

Nobody Uses Files Any More do They? New Technologies for Old Technology, File Processing in MQ MFT and IIB

Dave Gorman : IBM Integration Bus Performance

Chris Leonard : MQ Level 3 Service

Session 17054



#SHAREorg



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



How do most organizations move files today?

Most organizations rely on a mix of homegrown code, several legacy products and different technologies ... and even people!

▪ FTP

- Typically File Transfer Protocol (FTP) is combined with writing and maintaining homegrown code to address its limitations

▪ Why is FTP use so widespread?

- FTP is widely available – Lowest common denominator
- Promises a quick fix – repent at leisure
- Simple concepts – low technical skills needed to get started
- FTP products seem “free”, simple, intuitive and ubiquitous

▪ Legacy File Transfer products

- A combination of products often used to provide silo solutions
- Often based on proprietary versions of FTP protocol
- Can't transport other forms of data besides files
- Usually well integrated with B2B but rarely able to work with the rest of the IT infrastructure – especially with SOA

▪ People

- From IT Staff to Business staff and even Security Personnel
- Using a combination of email, fax, phone, mail, memory keys...

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

Shortcomings of Basic FTP

Limited Reliability

- ❑ Unreliable delivery – Lacking checkpoint restart – Files can be lost
- ❑ Transfers can terminate without notification or any record – corrupt or partial files can be accidentally used
- ❑ File data can be unusable after transfer – lack of Character Set conversion



Limited Security

- ❑ Often usernames and passwords are sent with file – as plain text!
- ❑ Privacy, authentication and encryption often not be available
- ❑ Non-repudiation often lacking

Limited Flexibility

- ❑ Changes to file transfers often require updates to many ftp scripts that are typically scattered across machines and require platform-specific skills to alter
- ❑ All resources usually have to be available concurrently
- ❑ Often only one ftp transfer can run at a time
- ❑ Typically transfers cannot be prioritized

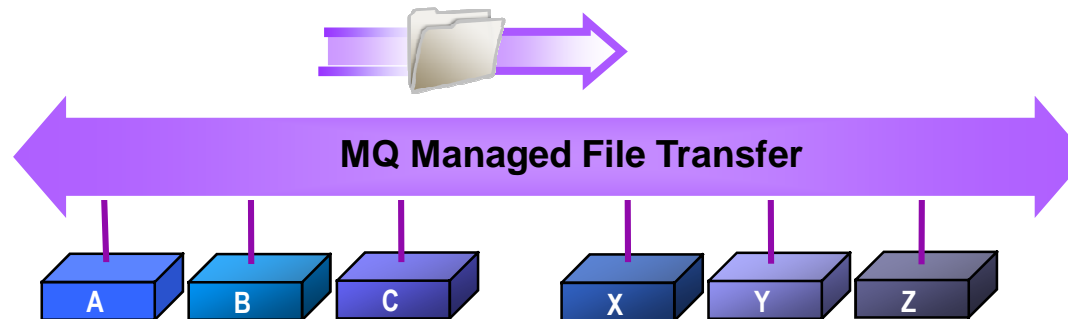


Limited visibility and traceability

- ❑ Transfers cannot be monitored and managed centrally or remotely
- ❑ Logging capabilities may be limited and may only record transfers between directly connected systems
- ❑ Cannot track the entire journey of files – not just from one machine to the next but from the start of its journey to its final destination



What is MQ Managed File Transfer?





- Auditable Full logging and auditing of file transfers + archive audit data to a database
- Reliable Checkpoint restart. Exploits solid reliability of MQ
- Secure Protects file data in transit using SSL. Provides end-to-end encryption using AMS
- Automated Providing scheduling and file watching capabilities for event-driven transfers
- Centralized Provides centralized monitoring and deployment of file transfer activities
- Any file size Efficiently handles anything from bytes to terabytes
- Integrated Integrates with MB, WSRR, ITCAMs for Apps, DataPower + Connect:Direct
- Cost Effective Reuses investment in MQ. Wide range of support (inc. z/OS and IBM i)

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

A consolidated transport for both files and messages

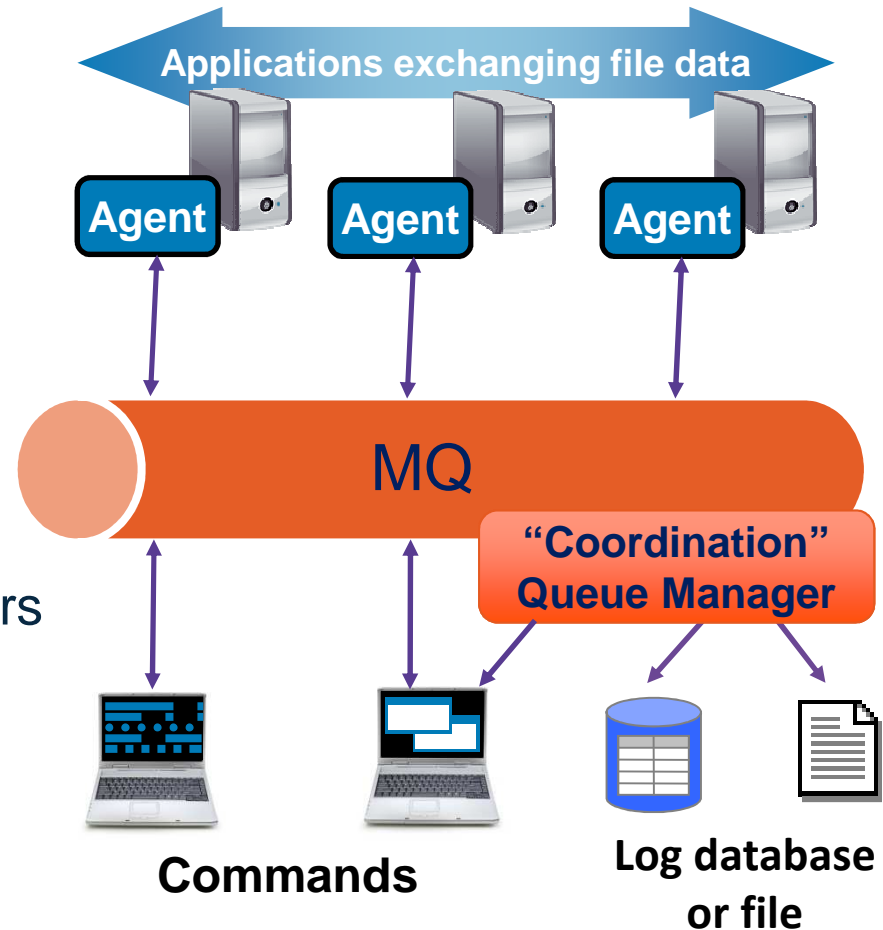
- Traditional approaches to file transfer result in parallel infrastructures
 - One for files – typically built on FTP
 - One for application messaging – based on MQ, or similar
- High degree of duplication in creating and maintaining the two infrastructures
- Managed File Transfer reuses the MQ network for managed file transfer and yields:
 - Operational savings and simplification
 - Reduced administration effort
 - Reduced skills requirements and maintenance



