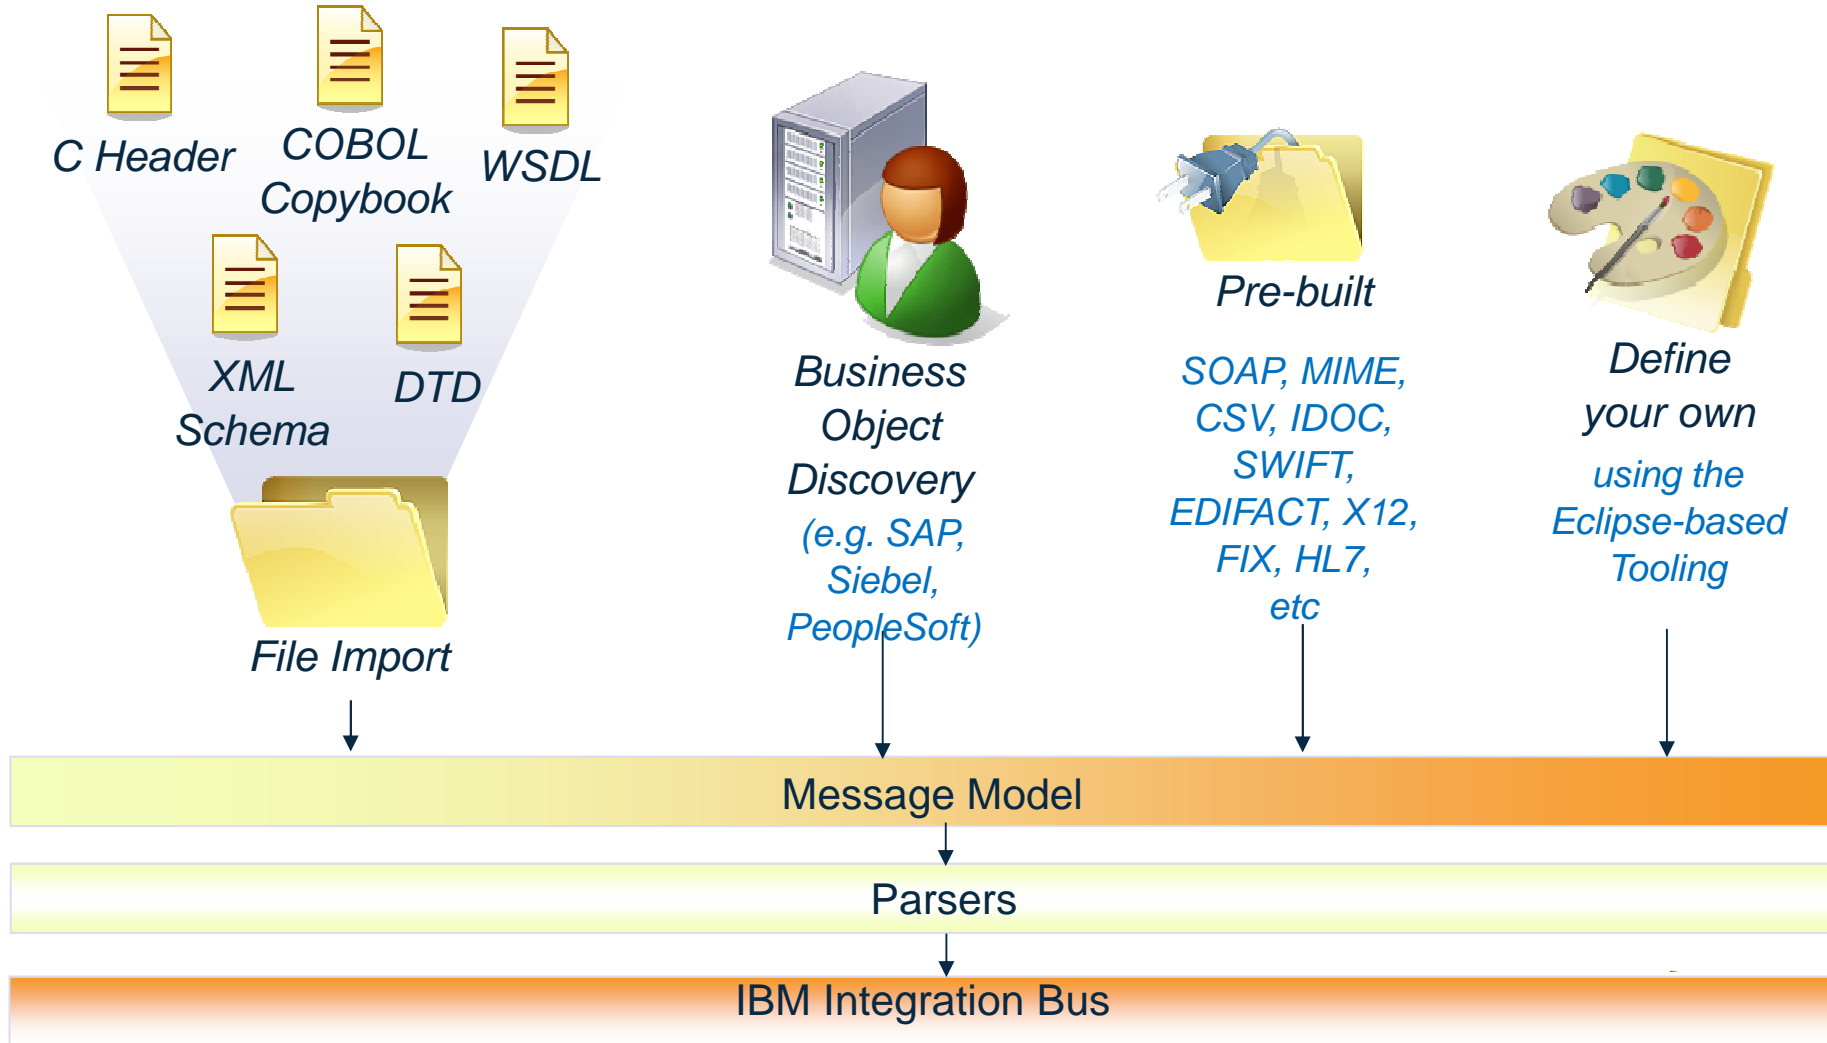


- Standards based model definitions
 - XML
 - DFDL

- Here is an example of how a physical data structure could be mapped to a logical tree.
 - Notice how multiple physical formats can correspond to the same logical tree. The first physical format is an XML structure that shows our Order message. The second is a comma separated value (CSV) structure of the same. The third comprises a set of fixed length fields in a custom wire format.
 - By manipulating the logical tree inside the Integration Broker rather than the physical bit-stream, the nodes can be completely unaware of the physical format of the data being manipulated. It also makes it easy to introduce new message formats into the broker.
- Applications have and require diverse data formats.
 - We all know that XML is the data format that's going to solve every data processing problem that exists! We also know that "XML++", the follow-on compatible meta format that someone in a research laboratory is working on will solve all the problems we don't even know we have today! The fact is that, without wanting to appear cynical, every generation goes through this process. Surely it was the same when COBOL superseded assembler.
 - The fact is, that for historic, technical, whimsical, political, geographical, industrial and a whole host of other reasons you probably never even thought of, a hugely diverse range of data formats exist and are used successfully by a myriad of applications every second of every day. It's something that we have to live with and embrace because it isn't going to get any better any time soon.
 - The advantage IBM Integration Bus brings by modelling all these messages is that we can rise above the message format detail; so that whether it's a tag delimited SWIFT or EDIFACT message, a custom record format closely mapping a C or COBOL data structure, or good old XML, we can talk about messages in a consistent, format independent way. Integration Bus can manage this diversity.

- The Logical Message Model.
 - Reconsider messages and their structure. When we architect messages (no matter what the underlying transport technology), we concern ourselves firstly with the logical structure. For example, a funds transfer message might contain an amount in a particular currency, a transaction date and the relevant account details of the parties involved. These are the important business elements of the message; when discussing the message, we refer to these elements.
 - However, when we come to realize the message, we have to choose a specific data format. This may be driven by many factors, but we have to choose one. You may be aware of the advantages of various message formats or have your own personal favourite, or may fancy inventing a new one, but the fact remains that you have to choose a physical *wire format*. So for our transfer message, we might decide to use XML, with its elements, attributes and PCDATA (and a DTD, if we're being really exact), or we might map more closely to a C data structure modelling our message with ints, shorts, chars etc. and worry about *their* various representations(!)
 - The Logical message model provided by Integration Bus allows one to describe a message in terms of a tree of elements, each of which has a (possibly user defined) type. At the message tree leaf nodes, the elements have simple types such as strings, integers, decimals, booleans etc. Moreover, elements can have various constraints and qualifiers applied to them that more fully describe them; e.g. elements might be optional, appear in a certain order or only contain certain values.

Creating message models



Complete your session evaluations online at www.SHARE.org/Seattle-Eval

Notes



- This slide describes some of the options available for creating message models.
- Firstly, if you have messages described by COBOL copybooks, C header files XML DTDs/Schemas or WSDL, use the WebSphere Message Broker supplied importers to generate your message model. A wide range of importers exist, so that you can kick start your message modelling.
- Secondly, if you wish to use the SAP, Siebel or PeopleSoft nodes inside WebSphere Message Broker V6.1, you can construct message models directly from the Business Objects on these systems.
- Thirdly, you can use pre-built models such as those for SWIFT messages.
- Finally, you can use graphical modelling available in the Message Broker Toolkit to model your messages. You've seen application connections and processing constructed using message flows and nodes; the Message Broker Toolkit provides a similarly visual approach to message modelling.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



Powerful Message Transformation Options



Mapping

- Graphical, easy to use
- Drag and Drop fields, apply functions



Java Compute

- Embed Java programs
- Ability to use XPath for tree access



Compute

- Describe powerful transformations quickly
- Uses SQL-based language (ESQL)




XSL Transform

- Convert XML to anything
- Uses standard XSL Style sheets



.NET Compute

- Use any of the 40+ .NET languages (e.g. C#, VB.NET)
- Access COM objects
- Windows only



- There are several options available to you out-of-the-box for transforming between message formats:
 - Use Graphical Mapping to visually represent messages and transform them. The Integration broker has a mapping node to allow users to visualize and transform messages. The mapping node presents input and output message views; i.e. visualisation of message definitions. Users can “map” elements from the input message to the output message using “drag and drop” operations. More complex operations are possible, such as field concatenation. Graphical mapping is most effective when you have relatively simple transformations to perform and you don’t want to use a programming language (ESQL or Java).
 - XSLT (eXtensible Stylesheet Language Transformations) is a language for describing message transformations. There is a node in IBM Integration Bus that allows you to convert XML messages using style sheets developed in this language.
 - The Compute node uses ESQL (Extended Structured Query Language). This is a language derived from SQL3, and is particularly suited to manipulating both database and message data. You do this with a single syntax (words) which has a common semantic (meaning). ESQL addresses message fields using a natural syntax. For example, the fourth traveller in a travel request message could be:
`InputRoot.Body.TravelRequestMessage.TravellerDetails[4].LastName`
 - ESQL has a rich set of basic data types and operators, as well as the kind of statements and functions you're used to from procedural programming languages, to allow you to perform powerful transformations.