Components of a typical MQ MFT network

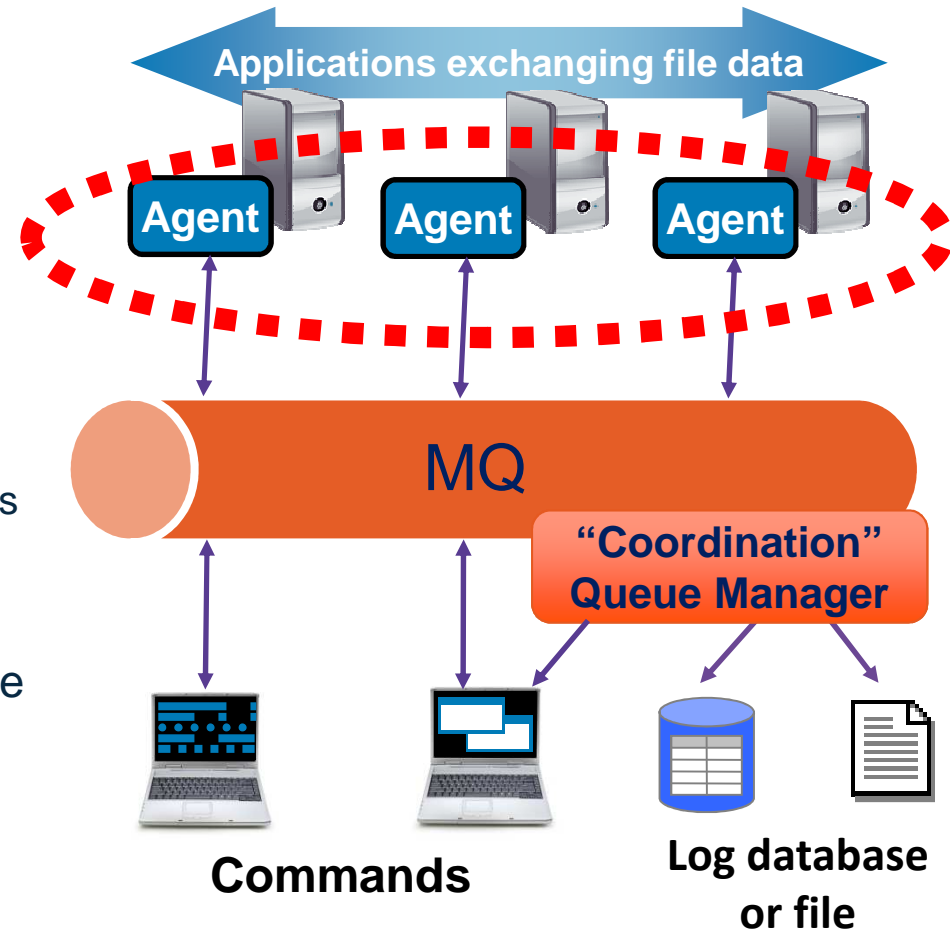


- Agents
 - The endpoints for managed file transfer operations
- Commands
 - Send instructions to agents
- Log database or file
 - A historical record of file transfers
- Coordination queue manager
 - Gathers together file transfer events



Agents

- Act as the end points for file transfers
- Long running MQ applications that transfer files by splitting them into MQ messages
 - Efficient transfer protocol avoids excessive use of MQ log space or messages building up on queues
- Multi-threaded file transfers
 - Can both send and receive multiple files at the same time
- Generate a log of file transfer activities which is sent to the “coordination queue manager”
 - This can be used for audit purposes
- Associated with one particular queue manager (any supported version)
 - Agent state on queues





Notes on: Agents



NOTES

- MFT agent processes define the end-points for file transfer. That is to say that if you want to move files off a machine, or onto a machine – that machine would typically need to be running an agent process
- Agent processes are long running MQ applications that oversee the process of moving file data in a managed way. Each agent monitors a ‘command’ queue waiting for messages which instruct it to carry out work, for example file transfers

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





Notes on: Agents



NOTES

- The MFT agent process needs connectivity to an MQ queue manager to do useful work. It can connect either directly to a queue manager running on the same system, or as an MQ client using an embedded version of the MQ client library (which is kept completely separate to any other MQ client libraries that may or may not already have been installed onto the system)*
 - Each agent requires its own set of MQ queues – which means that an agent is tied to the queue manager where these queues are defined
 - However – one queue manager can support multiple agents

* Note: availability of direct (bindings) connectivity or MQ client based connectivity is dependent on the version of MQ MFT in use

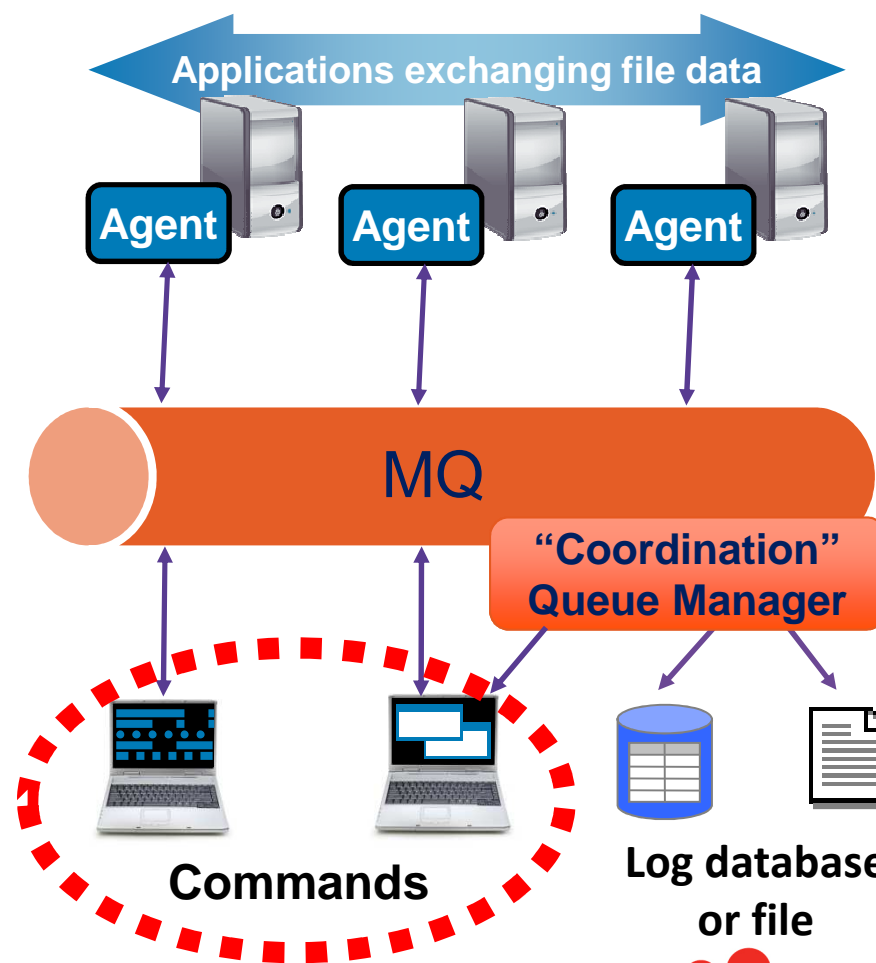
- MQ Managed File Transfer on z/OS does not support the MQ client style of connectivity
- Managed File Transfer on distributed platforms has a 'server' and 'client' offering. The agent component of the 'client' offering is restricted to only supporting MQ client style connectivity. The agent component of the 'server' offering may be used either connectivity options

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



Commands

- Send instructions to agents and display information about agent configuration
 - Via MQ messages
- Many implementations of commands:
 - MQ Explorer plug-in
 - Command line programs
 - Open scripting language
 - JCL
 - Documented interface to program to





Notes on: Commands

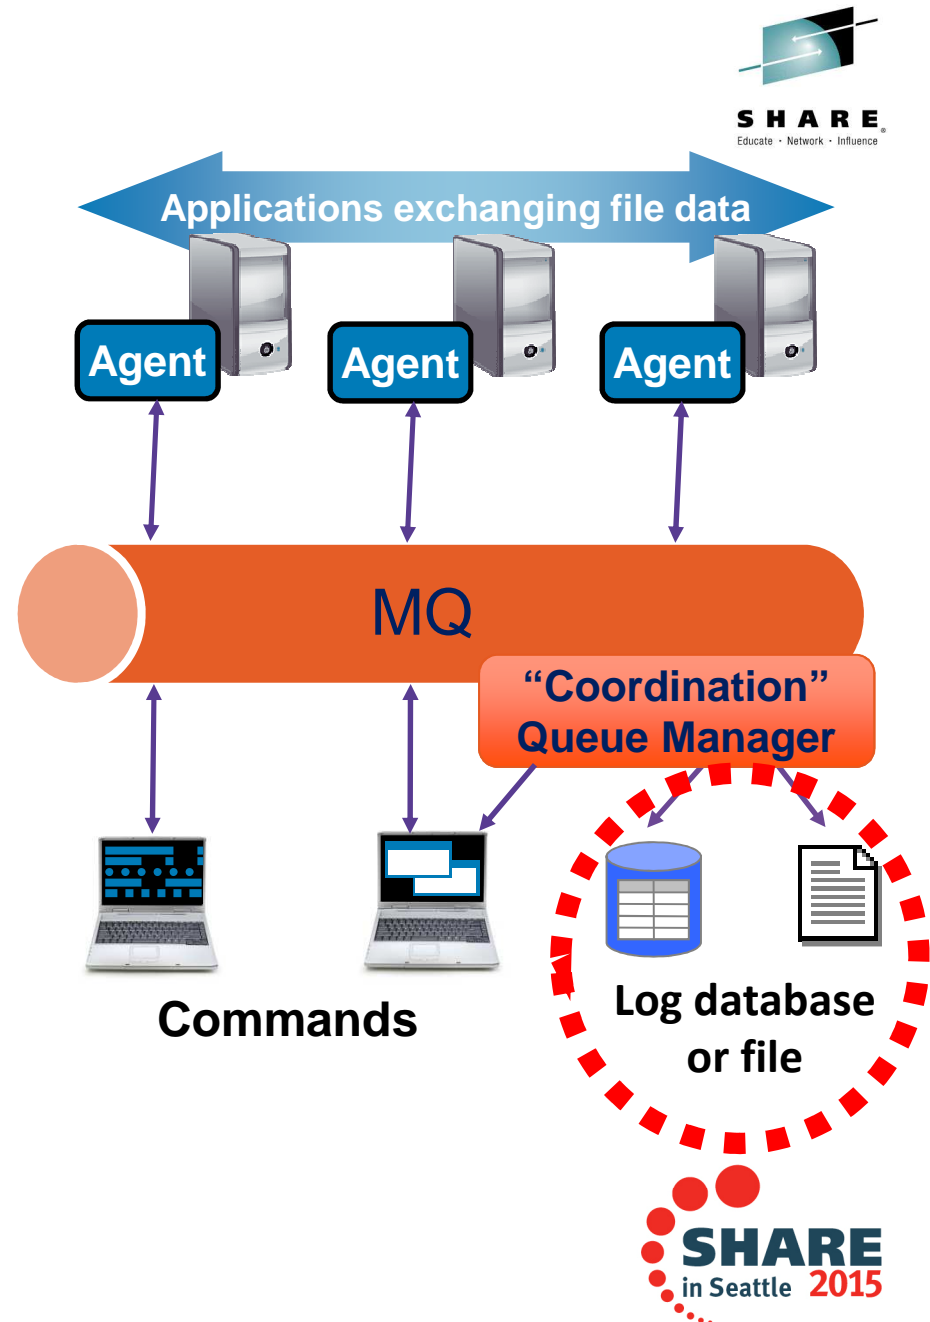


NOTES

- “Commands” is the name we have given to anything which instructs the MFT agent. As described on the previous slide, there are a wide range of command implementations including graphical and non-graphical command-line based commands
- Commands instruct the MFT agent by sending it messages. The messages themselves use a documented format which can easily be incorporated into your own applications
- The commands that are supplied with MFT can connect either as an MQ client (again based on embedded client libraries) or directly to a queue manager located on the same system

Log Database & File

- Keeps a historical account of transfers that have taken place
 - Who, where, when... etc.
- Implemented by the 'logger' component which connects to the coordination queue manager
 - Stand alone application
 - Can log to database or file
 - Or JEE application
 - Can log to database only





Notes on: Log Database



NOTES

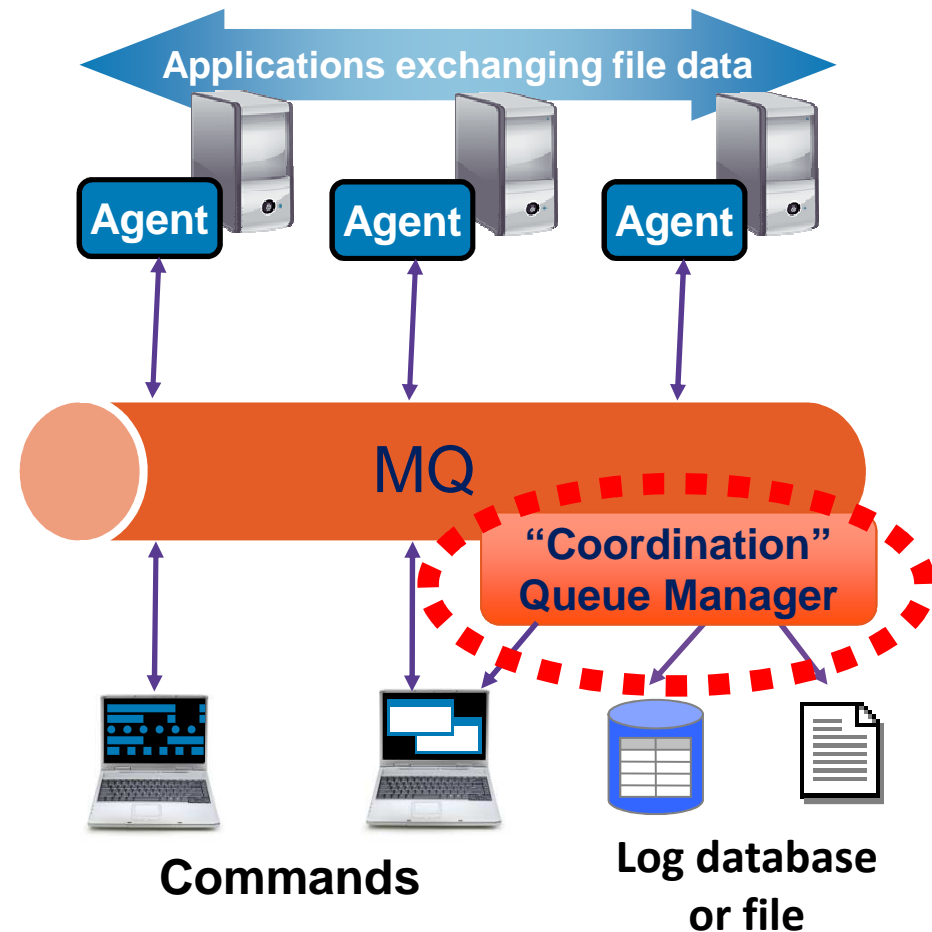
- Managed File Transfer can record a historical account of file transfers to a database using the 'database logger' component.
 - This component is available to run as a stand alone process or as a JEE application
- The information used to populate the log database is generated, as MQ messages, by the MFT agents participating in file transfers. This is routed to a collection point in the MQ network, referred to as the 'coordination queue manager' (see next slides). The database logger component subscribes to the messages produced by agents and reliably enters them into a database.

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



Coordination Queue Manager

- Gathers together information about events in the file transfer network
- Not a single point of failure
 - Can be made highly available
 - Messages stored + forwarded
- MQ v7 publish / subscribe
 - Allows multiple log databases, command installs
 - Documented interface