- Continuing the options available to you for transforming between message formats:
 - You can also use the power of Java to route and transform your messages. The JavaCompute node is fully integrated into the Eclipse Development Environment to providing usability aids such as content assist and incremental compilation. You can refer to elements using an expressive XPATH syntax, so that message navigation and element creation and modification are vastly simplified, and comparable in simplicity to ESQL field references. JDBC type 4 support allows you to perform database and message tree operations in the Java Compute node. On z/OS Java workload is eligible for offload onto the zSeries Application Assist Processors zAAP.
 - .NET is available as an option on Windows platforms, and allows you to transform data and invoke general purpose integration logic on any of the .NET supported platforms. Templates for Visual Studio and an integrated Visual Studio debugger make this a natural choice for Windows programmers.
 - for example, you can use this node to integrate your Excel spreadsheets!
- You can mix and match your transformation styles, or use just one throughout your enterprise. It will probably be your skill set which determines whether you chose to use Graphical Mapping, Java, XSLT or ESQL to perform your message routing and transformation.
- Note that as a separate product, you can also use WebSphere Transformation Extender maps to describe transformations.

Easily Address Message Elements



JavaCompute

```
public class jcn extends MbJavaComputeNode {
    public void evaluate(MbMessageAssembly assembly) throws MbException {
        ...
        String lastName =
            (String)assembly.getMessage().evaluateXPath("/Body/Order/Name/Last");
        ...
    }
}
```



Route

Route Node Properties - Route

Filter table*

Filter pattern	Routing output terminal
\$Body/Order/Price > 1000	BigSpenders
\$Body/Order/Price < 200	Cheapskates



DataInsert

```
IF Body.Order.Date < '2008/01/01' THEN
    INSERT INTO Database.OldOrders (LastName,Item,Quantity)
    VALUES (Body.Order.Name.Last,
            Body.Order.Item,
            Body.Order.Quantity);
ENDIF;
```

- Each node's configuration (which includes some Java logic in the case of the Java Compute node) dictates what you want the node to do, and this may include manipulation of one or more elements in the message tree.
- Here are some examples of node configurations that address elements in the logical tree.
- In most cases, elements can be addressed using either XPATH (as shown in the JavaCompute and Route) or ESQL (as shown in the DataInsert node).

Development

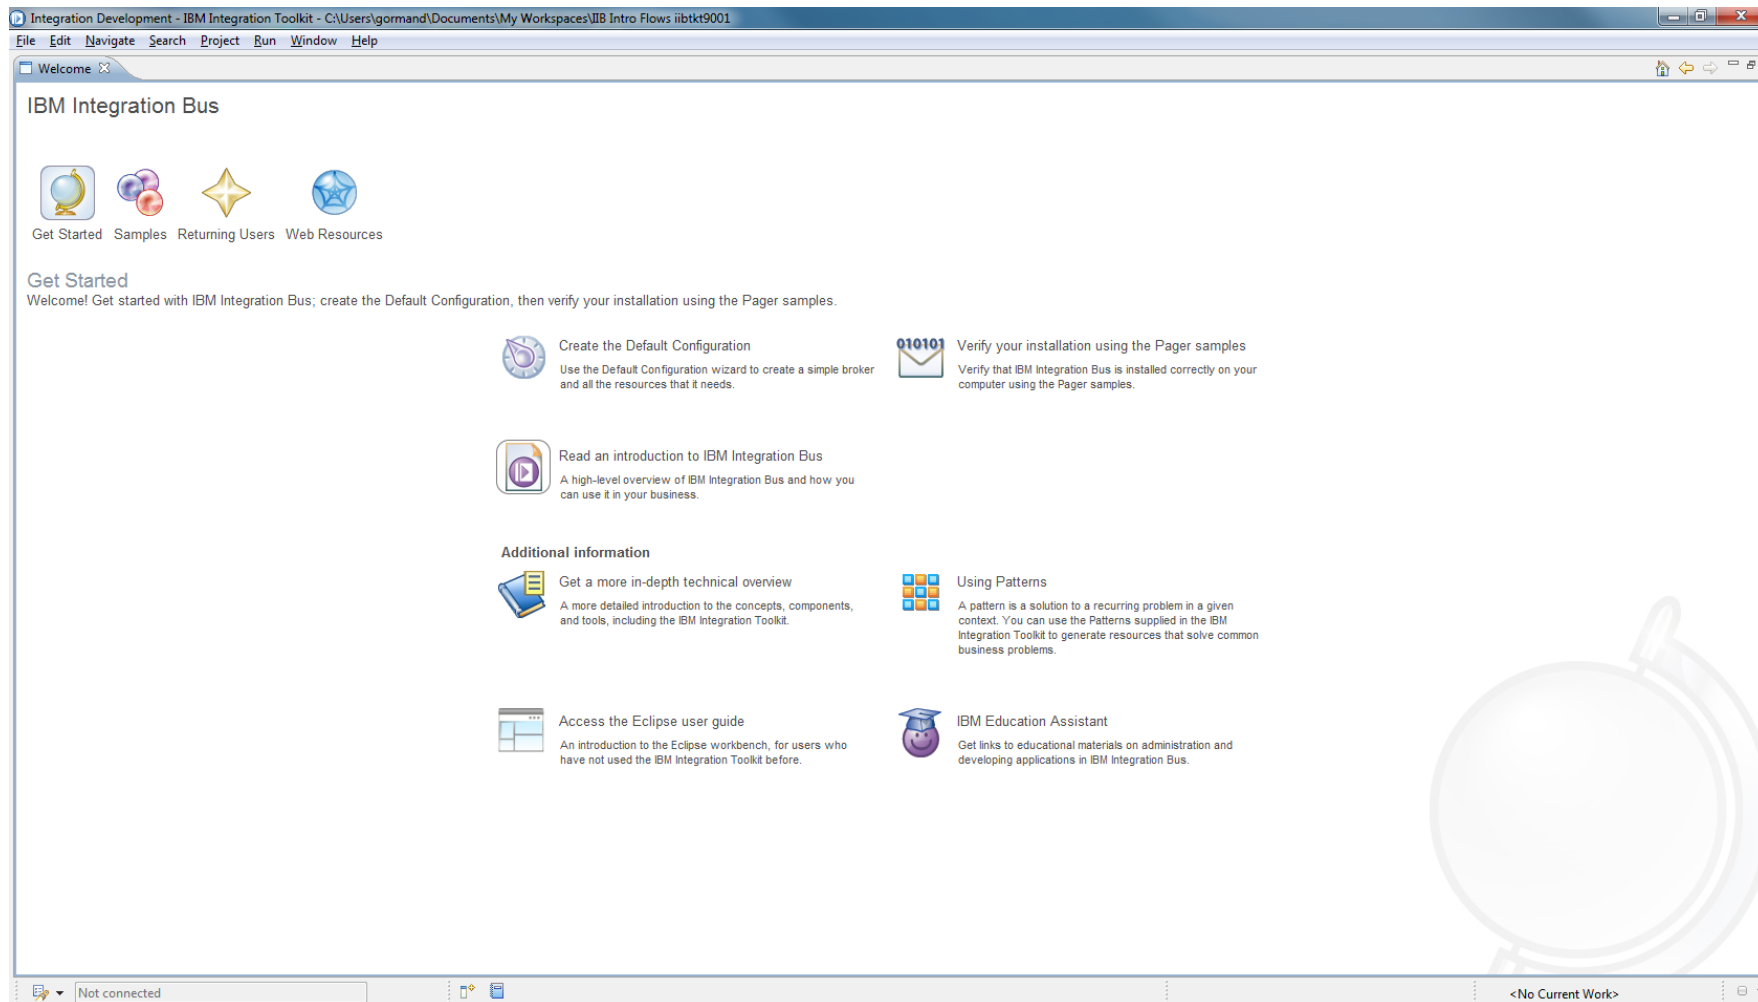
Complete your session evaluations online at www.SHARE.org/Seattle-Eval