Notes on: Coordination queue manager



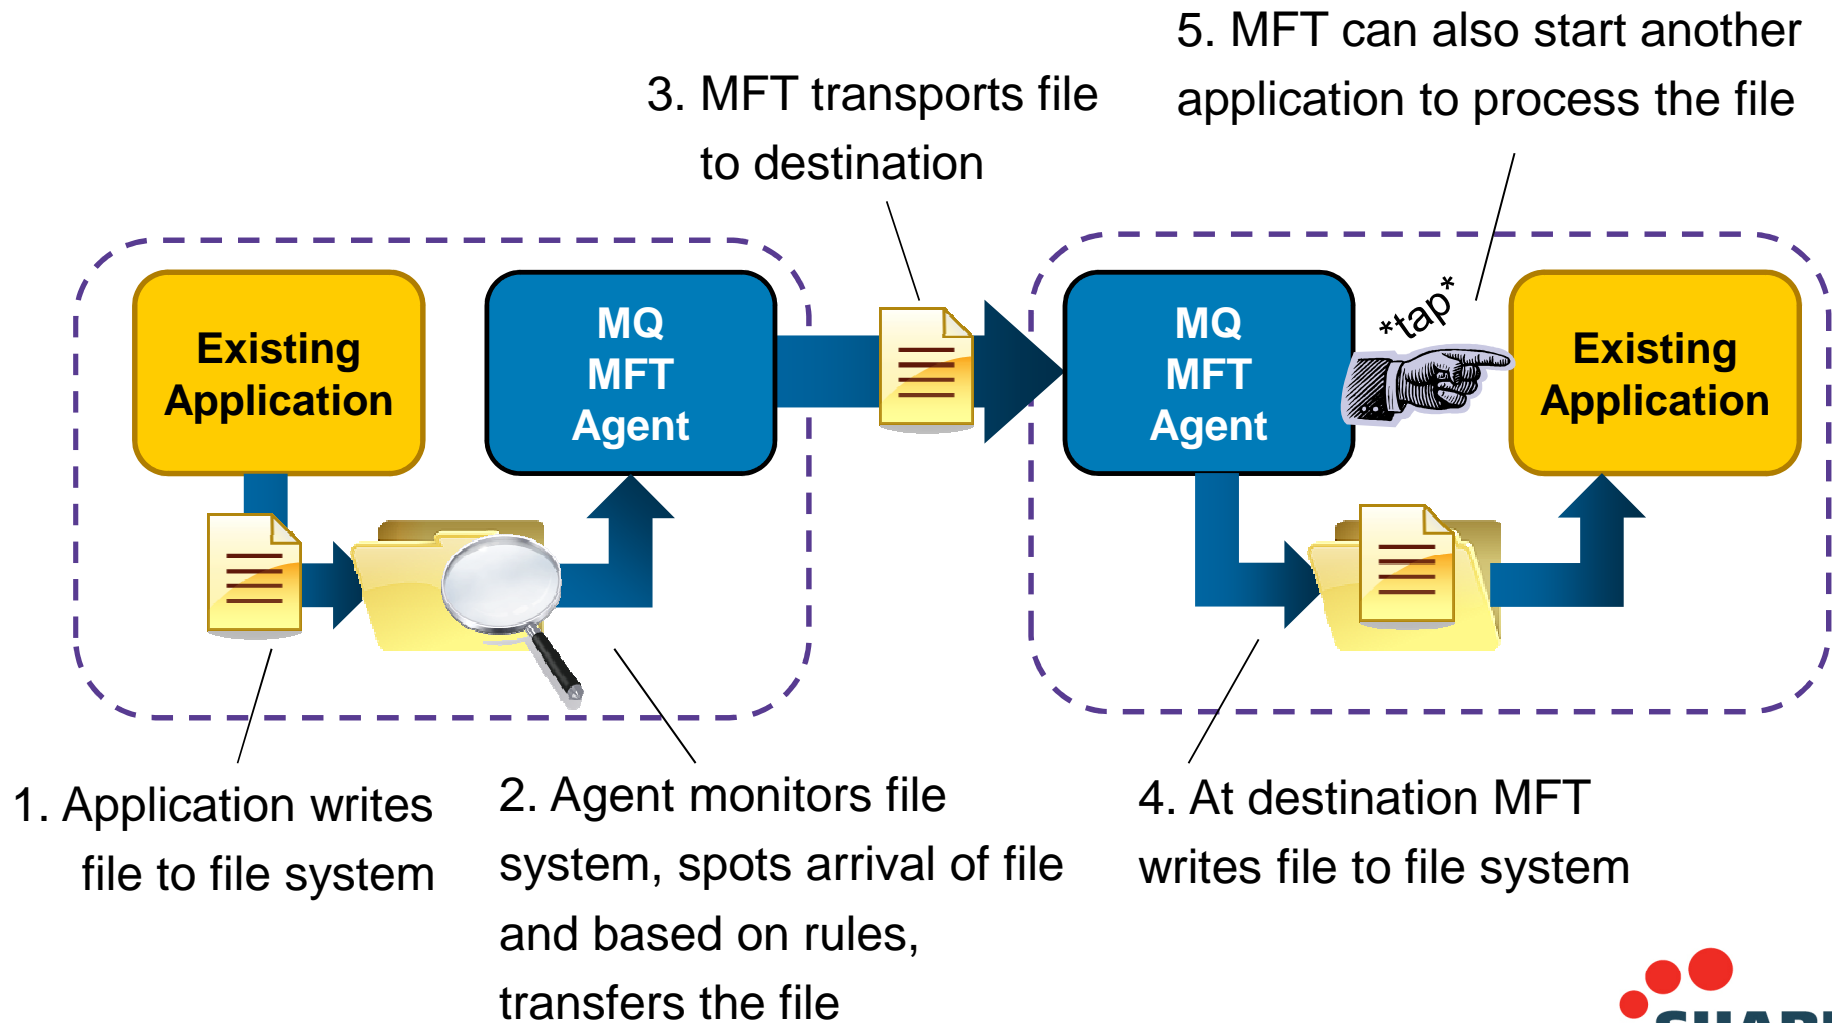
NOTES

- The coordination queue manager is used as the gathering point for all the information about file transfers taking place between a collection of MFT agents
- The queue manager uses publish/subscribe (so it must be MQ version 7) to distribute this information to “interested parties” which typically include:
 - The MQ Explorer plug-in, which provides a graphical overview of MFT activity
 - The database logger component, which archives the information to a database
 - Some of the command line utilities which are part of the MFT product
- The format used to publish information is documented and can be used to develop 3rd party applications which process this data
- Although there is only a single ‘coordination’ queue manager for a given collection of agents it does not represent a point of failure:
 - MQ stores and forwards messages to the coordination queue manager when it is available – so if the coordination queue manager is temporarily unavailable no log data is lost
 - The ‘coordination’ queue manager can be made highly available using standard HA techniques such as MQ multi-instance queue manager or via a HA product such as PowerHA

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



Example usage of monitoring & program execution





Notes on: Monitoring & Program Execution

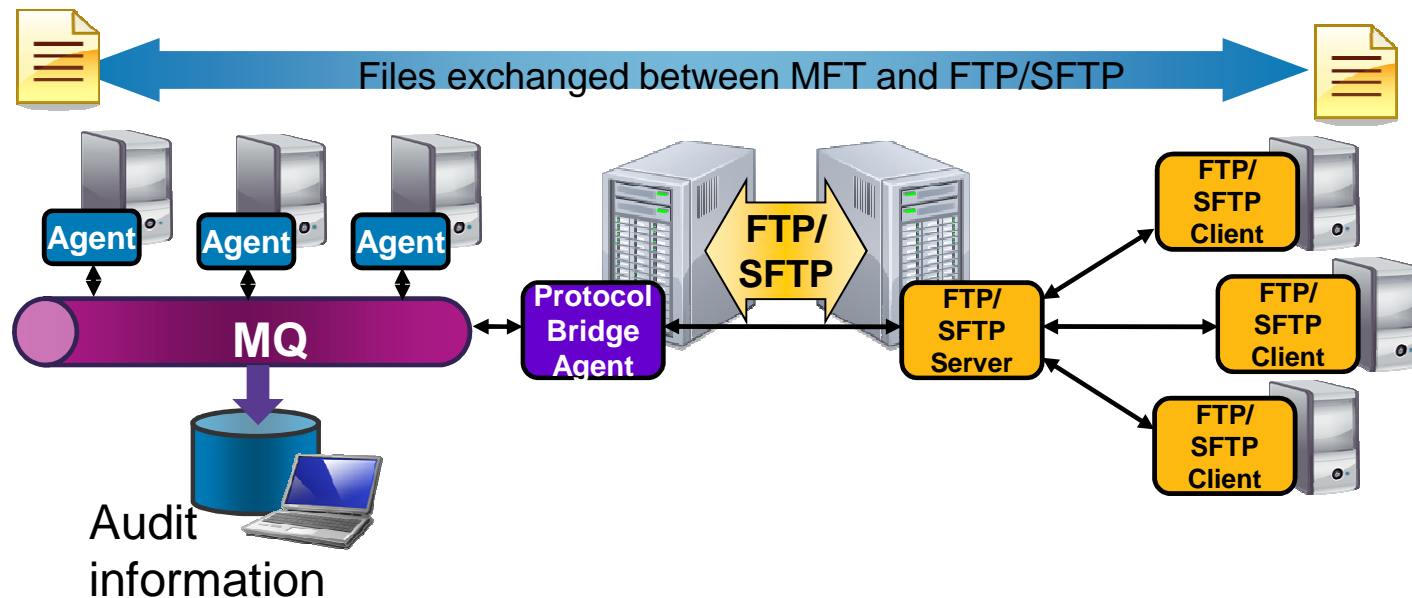


NOTES

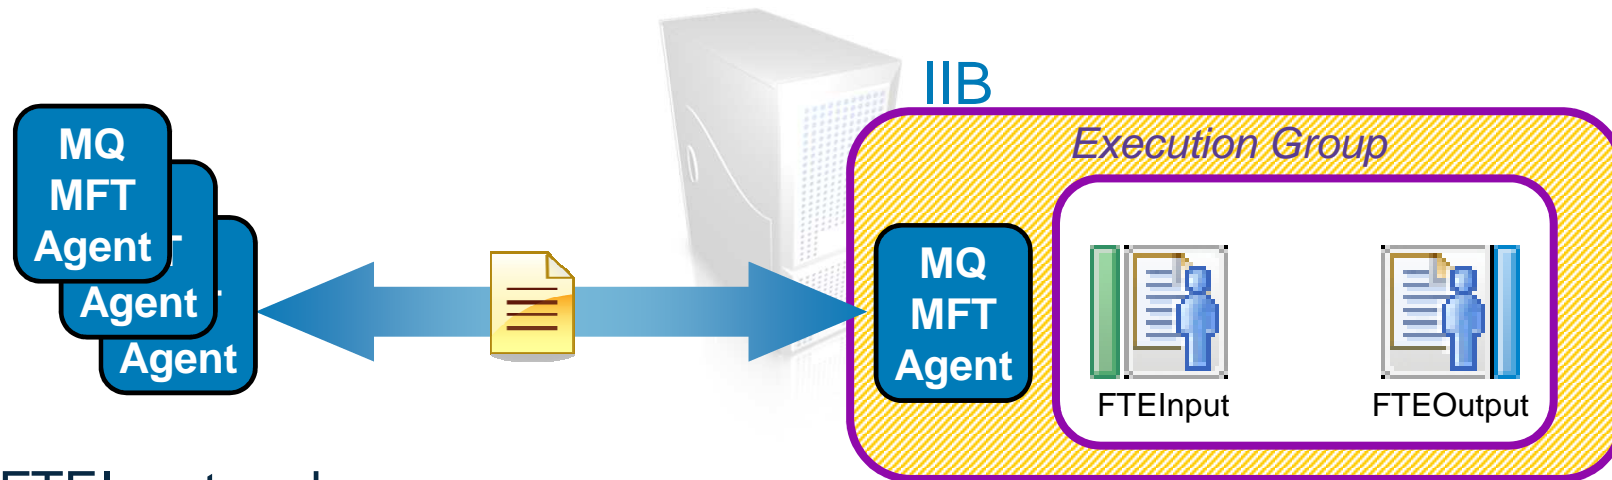
- Resource monitors work in two stages:
 1. Poll a resource (in this case the file system – but as we’ll see later the ‘resource’ can also be an MQ queue) and identify that a condition has been met (perhaps the appearance of a file matching a particular pattern)
 2. Perform an action (which can include starting a managed file transfer, or running a script), optionally propagating information about the resource (for example the name of the file triggered on) into the action
- As shown on the previous slide, resource monitors are typically used to provide integration with an existing system without needing to make changes to the system
- Another function of MFT, used for integration with existing systems is the ability to execute programs or scripts both on the source or destination systems for a file transfer. This can be used to:
 - Start a program, on the source system, which generates the file data to be transferred prior to performing the managed file transfer
 - Start, or notify, a program on the destination system when the file data has been transferred – allowing it to process the data without having to poll

Protocol bridge agents

- Support for transferring files located on FTP and SFTP servers
 - The source or destination for a transfer can be an FTP or an SFTP server
- Enables incremental modernization of FTP-based home-grown solutions
 - Provides auditability of transfers across FTP/SFTP to central audit log
- Ensures reliability of transfers across FTP/SFTP with checkpoint restart
 - Fully integrated into graphical, command line and XML scripting interfaces
 - Just looks like another MFT agent...



IBM Integration Bus Nodes

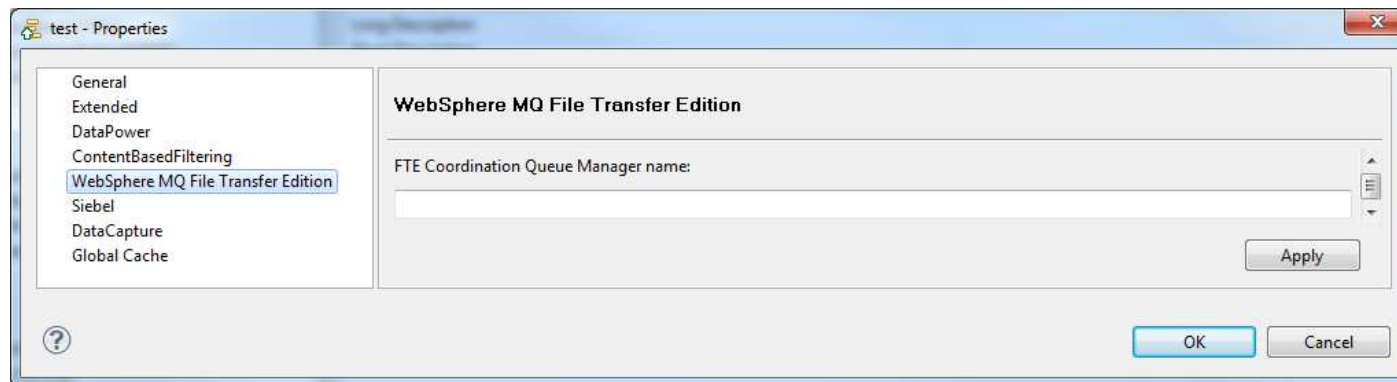


- FTEInput node
 - Build flows that accept file transfers from the MQ MFT network
- FTEOutput node
 - Build flows that are designed to send a file across an MQ MFT network
- When MQ MFT nodes are used in a flow an MFT agent is automatically started in the IIB integration server

Creating and running an agent in Integration Bus



- Install
 - FTE code is installed as part of Integration Bus install. No need for separate install
- Deployment
 - Agents run in the integration server JVMs. One agent per integration server
 - Coordinating queue manager defined as an integration server property



- Agent name is derived from integration node and integration server name:
 - <integration node name>.<integration server name>
 - FTE Agents are created automatically by integration node including required queues
- Starting/stopping
 - Agent started when first FTE flow node using it is deployed
 - Agent stopped when last FTE flow node using it is un-deployed or stopped

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





Notes on: Creating and running an agent in IIB



NOTES

The core FTE product is installed as part of the normal IIB install. No additional components need to be installed apart from the MQExplorer add on which can be used to monitor transfers. It is not required for the transfers to work but is useful for monitoring what is happening with transfers.

Each integration server can be configured to run an FTE agent. This is done by setting the coordination queue manager property either using MBExplorer or the mqsichangeproperties command.

The agent name is derived by concatenating the integration node and integration server name together. If the name is too long to be a valid agent name then it is truncated. If it contains non-valid characters (any characters not supported in MQ queue names) an error is written to the local system log.

All the configuration for the agent is created when the coordination queue manager is set and also deleted when it is unset.

As well as creating the required config files, all the required queues are also created.

The actual agent is started when the first FTE flow node using it is started and stopped when the last FTE flow node using it is stopped.

If a node is deployed without setting the coordination queue manager then an agent will be created at deployment time using the integration node queue manager as the coordination queue manager. This is a temporary agent which is deleted when the last FTE flow node using it has been stopped. It is recommended that the integration server property is set even when the integration node queue manager is the coordination queue manager.



Notes on: Administering an agent in IIB



NOTES

The administration of the agent is very simple with almost all tasks being done automatically by the integration server.

All the config files created are identical to those used in a standalone agent.

Logging by the agent is written to the standard integration node user and service trace but the log file in the config directory is still written to as well.


A location under the workpath is used as the default transfer directory and also to stage files before they are transferred. To separate files being transferred by different nodes a directory structure is created which includes the flow name and node name.

As with the file nodes a mqsitransmit directory is used to build files up before transferring them. The names of files in these directories are mangled to stop clashes of files with the same name which are going to be transferred to different agents or different directories. It is recommended not to delete files from the transmit directory unless the message flow using them has been stopped and all files are deleted.

The FTE MQExplorer plugin is very useful to monitor the progress of any transfers which have been done.


Accessing agent from a message flow

- Messages received and sent from agent using the following flow nodes:
 - FTEInput node: receives any file transferred to integration server agent
 - FTEOutput node: constructs a file and sends a request to the integration server agent to transfer a file

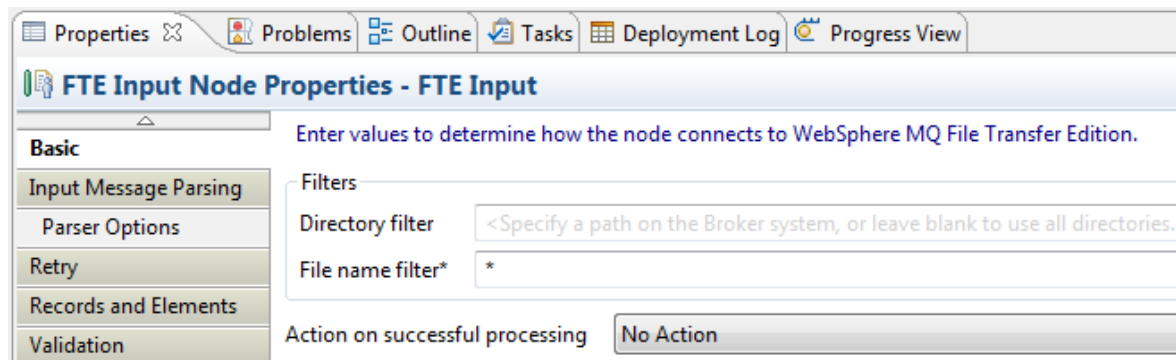


- FTE nodes based on the existing file nodes allowing:
 - Record based processing
 - Stream parsing for large file support
 - File filtering
 - Can be used with all other flow nodes
- Agent is started or stopped based on whether any FTE nodes are running

FTE Input Node



- Consistent with file input node but makes full use of the power of MQ MFT
- Timely
 - The FTE node is notified by the FTE agent when an inbound transfer is complete
 - The node processes the files in the transfer immediately.
 - Each file is processed independently
 - Can leave file unchanged after processing and just delete notification message
- Metadata
 - Metadata associated with the transfer is sent with the notification
 - Includes user defined metadata
- Filter
 - The node by default receives all transfers
 - Can specify which files to receive using a filter



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





Notes on: FTEInput Node



NOTES

The FTEInput node has all the core function of the standard File input node but has been enhanced to make use of the powerful function provided by MQ MFT.