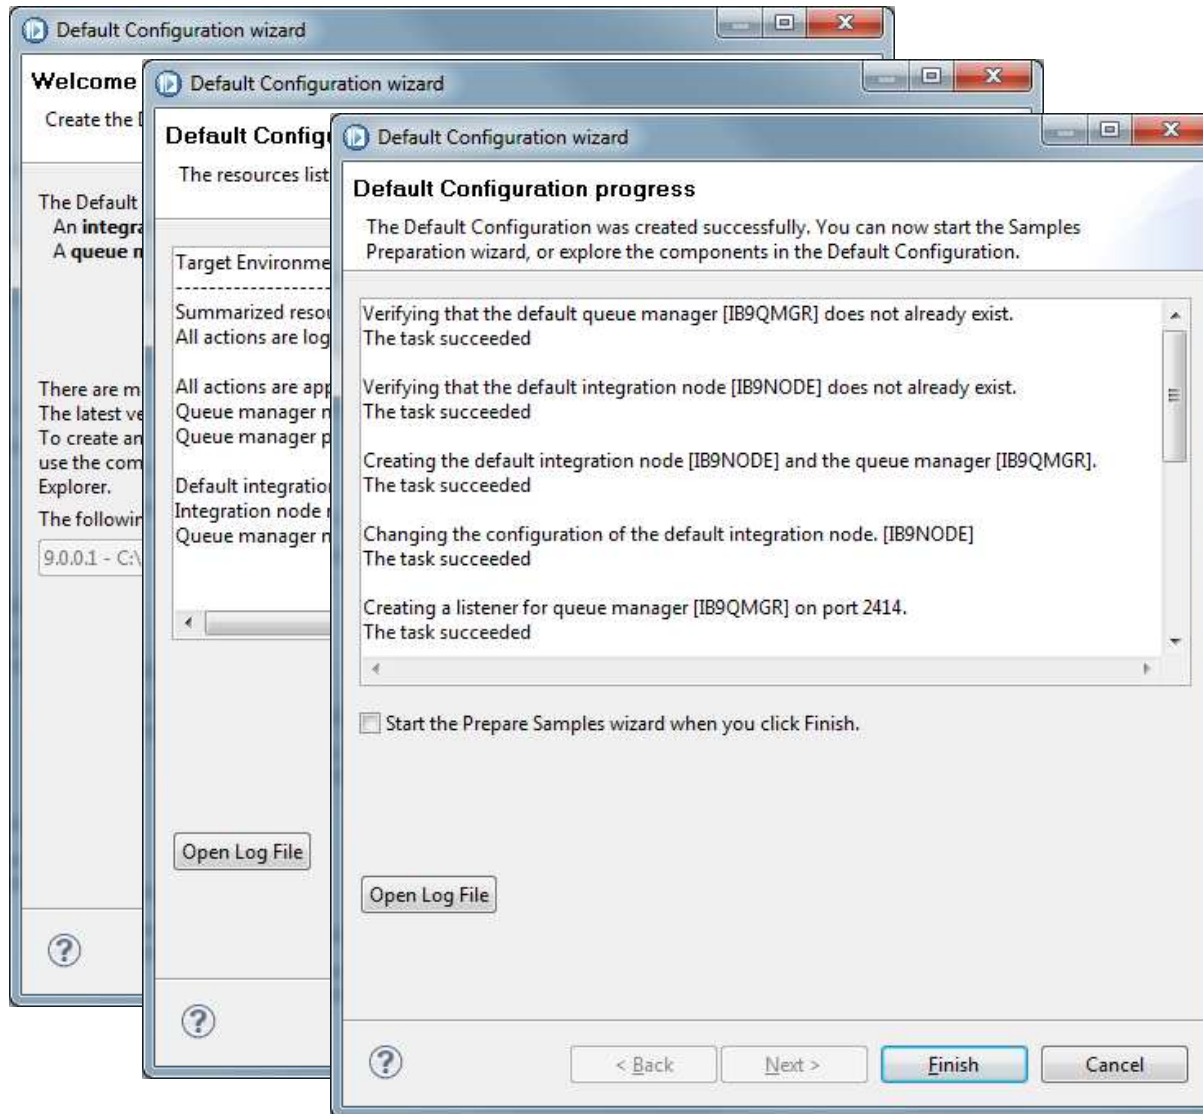
Integration Toolkit



- Eclipse based tooling for developing integrations
- Easy “getting started” accelerators

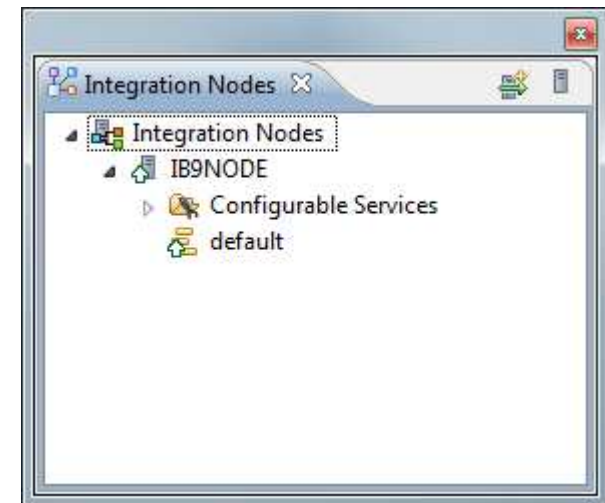


Default Configuration Wizard



Use the DCW to easily create:

- Default Queue Manager
- Default Integration Node
- Default Integration Server



Quick Starts

Start building your application with one of the following tasks.

[Start by creating an application](#)

An **Application** is a container for all the resources that are required to create a solution. [More...](#)

[Start by creating an integration service](#)

An **Integration Service** is an application with a well-defined interface and structure. [More...](#)

[Start by creating a library](#)

A **Library** is a logical grouping of related code, data, or both. [More...](#)

[Start from WSDL and/or XSD files](#)

Use this task to create an **Integration Service, Application or Library** which includes your WSDL and/or XSD files.

[Start by discovering a service](#)

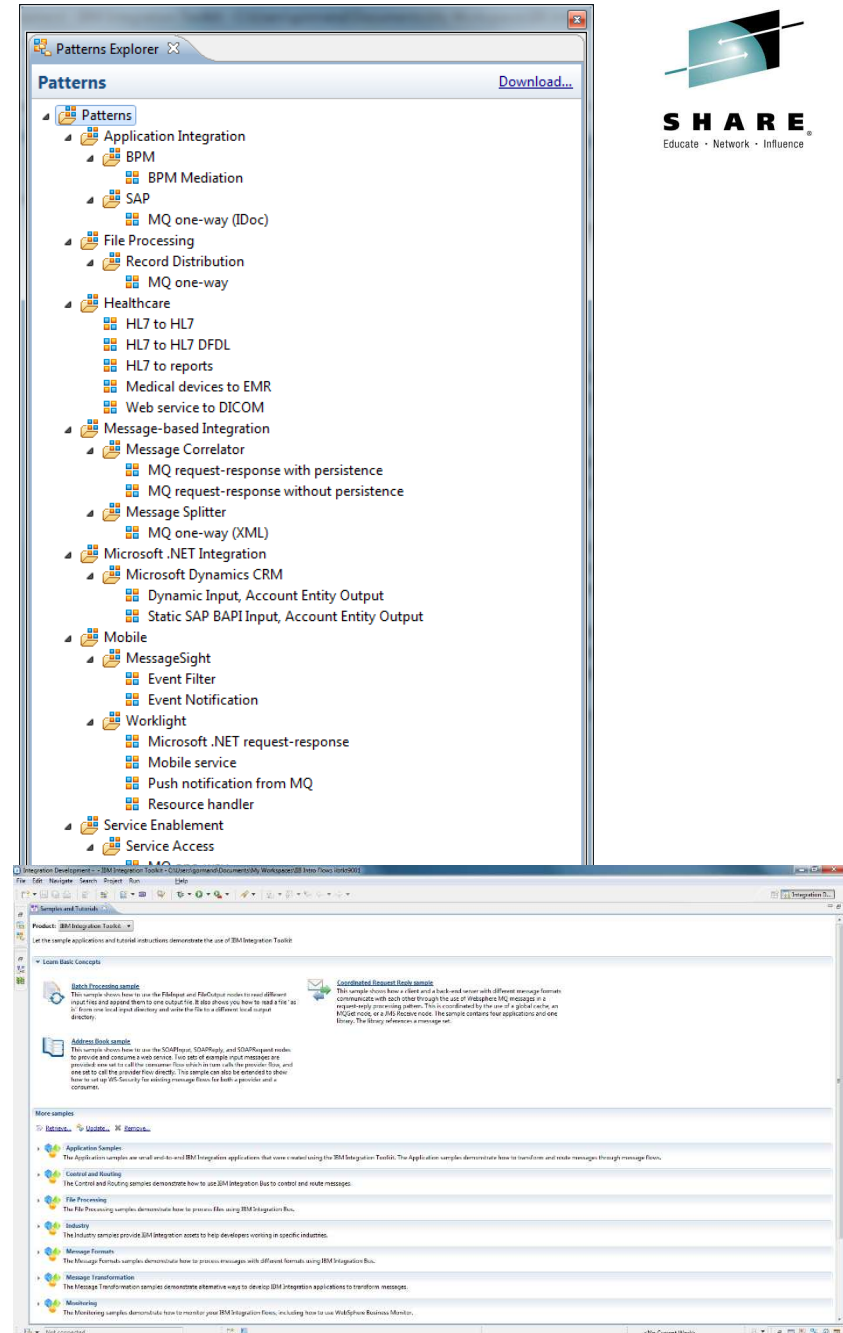
A **Service** allows you to invoke the remote system using discovered resources.

[Start from patterns](#)

A **Pattern** is a reusable solution that encapsulates a tested approach to solving a common architecture, design, or deployment task. [More...](#)

[Start from samples](#)

Use the **Samples** to learn more about the features in IBM Integration Bus. [More...](#)



The screenshot displays two windows from the IBM Integration Bus interface. The top window, titled "Patterns Explorer", shows a hierarchical tree of patterns under the "Patterns" folder. The tree includes categories such as Application Integration, File Processing, Healthcare, Message-based Integration, Microsoft .NET Integration, Mobile, and Service Enablement. Each category contains specific pattern names, such as "BPM Mediation", "MQ one-way (IDoc)", "HL7 to HL7", and "Message Correlator". A "Download..." link is visible in the top right corner of the Patterns Explorer window.

The bottom window, titled "Samples and Catalogs", displays a list of sample applications and their descriptions. The list includes:

- Back Processing sample**: This sample shows how to use the FileInput and FileOutput nodes to read different input files and append them to one output file. It also shows how to read a file as an event and use the event data to write the file to a different local output directory.
- Coordinated Request Reply sample**: This sample shows how a client and a back-end server with different message formats communicates with each other through the use of WebSphere MQ messages in a request-reply processing pattern. This is coordinated by the use of a global entity, an MQCorr node, or a JMS Receive node. The sample contains four applications and one library. The library references a message map.
- Address Book sample**: This sample shows how to use the SOAPRequest, SOAPReply, and SOAPRequest nodes to process and generate a web service. This set of samples (SOAPRequest and SOAPReply) are used to call the webservice flow which in turn calls the provider flow, and one or more call the provider flow directly. This sample can also be extended to show how to set up WS-Security for existing message flows for both a provider and a consumer.


Below the list, there is a section titled "More samples" with a list of categories and their descriptions:

- Application Samples**: The Application samples are small and IBM Integration applications that were constructed using the IBM Integration Toolkit. The Application samples demonstrate how to transform and route messages through message flows.
- Control and Routing**: The Control and Routing samples demonstrate how to use IBM Integration Bus to control and route messages.
- File Processing**: The File Processing samples demonstrate how to process files using IBM Integration Bus.
- Industry**: The Industry samples provide IBM Integration assets to help developers working in specific industries.
- Message Formats**: The Message Formats samples demonstrate how to process messages with different formats using IBM Integration Bus.
- Message Transformation**: The Message Transformation samples demonstrate alternative ways to develop IBM Integration applications to transform messages.
- Monitoring**: The Monitoring samples demonstrate how to monitor your IBM Integration flows, including how to use WebSphere Business Monitor.