It does not require any polling mechanism to scan directories because it is directly trigger by the embedded agent when a file has arrived and the transfer is complete.

It processes each file in a transfer separately and can process each file in parallel.

As well as receiving the data from the transfer it also receives all the meta data associated with the transfer. This includes lots of information from MQ MFT but also can include user defined data.

Each FTE Input node can specify a filter of which files it wants to process. By default it processes all files. If two nodes both have a filter which matches a file then only one will get given it to process. By default, if no filter expression is given, then the node we accept any transferred file.

The file name filter accepts wild cards but the directory can either be blank (accept any files) or a string that must match exactly. Relative paths are allowed which are taken relative to the default transfer directory.

FTE Input Node – metadata




LocalEnvironment	
FTE	
Transfer	
Directory	C:\\temp\\perf1
JobName	
Name	test.txt
LastModified	2010-07-05 13:37:55.796
SourceAgent	MB7BROKER.FTESAMPLEIN
DestinationAgent	MB7BROKER.FTESAMPLEIN
OriginatingHost	jreeve6
TransferId	414d51204d4237514d47522020202045c1314c2000780c
MQMDUser	SYSTEM
OriginatingUser	SYSTEM
TransferMode	binary
TransferStatus	0
FileSize	364
ChecksumMethod	MD5
Checksum	b0b492e8f7385747f3335276615ec5c0
DestinationAgentQmgr	MB7QMGR
SourceAgentQmgr	MB7QMGR
OverallTransferStatus	0
TotalTransfers	1
TransferNumber	1
UserDefined	
com.ibm.jreeve.Test1	Value 1
SourceAgent	MB7BROKER.FTESAMPLEIN
TimeStamp	20100705_123755_796108
Offset	0
Record	1
Delimiter	
IsEmpty	false

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



FTE Output Node

- Consistent with file output node but makes full use of the power of MQ MFT
- Transfer details
 - The destination agent, directory etc are defined on the node



Properties Problems Outline Tasks Deployment Log Progress View

FTE Output Node Properties - FTE Output

Settings for working with WebSphere MQ File Transfer Edition.

Description

Basic

Request

Records and Elements

Validation

Monitoring

Metadata

Job name <If you leave the job name blank, the job name does not appear in the transfer log.>

Destination

Agent <If you leave the agent blank, the agent embedded in the execution group will be used.>

Queue manager <If you leave the queue manager blank, the broker's queue manager will be used.>

File directory <Leave blank to use the agent's default directory.>

File name* *.txt
For example: test.txt or *.txt

Options

Mode Binary transfer (no conversion)

Disable computation of MD5 check sum

Overwrite files on destination system

- All details can be overridden using Local environment
- Support also for wild card file names
- Staging
 - Uses a local staging directory to build up a file record by record for transfer
 - Once a file is finished a request is sent to the FTE Agent to transfer the file
- Metadata
 - User data – if provided in the Local Environment - is sent with the transfer
 - Other metadata is generated by the FTE Agent or by the integration node.
 - Ant scripts – give details of ant scripts to run on remote agent

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





Notes on: FTEOutput Node



NOTES

The FTEOutput node makes use of the core function provided by File nodes.

The main properties on the node are details of where to transfer the files to. These can all be overridden using the local environment.

The main difference with an FTEOutput node is that the destination the file is to be written to is not local to the integration node file system but instead is a file system on a remote agent. The node first writes it to a staging directory on the local file system and then sends a request to the embedded agent to transfer it.

Files are built up in the mqsitransmit directory just like with a file node before being moved to the final file name once the file is complete and the transfer request is sent to the FTE agent. The name in the staging area is the same as the name the file will be transferred to unless a file with that name already exists on the local file system. If it exists than a number is appended to the file. The file name it is transferred to on the remote system is not effected by this and will not have the number appended.

FTE Output Node – Overrides and metadata



Overrides:

LocalEnvironment	
Destination	
FTE	
Directory	c:\\temp\\perf1
Name	test.txt
DestinationQmgr	MB7QMGR
DestinationAgent	MB7BROKER.FTESAMPLEIN
JobName	Job 1
Overwrite	true
UserDefined	
com.ibm.jreeve.Test1	Value 1
PreDestinationCall	
Name	john1.xml
PostDestinationCall	
Name	john2.xml

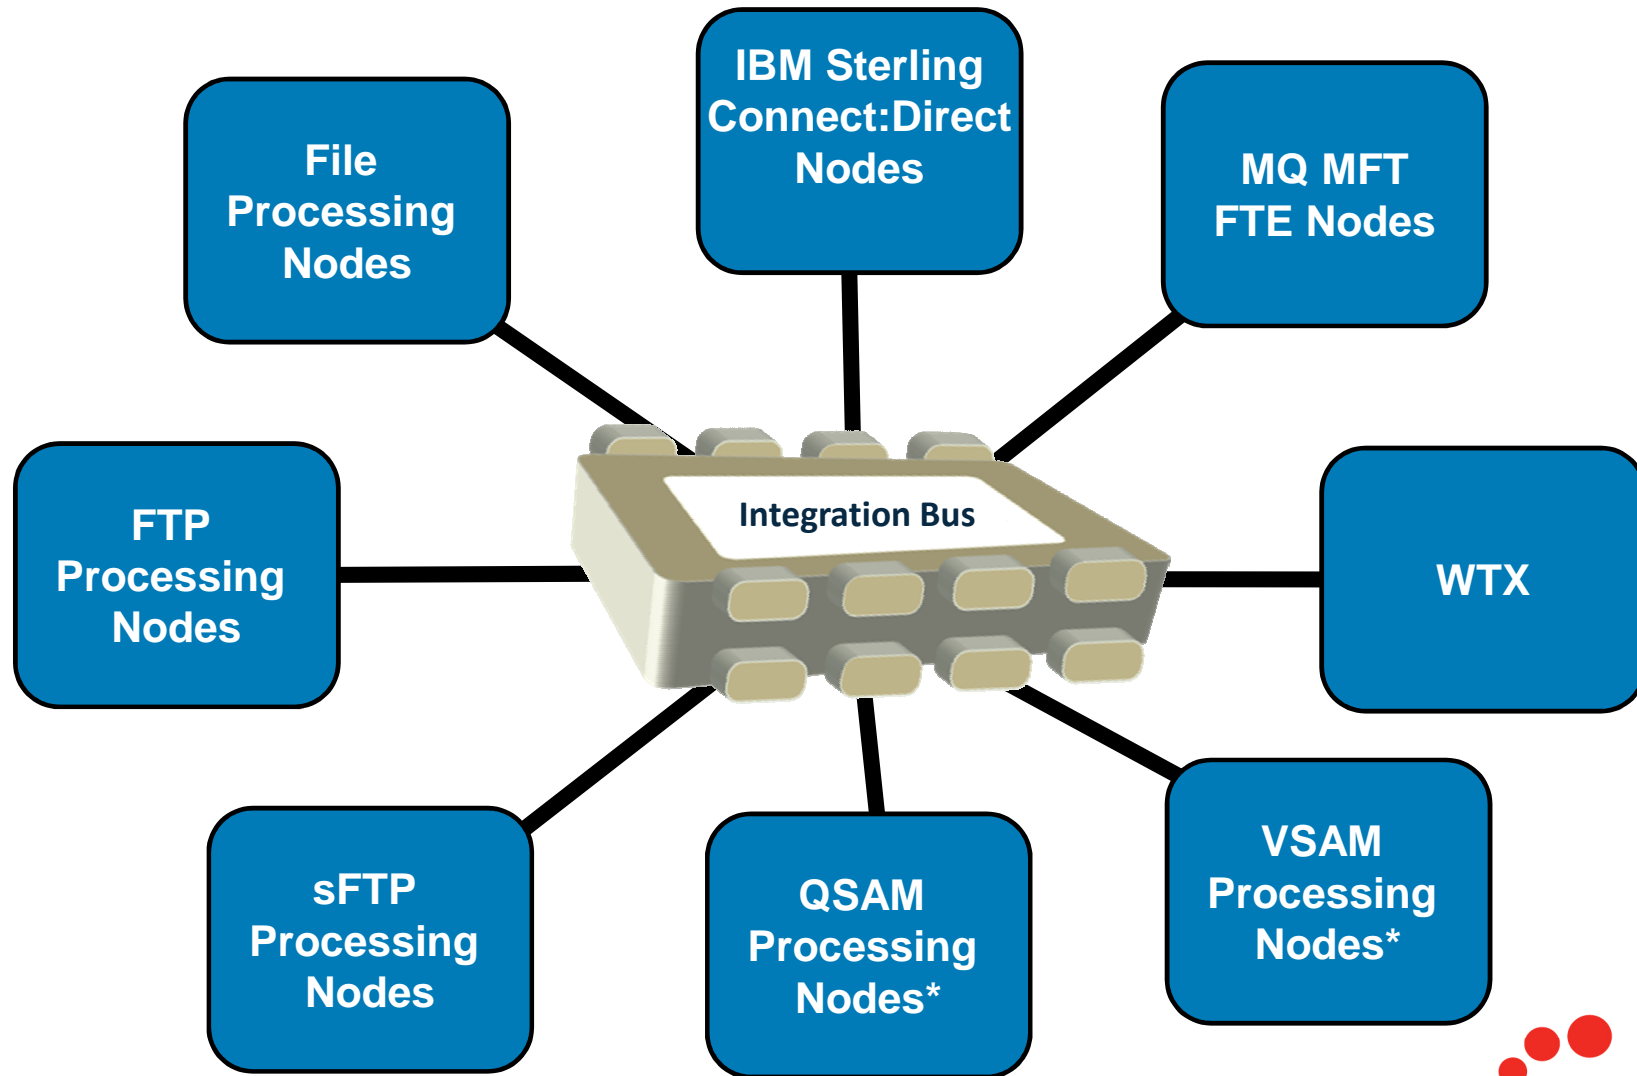
Metadata:

WrittenDestination	
FTE	
Name	test.txt
Directory	c:\\temp\\perf1
DestinationAgent	MB7BROKER.FTESAMPLEIN
DestinationQmgr	MB7QMGR
JobName	Job 1
Overwrite	true
TransferId	414D51204D4237514D4752202020202045C1314C2000780B
UserDefined	
PreDestinationCall	
Name	john1.xml
PostDestinationCall	
Name	john2.xml

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




File integration in IBM Integration Bus



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

* Support pac

IBM Integration Bus : FileInput Node



File Input

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





Notes on: FileInput Node



NOTES

- The FileInput nodes reads data from a file and triggers the start of processing in the flow. Three main areas will be covered in detail:
 - The mechanism used to detect a file is ready to be processed
 - The splitting of the file up into records
 - The archiving or deleting of the file once processing has been finished

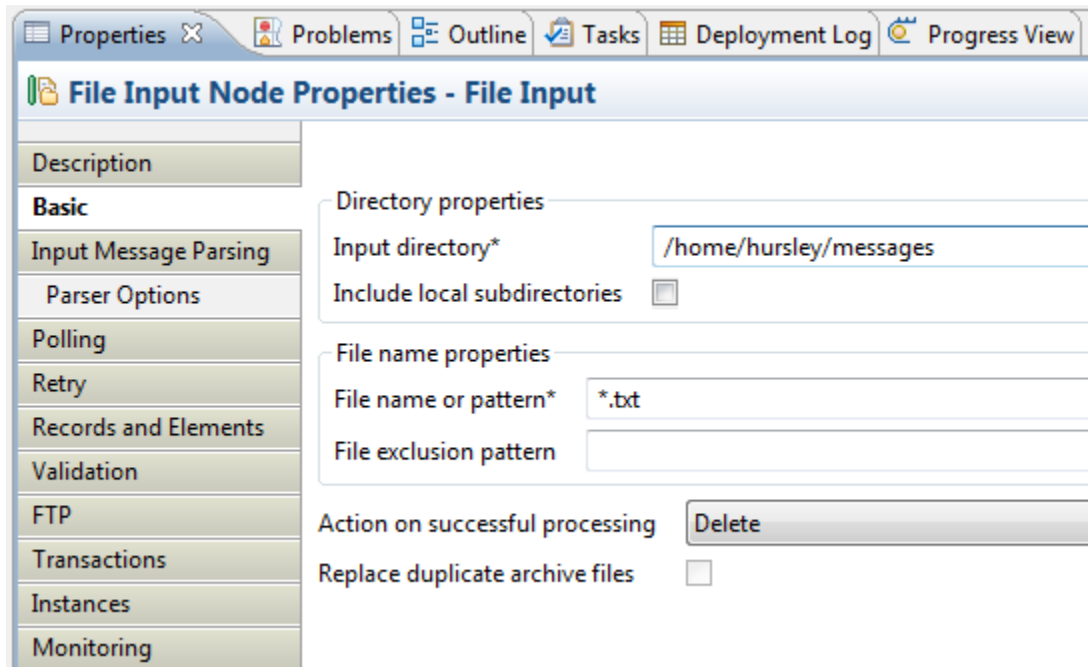
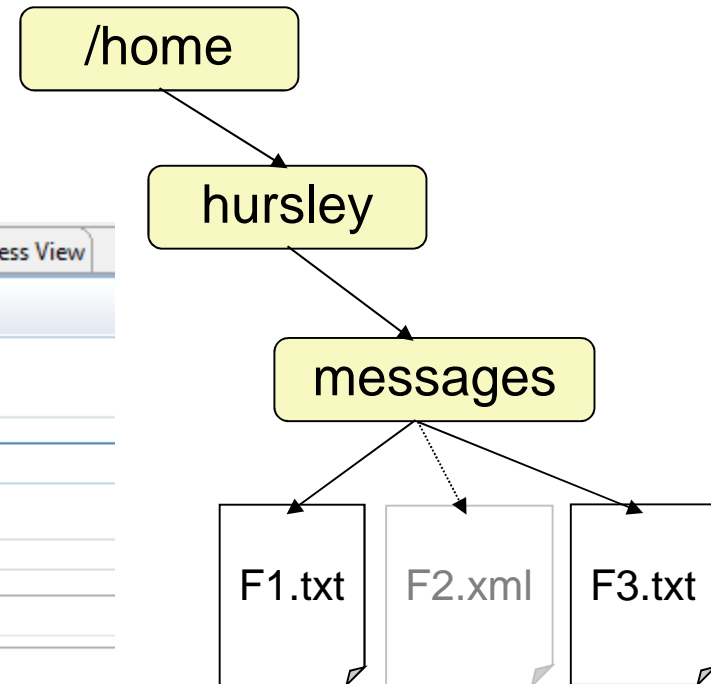
FileInput node – Operation



File Input



- Scans a pre-configured directory (relative or absolute) for files that match a given specification
- Locked files are ignored until they become unlocked



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



FileInput node – Record Detection



- Handling options (on the Records and Elements tab):
 - Whole file
 - Fixed Length *
 - Delimited *
 - Parsed Record Sequence *

Note * - results in separate records - message flow is invoked multiple times

- Only requires one record to be in memory at any one time
 - Allows very large files (Gigabyte) to be streamed efficiently
 - Streaming possible with DFDL, MRM (CWF and TDS) and XMLNSC parsers only

- If connected, 'End Of Data' terminal is triggered at end of file
 - Empty BLOB message and a LocalEnvironment.File structure



Record Detection Examples

- With input file `Boston|10|San Francisco|9|Seattle|8|`

- Whole file

- Propagates one message:

`Boston|10|San Francisco|9|Seattle|8|`

- Fixed length (size = 10 bytes)

- Propagates four messages:

`Boston|10|`

`San Franci`

`sco|9|Seat`

`tle|8|`

- Delimited (character = '|')

- Infix: propagates seven messages

`Boston`

`10`

`San Francisco`

`9`

`Seattle`

`8`

`“”`

- Postfix: propagates six messages

`Boston`

`10`

`San Francisco`

`9`

`Seattle`

`8`

Record Detection Examples 2

- With input file:

```
<cities><city name="Boston" rating="10"></cities>  
<cities><city name="San Francisco" rating="9"></cities>  
<cities><city name="Seattle" rating="8"></cities>
```