Logically separate your applications

- You can separate your applications during development and runtime
- Application
 - means of encapsulating resources to solve a specific connectivity problem
 - application can reference one or more libraries



- Library
 - a logical grouping of related routines and/or data
 - libraries help with reuse and ease of resource management
 - library can reference one or more libraries



- Feature of Application and Libraries
- Applications promote encapsulation and isolation
 - Typically contain “main” message flows and dependent resources
 - ESQL, Java, Maps, Message models, subflows, Adapter files, etc.
 - dependent resources could live in referenced libraries
 - Multiple applications can be packaged into a single BAR
 - Multiple applications can be deployed to an execution group
 - Visibility of resource restricted to containing application
 - Referenced libraries are deployed inside application container (by copy)
- Libraries facilitate re-use and simplify resource management
 - Typically contain reusable helper routines and resources
 - Subflows, ESQL, Java, Maps, Message models, Adapter files, etc.
 - Use multiple libraries to group related resources (e.g. by type or function)
 - Multiple applications can reference the same library
 - Each application gets its own copy of the library during package/deploy
 - Libraries are packaged as part of referencing application in the BAR
 - Library can reference other Libraries

Samples



The screenshot displays the IBM Integration Toolkit interface. The main window shows a list of sample applications under the heading "Learn Basic Concepts".

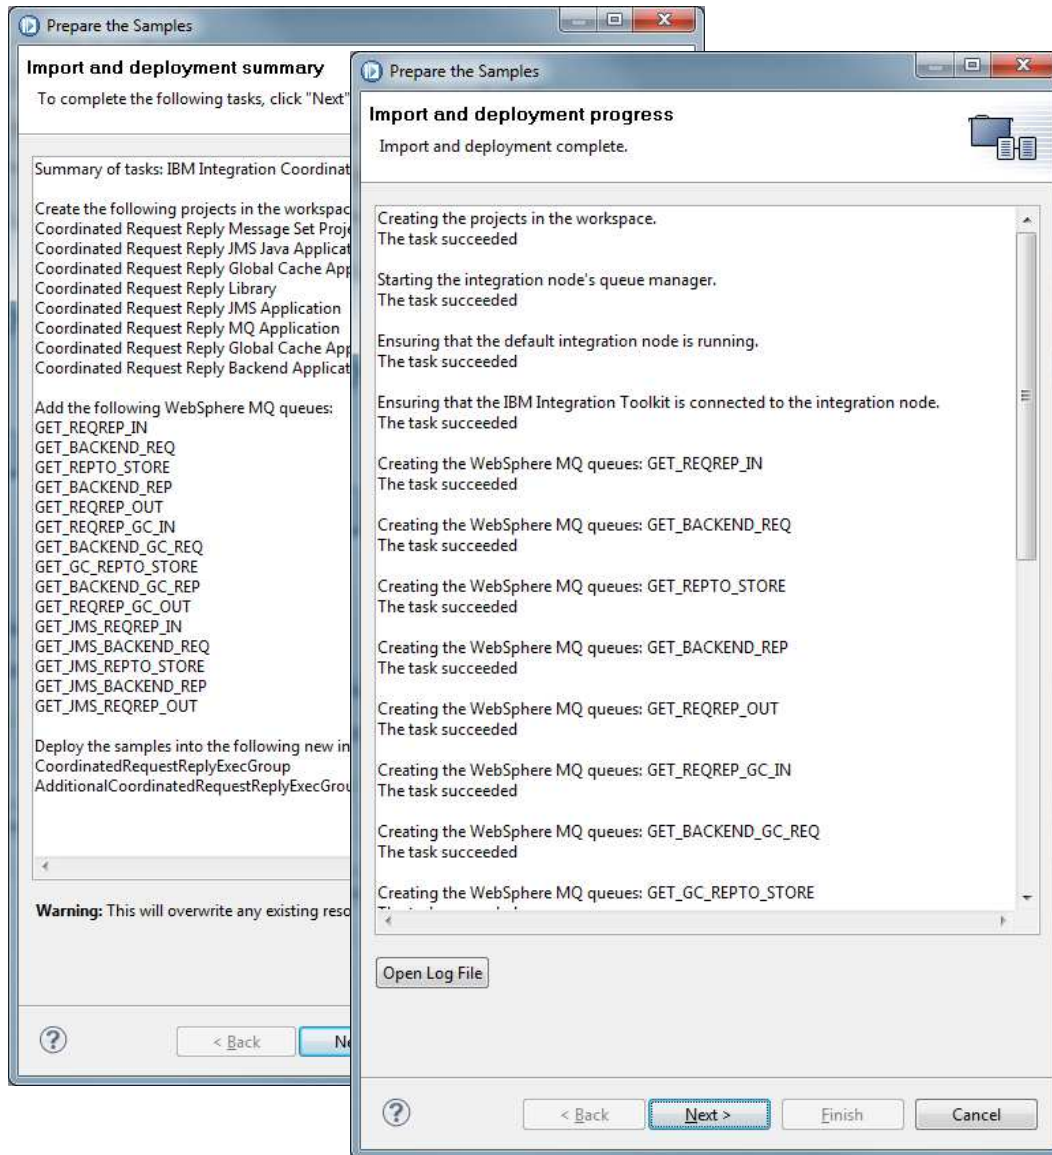
- Batch Processing sample**: This sample shows how to use the FileInput and FileOutput nodes to read different input files and append them to one output file. It also shows you how to read a file "as is" from one local input directory and write the file to a different local output directory.
- Address Book sample**: This sample shows how to use the SOAPInput, SOAPReply, and SOAPRequest nodes to provide and consume a web service. Two sets of example input messages are provided: one set to call the consumer flow which in turn calls the provider flow, and one set to call the provider flow directly. This sample can also be extended to show how to set up WS-Security for existing message flows for both a provider and a consumer.
- Coordinated Request Reply sample**: This sample shows how a client and a back-end server with different message formats communicate with each other through the use of WebSphere MQ messages in a request-reply processing pattern. This is coordinated by the use of a global cache, an MQGet node, or a JMS Receive node. The sample contains four applications and one library. The library references a message set.

Below the "Learn Basic Concepts" section, there is a "More samples" section with a list of application samples:

- Application Samples**: The Application samples are small end-to-end IBM Integration applications that were created using the IBM Integration Toolkit. The Application samples demonstrate how to transform and route messages through message flows.
- Airline Reservations sample**: This sample shows how to use a range of nodes, including nodes for aggregation, routing, tracing, filtering, and updating database tables.
- Data Warehouse sample**: This sample contains a message flow that archives data, such as sales data, to a database.
- DatabaseInput Node sample**: This sample demonstrates how to take data from a database, as it is being updated, and process the data within IBM Integration Bus.
- Error Handler sample**: This sample contains a message flow and a subflow to show error handling in message flow applications.
- Graphical Data Mapping Loyalty Data Warehousing sample**: This sample contains a message flow that inserts and updates rows in a database.
- Large Messaging sample**: This sample shows how to process messages that contain repeating structures, and how to minimize the virtual storage requirements for the message flow.
- Message Routing sample**: This sample shows how to use a message flow to route messages to different WebSphere MQ queues based on data stored in a database table or a file.
- Pager samples**: This sample shows simple point-to-point messaging and publish/subscribe messaging. Use graphical interfaces to send text messages to a pager application, or to subscribe to publications about the surf on selected beaches.
- Scribble sample**: This sample is a small, graphical whiteboard application on which you draw using your mouse pointer. Depending on the options you choose, you can see the effects of message transformation using WebSphere MQ transport or RealTime transport.
- Solar Pattern Authoring sample**: This sample shows how to build an IBM Integration pattern. The sample provides an example integration project that calculates the sunrise and sunset times in a PHPCompute node. The sample also provides a pattern authoring project that

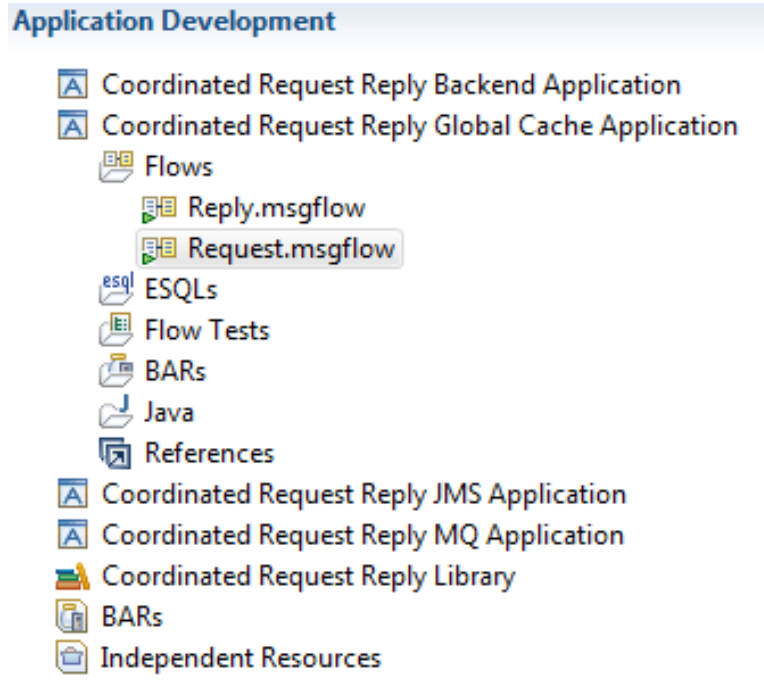
Complete your session evaluations online at www.SHARE.org/Seattle-Eval

Import and Deploy a product sample

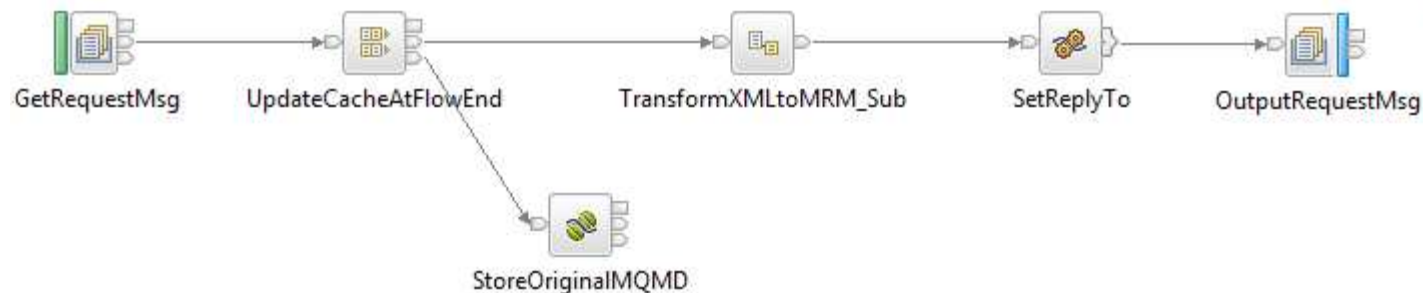


- There are many product samples that show how to use IIB
- These are easily imported into the development toolkit workspace, and deployed to the integration server runtime

Sample artefacts



- This sample has imported several integration data flows to demonstrate a coordinated request/reply scenario
- The message flows are grouped in applications and libraries as can be seen to the left
- Below is one of the message flows for this sample. It is a complete transaction



Generate Pattern Instances



New Pattern Instance

Create New Instance

The instance name uniquely identifies the pattern instance project that is created when the pattern instance is saved or generated. The name ideally reflects the business and integration functions that the pattern supports.

Pattern instance name:

Configure Pattern Parameters

Provide values for pattern parameters. Click the "Generate" button or click [here](#) to generate a pattern instance.

❗ The pattern parameter "Service WSDL" is mandatory but a value is not set. [more...](#)

Pattern Parameters	Pattern Parameters Details						
<p>Service WSDL * <Select deployable service WSDL> <input type="button" value="Browse..."/></p> <p>Validation of SOAP request: <input type="text" value="None"/></p> <p>Validation of SOAP response: <input type="text" value="None"/></p> <p>Provider information <input checked="" type="checkbox"/></p> <p>Provider address: <input type="text"/></p> <p>Provider queue manager: <input type="text"/></p> <p>Provider request queue *: <input type="text" value="PROVIDER"/></p> <p>Response information <input checked="" type="checkbox"/></p> <p>Response queue and validation requirements:</p> <p>Response queue *: <input type="text" value="RESPONSE"/></p> <p>Store queue *: <input type="text" value="STORE"/></p> <p>Expiry of store queue messages in tenths of a second *: <input type="text" value="No timeout"/></p> <p>Logging <input checked="" type="checkbox"/></p> <p>Enable logging messages and specify their destination:</p> <p>Logging required: <input type="checkbox"/></p> <p>Log queue manager: <input type="text"/></p> <p>Log queue *: <input type="text" value="LOG"/></p>	<p>Service information</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th>Pattern parameter</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Service WSDL</td> <td> <p>This pattern parameter defines the deployable WSDL file that defines the service interface presented by the broker to requesters.</p> <p>This value must include the path name starting at the workspace message set project.</p> <p>To navigate to the required WSDL, click Browse.</p> </td> </tr> <tr> <td>Validation of SOAP request</td> <td> <p>Incoming service requests are validated as defined by this parameter. Valid values are:</p> <ul style="list-style-type: none"> • None • Content and value </td> </tr> </tbody> </table>	Pattern parameter	Description	Service WSDL	<p>This pattern parameter defines the deployable WSDL file that defines the service interface presented by the broker to requesters.</p> <p>This value must include the path name starting at the workspace message set project.</p> <p>To navigate to the required WSDL, click Browse.</p>	Validation of SOAP request	<p>Incoming service requests are validated as defined by this parameter. Valid values are:</p> <ul style="list-style-type: none"> • None • Content and value
Pattern parameter	Description						
Service WSDL	<p>This pattern parameter defines the deployable WSDL file that defines the service interface presented by the broker to requesters.</p> <p>This value must include the path name starting at the workspace message set project.</p> <p>To navigate to the required WSDL, click Browse.</p>						
Validation of SOAP request	<p>Incoming service requests are validated as defined by this parameter. Valid values are:</p> <ul style="list-style-type: none"> • None • Content and value 						

- Generate message flows based on patterns which define best practice
- Properties allow the pattern instance to be customised
- User-Defined patterns can be authored and distributed to other developers to enforce best practices



Pattern Artefacts

Summary for pattern instance MQRequestResponsePattern

To complete pattern application MQRequestResponsePattern, review the actions in this summary file.

Flow generation

This service facade has been generated from the Service Facade to WebSphere MQ: request-response pattern, which creates a bridge between the synchronous HTTP protocol that is typically used with Web services and existing applications with WebSphere MQ interfaces.

Project MQRequestResponsePattern_Flows has been created. This project includes the following message flows:

- [Request](#)
- [Response](#)

and subflows:

- [RequestProcessor](#)
- [ResponseProcessor](#)
- [Error](#)

Tasks to complete

Create queues

The following queue managers and queues must exist before you can run the pattern. Broker queues on the broker queue manager:

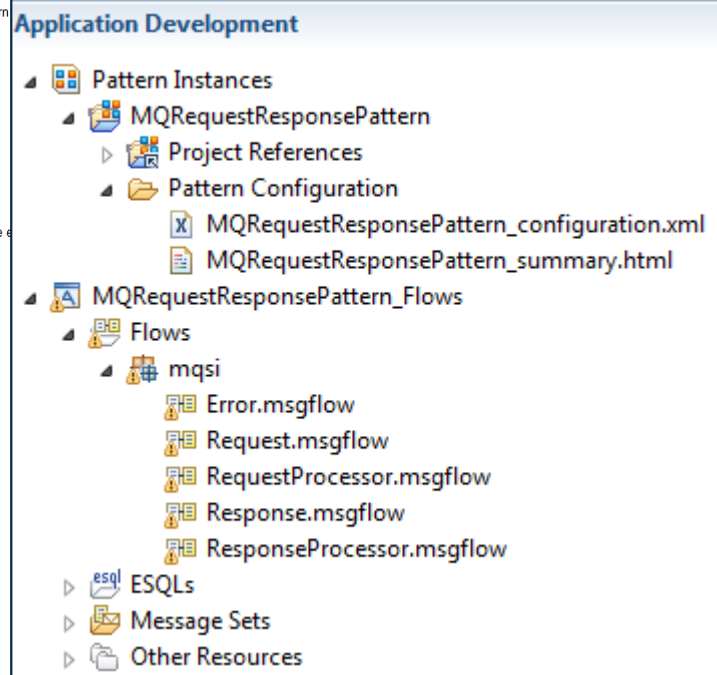
- Response queue: RESPONSE
- Store queue: STORE

Other queues:

- Error queue: ERROR on the broker queue manager
- Provider request queue: PROVIDER on the broker queue manager

Broker archive

Add this pattern instance to a broker archive file for deployment. In the Broker Archive e



- Message flows are generated in the pattern instance
- They are ready to be deployed to the integration server runtime
- Any tasks required to run the message flows are listed, such as creating MQ Queues

Services



- Within IIB, easily create new services that have a well defined interface and structure:
 - A whole new service from scratch including the associated WSDL and message flow
 - A service including the message flows, from an existing WSDL file
 - A service, including the message flows, from an existing IBM BPM service
 - A service based on an existing database
 - A service based on an existing MQ queue manager and queue definition

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

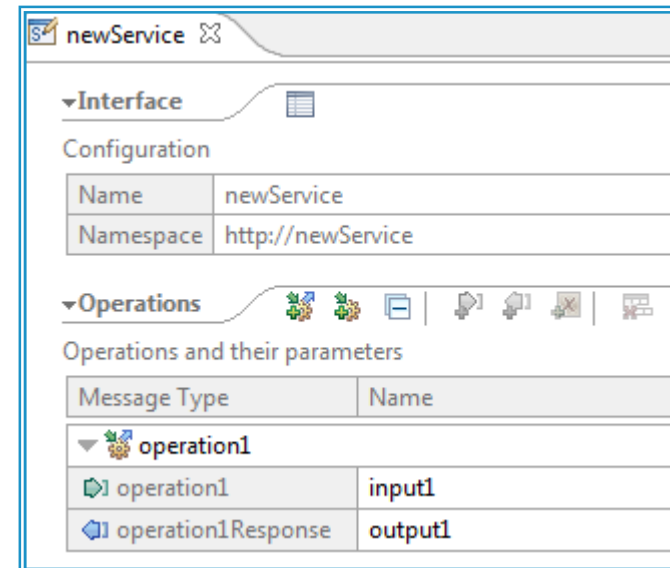
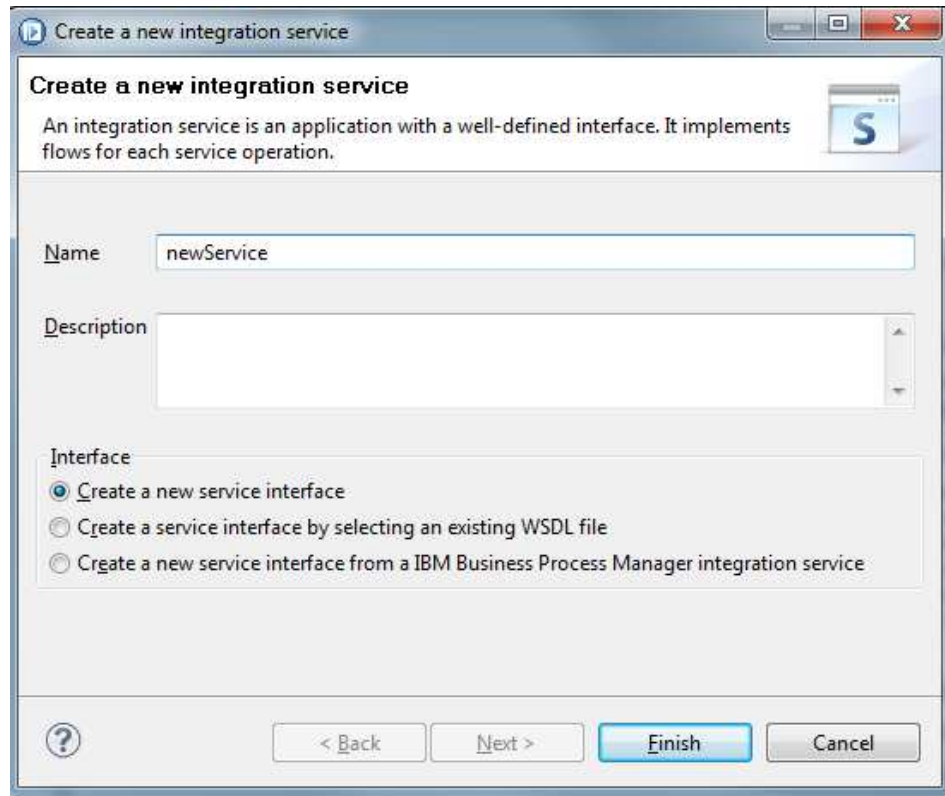


Creating a new Service

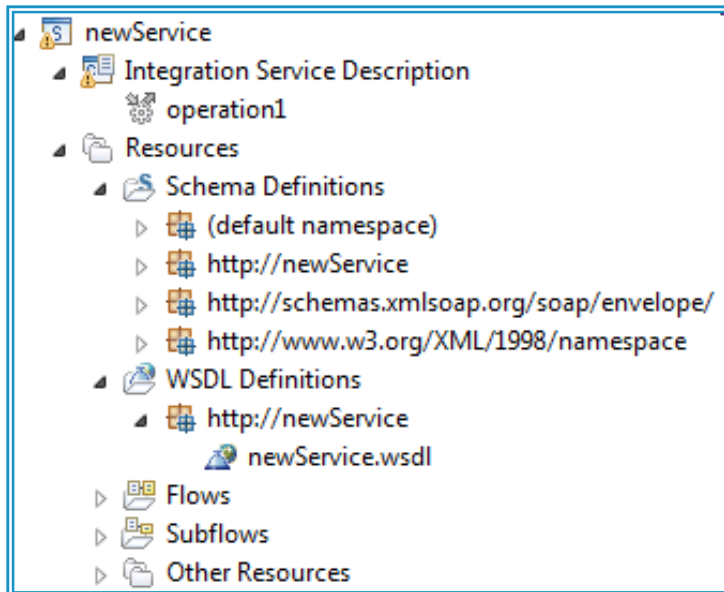
[Start by creating an integration service](#)

An **Integration Service** is an application with a well-defined interface and structure. Like an application, it is a container for resources required to create a Web Services solution. These resources provide the implementation of the interface with specified operations that you implement as separate message subflows. You can deploy a service. A web service consumer can interrogate the deployed service to return its interface. [Collapse](#)

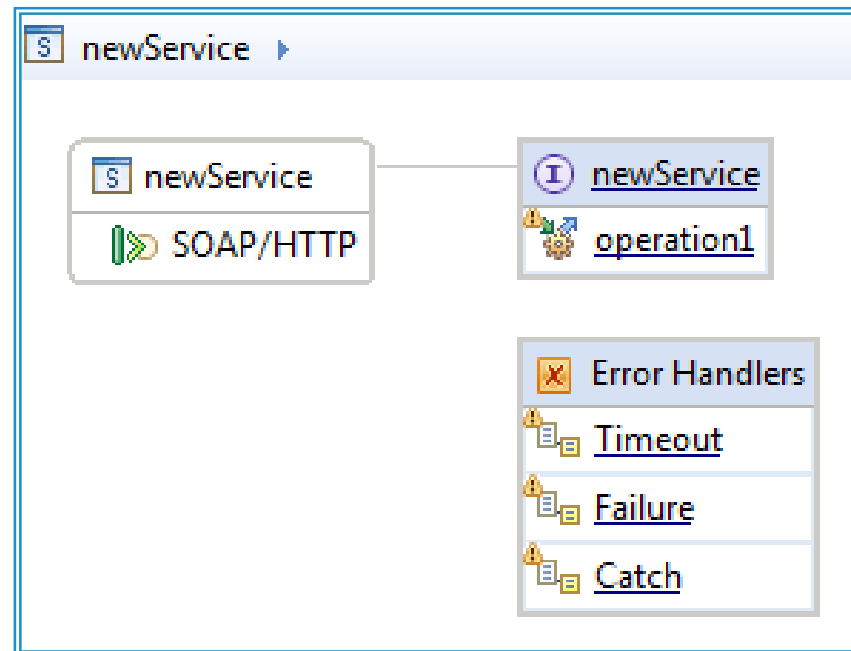
- Use the wizard to create a brand new service



Service Artefacts



- All artefacts are created for the new service, including the WSDL and a SOAP based flow
- A subflow is included for the operation and error handlers which the developer can then update



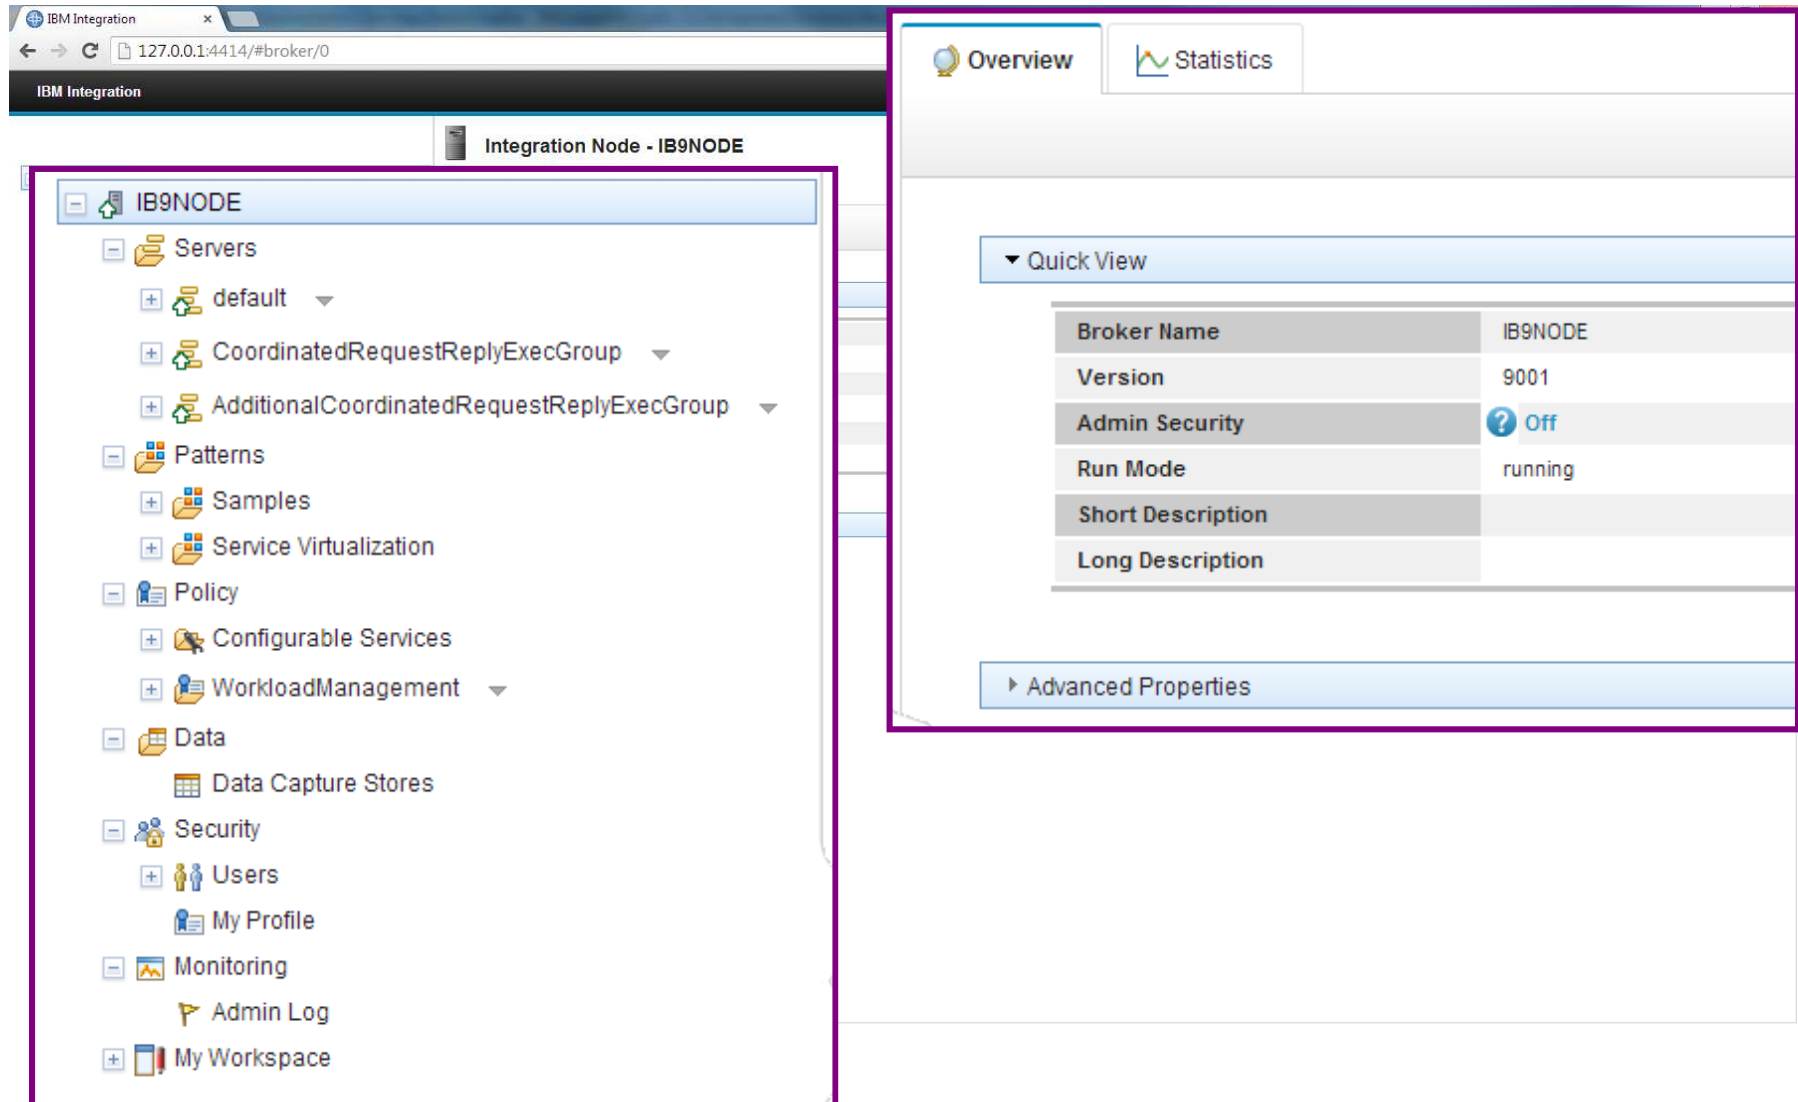


Administration

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



Administration using the WebUI



The screenshot displays the IBM Integration WebUI interface. The browser address bar shows the URL `127.0.0.1:4414/#broker/0`. The main content area is titled "Integration Node - IB9NODE".

Navigation Tree (Left Panel):

- IB9NODE
 - Servers
 - default
 - CoordinatedRequestReplyExecGroup
 - AdditionalCoordinatedRequestReplyExecGroup
 - Patterns
 - Samples
 - Service Virtualization
 - Policy
 - Configurable Services
 - WorkloadManagement
 - Data
 - Data Capture Stores
 - Security
 - Users
 - My Profile
 - Monitoring
 - Admin Log
 - My Workspace

Quick View (Right Panel):

Quick View	
Broker Name	IB9NODE
Version	9001
Admin Security	Off
Run Mode	running
Short Description	
Long Description	

Advanced Properties

View runtime statistics using the WebUI



- Using the WebUI in Integration Bus v9:
 - Control statistics at all levels
 - Easily view and compare flows, helping to understand which are processing the most messages or have the highest elapsed time
 - Easily view and compare nodes, helping to understand which have the highest CPU or elapsed times.
 - View all statistics metrics available for each flow
 - View historical flow data



Coordinated Request Reply MQ Application

Start

Stop

Statistics on (highlighted)

Statistics off

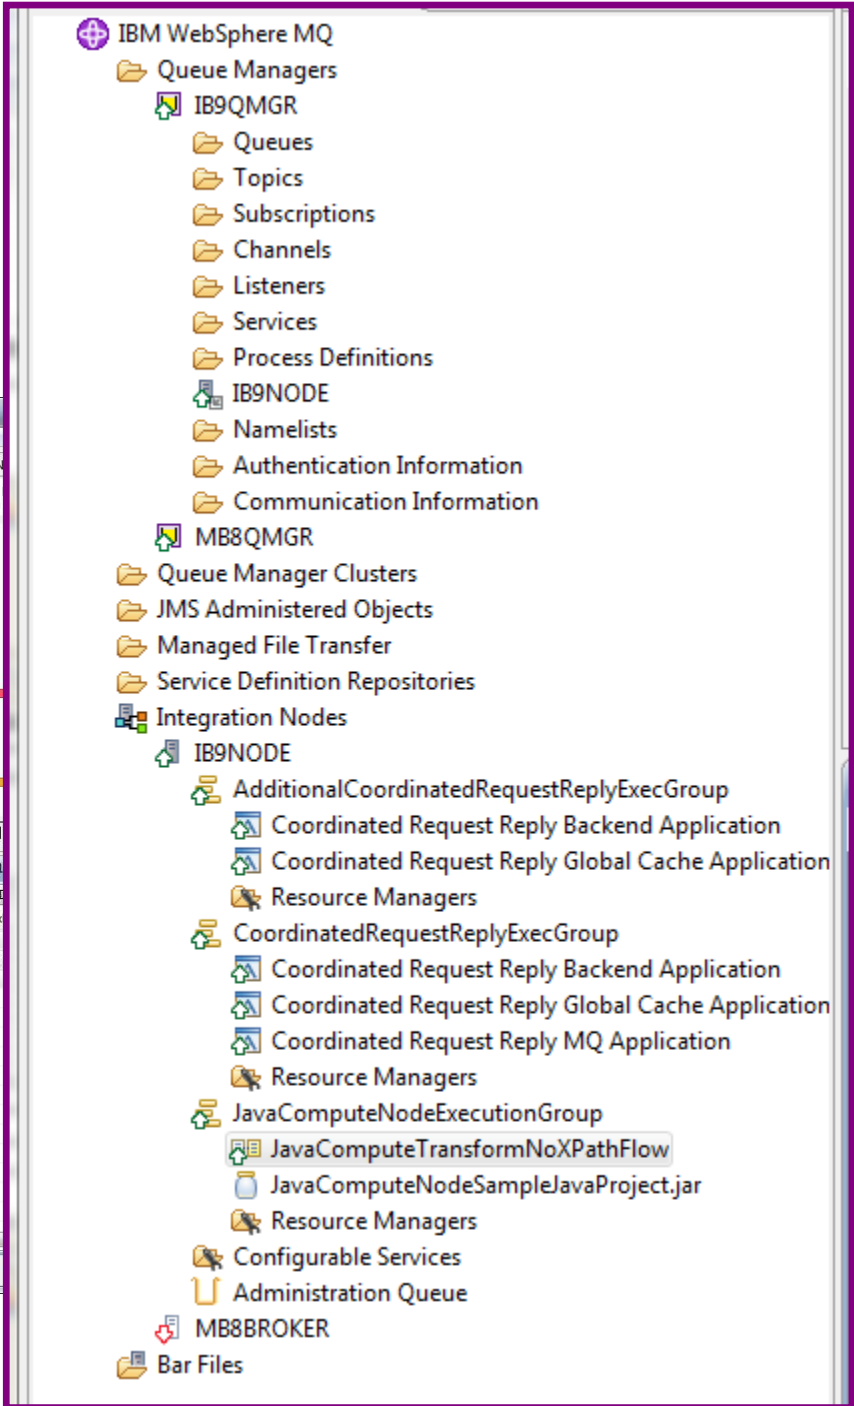
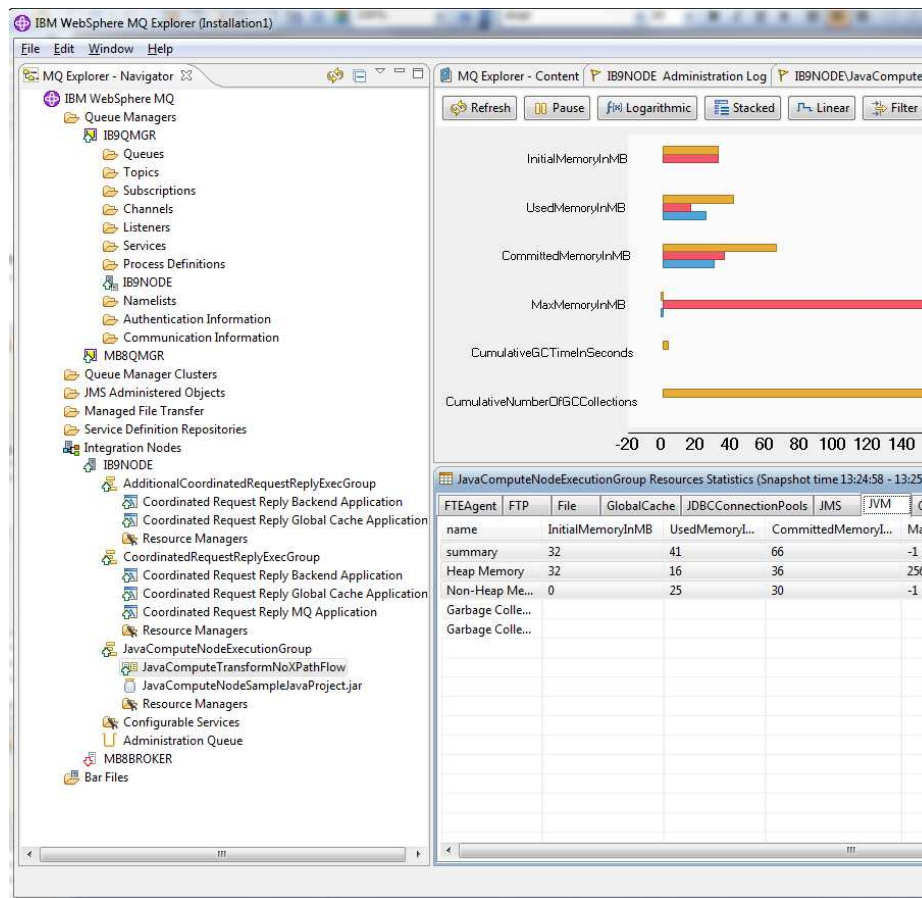
Statistics on (tooltip)

Throughput per message flow for last 15 seconds. Last updated at 09:50:29 GMT Daylight Time.

Flow name	Message Rate (messages/s)	Average Elapsed Time/Invocation (ms)	Average CPU Time/Invocation (ms)
Request	1.00	4.4	0.8
Reply	1.00	3.9	0.5
BackendReplyApp	1.00	1,001.7	1.1

Integration Explorer

- Fully functional administration tool based on MQ Explorer



Integration Explorer & Resource Statistics



- View resource statistics for resource managers in IIB such as JVM, ODBC, JDBC etc

The screenshot shows the IBM WebSphere MQ Explorer interface. On the left is a tree view of the system hierarchy. The main area displays resource statistics for a selected component. A bar chart visualizes memory usage (Initial, Used, Committed, Max) and garbage collection metrics. Below the chart is a table with the following data:

name	InitialMemoryInMB	UsedMemoryInMB	CommittedMemoryInMB	MaxMemoryInMB	CumulativeGCTimeInSeconds	CumulativeNumberOfGC
summary	32	41	66	-1	3	277
Heap Memory	32	16	36	256		
Non-Heap Me...	0	25	30	-1		
Garbage Colle...					2	262
Garbage Colle...					1	15

- Statistics
- User Trace All Flows
- Trace Nodes All Flows
- Service Trace
- Flow Debug Port
- Properties...

- View Message Flow Statistics
- Start Resource Statistics
- Stop Resource Statistics
- View Resource Statistics



Integration Explorer & Activity Log



- Start
- Stop
- Refresh
- Delete
- Open Activity Log**

- View activity as it happens using the Integration Explorer
- Filter by resource managers

MQ Explorer - Content | IB9NODE Administration Log | IB9NODE\JavaComputeNodeExecutionGroup\JavaComputeTransformNoXPathFlow - Activity Log

All Columns | Apply filter | Clear | All Threads | Select columns... | Previous | Next

1000 entries

Message ...	Timestamp ^	RM	MSGFLOW	Message Summary
i BIP11506I	15-Jul-2013 12:50:07.000...		JavaComputeTransform...	Committed a local transaction.
i BIP11501I	15-Jul-2013 12:50:07.000...		JavaComputeTransform...	Received data from input node 'JCTransformNoXPathInput'.
i BIP11506I	15-Jul-2013 12:50:08.000...		JavaComputeTransform...	Committed a local transaction.
i BIP11501I	15-Jul-2013 12:50:08.000...		JavaComputeTransform...	Received data from input node 'JCTransformNoXPathInput'.
i BIP11506I	15-Jul-2013 12:50:09.000...		JavaComputeTransform...	Committed a local transaction.
i BIP11501I	15-Jul-2013 12:50:09.000...		JavaComputeTransform...	Received data from input node 'JCTransformNoXPathInput'.
i BIP11506I	15-Jul-2013 12:50:10.000...		JavaComputeTransform...	Committed a local transaction.
i BIP11501I	15-Jul-2013 12:50:10.000...		JavaComputeTransform...	Received data from input node 'JCTransformNoXPathInput'.
i BIP11506I	15-Jul-2013 12:50:11.000...		JavaComputeTransform...	Committed a local transaction.
i BIP11501I	15-Jul-2013 12:50:11.000...		JavaComputeTransform...	Received data from input node 'JCTransformNoXPathInput'.
i BIP11506I	15-Jul-2013 12:50:13.000...		JavaComputeTransform...	Committed a local transaction.
i BIP11501I	15-Jul-2013 12:50:13.000...		JavaComputeTransform...	Received data from input node 'JCTransformNoXPathInput'.
i BIP11506I	15-Jul-2013 12:50:14.000...		JavaComputeTransform...	Committed a local transaction.
i BIP11501I	15-Jul-2013 12:50:14.000...		JavaComputeTransform...	Received data from input node 'JCTransformNoXPathInput'.
i BIP11506I	15-Jul-2013 12:50:15.000...		JavaComputeTransform...	Committed a local transaction.
i BIP11501I	15-Jul-2013 12:50:15.000...		JavaComputeTransform...	Received data from input node 'JCTransformNoXPathInput'.
i BIP11506I	15-Jul-2013 12:50:16.000...		JavaComputeTransform...	Committed a local transaction.
i BIP11501I	15-Jul-2013 12:50:16.000...		JavaComputeTransform...	Received data from input node 'JCTransformNoXPathInput'.
i BIP11506I	15-Jul-2013 12:50:17.000...		JavaComputeTransform...	Committed a local transaction.
i BIP11501I	15-Jul-2013 12:50:17.000...		JavaComputeTransform...	Received data from input node 'JCTransformNoXPathInput'.
i BIP11506I	15-Jul-2013 12:50:19.000...		JavaComputeTransform...	Committed a local transaction.
i BIP11501I	15-Jul-2013 12:50:19.000...		JavaComputeTransform...	Received data from input node 'JCTransformNoXPathInput'.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

Other forms of administration



- IIB can also be administered using:
 - Command line
 - REST interface
 - CMP Java API

The screenshot shows the IBM Integration Console 9.0.0.1 interface. The top window is a command prompt with the following text:

```
IBM Integration Console 9.0.0.1
MQSI 9.0.0.1
C:\Program Files\IBM\MQSI\9.0.0.1

C:\Program Files\IBM\MQSI\9.0.0.1>mqsilist
BIP1284I: Broker 'IB9NODE' on queue manager 'IB9QMGR' is running.
BIP8071I: Successful command completion.
C:\Program Files\IBM\MQSI\9.0.0.1>
```

Below the command prompt is a browser address bar showing the URL: localhost:4414/apiv1/executiongroups. Below the browser is a snippet of XML output:

```
<executionGroups type="executionGroup">
  <executionGroup type="executionGroup"
    propertiesUri="/apiv1/executiongroup">
  <executionGroup type="executionGroup"
    propertiesUri="/apiv1/executiongroup">
  <executionGroup type="executionGroup"
    propertiesUri="/apiv1/executiongroup">
  </executionGroups>
```

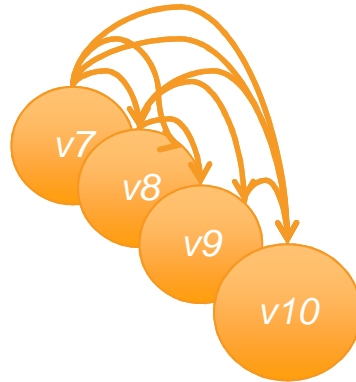


At the bottom is the IBM Integration API (CMP) Exerciser window. It has a table with columns 'Method' and 'Result'. Below the table is a log area with the following text:

```
15:39:00 This is a sample application that demonstrates the capabilities of the IBM Integration API, also
15:39:00 known as the CMP, which allows you to manage brokers programmatically.
15:39:00
15:39:00 Program started
15:39:00
```

The status bar at the bottom of the window says 'Disconnected.'




Some other useful features




Data Capture Store

Data viewer Replay list

DefaultCaptureStore  Mark for replay

Event time	Local Transaction ID	Pare Tran
<input checked="" type="checkbox"/> 2012-01-24 16:16:48.716	TESTDATA0123456789@!	CDA

GitHub Search or type a con



Chef cookbooks for ot4i
ot4i-cookbooks

Policies

Overview

WorkloadManagement

Values that you do not define on this page are inherited from t

Policy Name: BatchWor

Targets and Limits

Notification Threshold: 100

Maximum Rate: 300

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



Download Today!

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



Developer Edition



- Free edition of IB with all nodes available and no time limitations
- Throughput rate limited to 1TPS per integration flow
- Assistance through user community (e.g. mqseries.net)
 - No formal IBM support
- Simple to download, install and use
 - Single installation package contains ALL required software:
 - MQ 7.5, Integration Bus (Runtime, Toolkit, Explorer)
 - Available on Windows and Linux platforms



Complete your session evaluations online at www.SHARE.org/Seattle-Eval





Industry Solutions


Complete your session evaluations online at www.SHARE.org/Seattle-Eval



Industry Specific Solutions



- Industry Specific Packs providing relevant support for applications in certain industries
 - Build on IBM Integration Bus functionality
 - Contain message models / patterns / nodes
 - Web-based operational monitoring

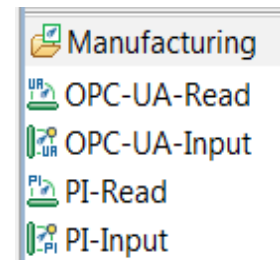


- Healthcare
 - HL7 Message Model
 - DICOM / ATNA / Medical Device Integration Nodes



- Retail
 - WebSphere Commerce and Sterling Order Management integration
 - TLog / POSLog / ARTS transformation and storage

- Manufacturing
 - OPC / PI / MQTT nodes



Complete your session evaluations online at www.SHARE.org/Seattle-Eval



Key Usage Scenarios

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

Top Integration Usage Patterns

- What are the top issues that people want to solve with integration solutions?

- **Extend the Reach** of Existing Applications
- Distribute **Database** information to where it's needed
- Create a **File Hub** to connect batch and online applications
- Get the most from **Packaged Applications**
- Take advantage of **.NET applications**
- Provide a Policy Enforcement Point for **Secure Connectivity**
- Extend Enterprise to **Devices and Mobile**
- **Monitor** your business activity and act intelligently
- Detect and Act Upon **Business Events and Rules**
- Provide Connectivity and Integration for **Business Processes**
- Make an inventory and enable **Policy based management**



Common
usage
patterns

Emerging
usage
patterns

- New usage patterns are continually emerging as business needs evolve!


IBM Integration Bus



- **IBM's Strategic Integration Technology**

- **Universal Connectivity FROM anywhere, TO anywhere**

- Simplify application connectivity for a flexible & dynamic infrastructure

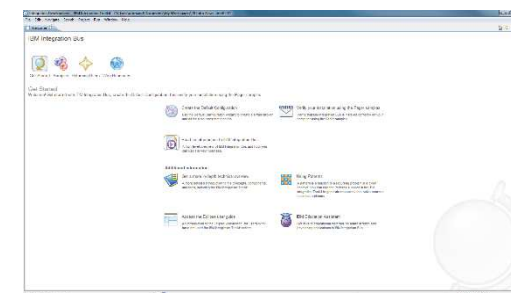


- **Protocols, Transports, Data Formats & Processing**

- Supports a wide range of built-in transports, protocols & systems
 - MQ, JMS 1.1, HTTP(S), SOAP, REST, File (incl. FTP & FTE), Database, TCP/IP, MQTT...
 - CICS, IMS, SAP, SEBL, PeopleSoft, JDEdwards, SCA, CORBA, email...
- Supports a broad range of data formats
 - Binary (C/COBOL), XML, CSV, JSON, Industry (SWIFT, EDI, HL7...), IDOCs, User Defined
- Message Processors
 - Route, Filter, Transform, Enrich, Monitor, Publish, Decompose, Sequence, Correlate, Detect...

- **Simple Programming with Patterns & Graphical Data Flows**

- Patterns for top-down, parameterized connectivity of common use cases
- Graphical data flows represent application & service connectivity
 - Custom logic via Graphical mapping, PHP, Java, ESQL, XSL & .NET
- Services, Applications and Libraries for structure and reuse



- **Extensive Management, Performance & Scalability**

- Extensive Administration & Systems Management facilities for developed solutions
- Wide range of operating system and hardware platforms supported, including virtual & cloud options
- High performance transactional processing, additional vertical & horizontal scalability
- Deployment options include Trial, Express, Standard and Advanced

- **Industry Packs for Industry Specific Content**



This was Session # 17041. The rest of the week



	Monday	Tuesday	Wednesday	Thursday	Friday
08:30			17060: Understanding MQ Deployment Choices and Use Cases	17051: Application Programming with MQ Verbs [z/OS & Distributed]	16544: Why Shouldn't I Be Able To Open This Queue? MQ and CICS Security Topics Room: Willow B
10:00	17036: Introduction to MQ - Can MQ Really Make My Life Easier? [z/OS & Distributed]		17052: MQ Beyond the Basics - Advanced API and Internals Overview [z/OS & Distributed]	17054: Nobody Uses Files Any More do They? New Technologies for Old Technology, File Processing in MQ MFT and IIB [z/OS & Distributed]	17057: Not Just Migrating, but Picking up New Enhancements as You Go - We've Given You the Shotgun, You Know Where Your Feet Are [z/OS & Distributed]
			17035: MQ for z/OS, Using and Abusing New Hardware and the New V8 Features [z/OS] Room: Willow B		
11:15	17041: First Steps with IBM Integration Bus: Application Integration in the New World [z/OS & Distributed]		16732: MQ V8 Hands- on Labs! MQ V8 with CICS and COBOL! MQ SMF Labs! Room: Redwood	17046: Paging Dr. MQ - Health Check Your Queue Managers to Ensure They Won't Be Calling in Sick! [z/OS]	17053: MQ & DB2 – MQ Verbs in DB2 & InfoSphere Data Replication (Q Replication) Performance [z/OS]
01:45	17037: All About the New MQ V8 [z/OS & Distributed]	17034: MQ Security: New V8 Features Deep Dive [z/OS & Distributed]	17040: Using IBM WebSphere Application Server and IBM MQ Together [z/OS & Distributed]	17062: End to End Security of My Queue Manager on z/OS [z/OS]	All sessions in Seneca unless otherwise noted.
03:15	17042: What's New in IBM Integration Bus [z/OS & Distributed]	17065: Under the hood of IBM Integration Bus on z/OS - WLM, SMF, AT-TLS, and more [z/OS]	17043: The Do's and Don'ts of IBM Integration Bus Performance [z/OS & Distributed]	17039: Clustering Queue Managers - Making Life Easier by Automating Administration and Scaling for Performance [z/OS & Distributed]	
04:30	17059: IBM MQ: Are z/OS & Distributed Platforms like Oil & Water? [z/OS & Distributed]	17055: What's the Cloud Going to Do to My MQ Network?	17044: But Wait, There's More MQ SMF Data Now?!?! - Monitoring your Channels Using V8's New Chinit SMF Data [z/OS]	17068: Monitoring and Auditing MQ [z/OS & Distributed]	



Questions?



Complete your session evaluations online at www.SHARE.org/Seattle-Eval



Trademark Statement



- IBM and the IBM logo are trademarks of International Business Machines Corporation, registered in many jurisdictions. Other marks may be trademarks or registered trademarks of their respective owners.
- Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
- Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Netezza® is a trademark or registered trademark of IBM International Group B.V., an IBM Company.
- Worklight® is a trademark or registered trademark of Worklight, an IBM Company.
- Other company, product and service names may be trademarks, registered marks or service marks of their respective owners.
- References in this publication to IBM products and services do not imply that IBM intends to make them available in all countries in which IBM operates.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