- Parsed record sequence

- When the parser is set to XMLNSC, this propagates three XML messages

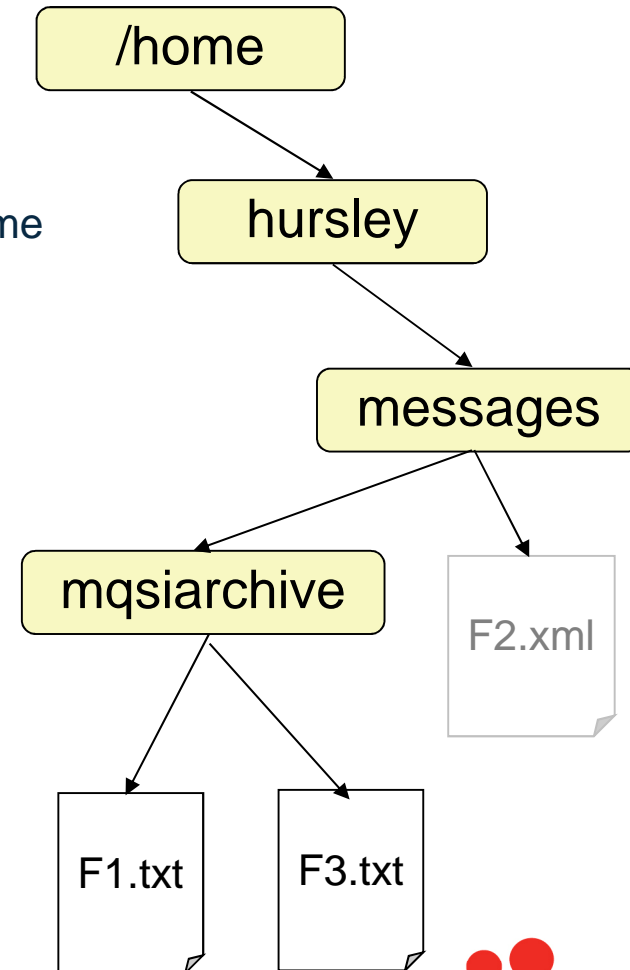
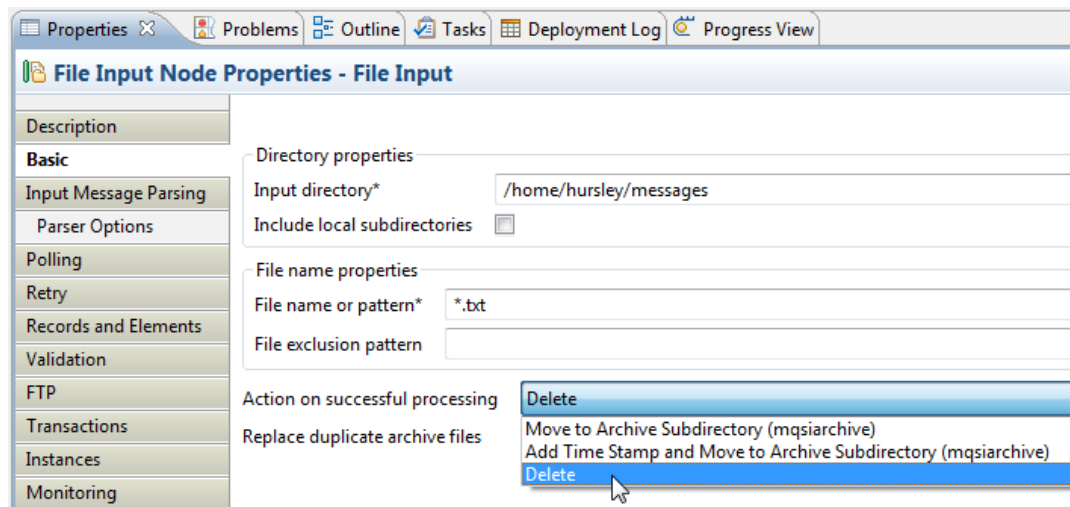
```
<cities><city name="Boston" rating="10"></cities>
```

```
<cities><city name="San Francisco" rating="9"></cities>
```

```
<cities><city name="Seattle" rating="8"></cities>
```

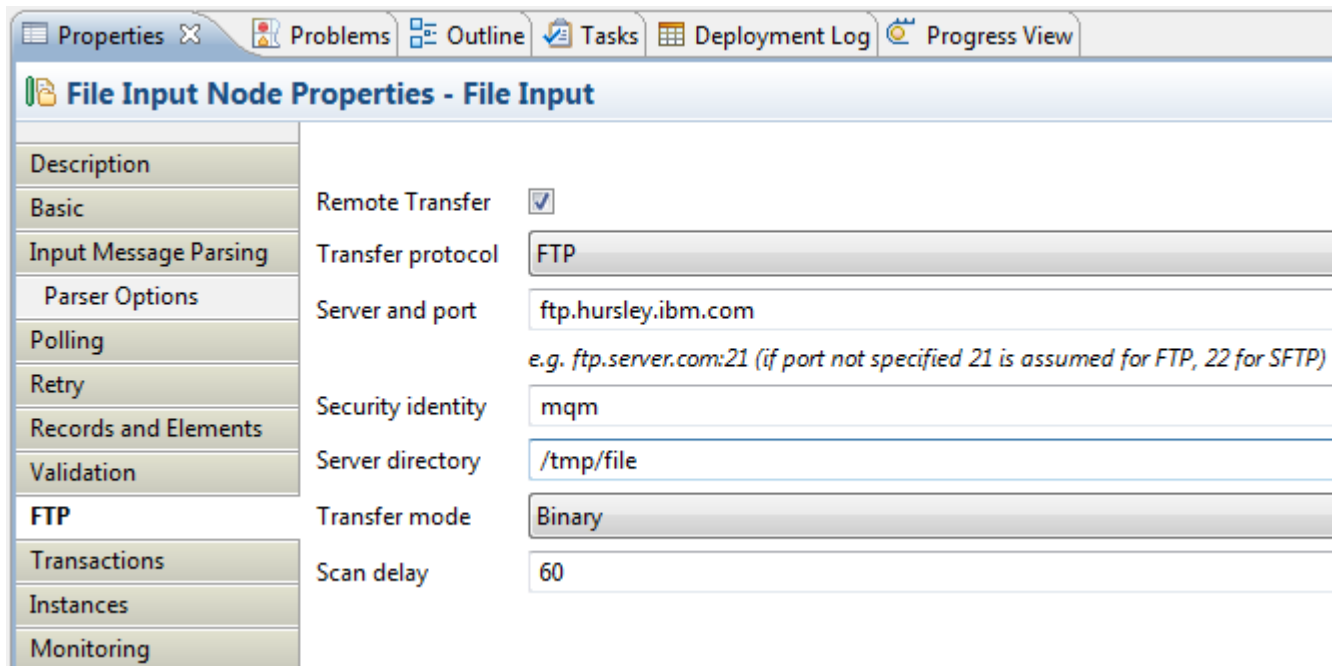
FileInput node – Archiving

- Upon successful processing, file is either deleted or moved to an *mqsiarchive* subdirectory
- Dealing with files with duplicate names:
 - Option to include timestamp in archived filename
 - Option to replace any existing file



FileInput node – using FTP and SFTP


- When active, FTP settings cause the node to periodically transfer files on a remote server to the local directory for input.
- Security Identity ‘UserMapping’ set using runtime command:
 - `mqsisetdbparms BROKER -n ftp::UserMapping -u USER -p PASS`



The screenshot shows the 'File Input Node Properties - File Input' dialog box. The 'FTP' tab is selected in the left-hand navigation pane. The main area displays the following configuration:

Remote Transfer	<input checked="" type="checkbox"/>
Transfer protocol	FTP
Server and port	ftp.hursley.ibm.com <small>e.g. ftp.server.com:21 (if port not specified 21 is assumed for FTP, 22 for SFTP)</small>
Security identity	mqm
Server directory	/tmp/file
Transfer mode	Binary
Scan delay	60

IBM Integration Bus : FileOutput Node



File Output

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






Notes on: FileOutput Node



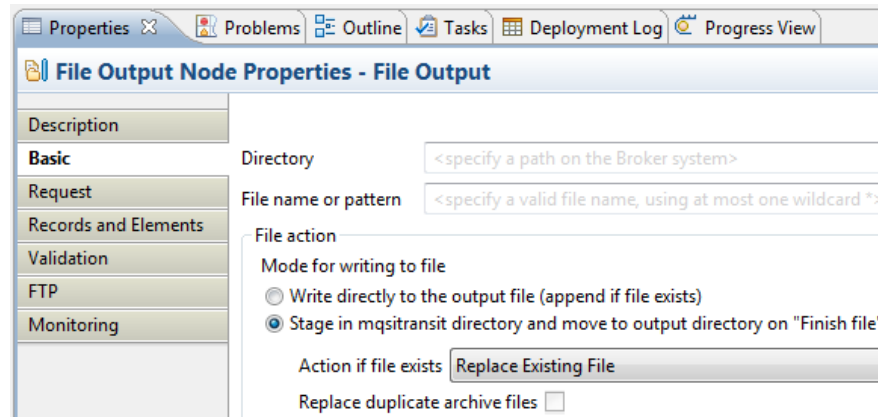
NOTES

- The file output node is used to create and write to a file anywhere in the middle of a flow:
 - How it decides which file to create and where
 - How a file is created using a series of records appended
 - What happens if the file it attempts to create already exists

FileOutput – Writing files

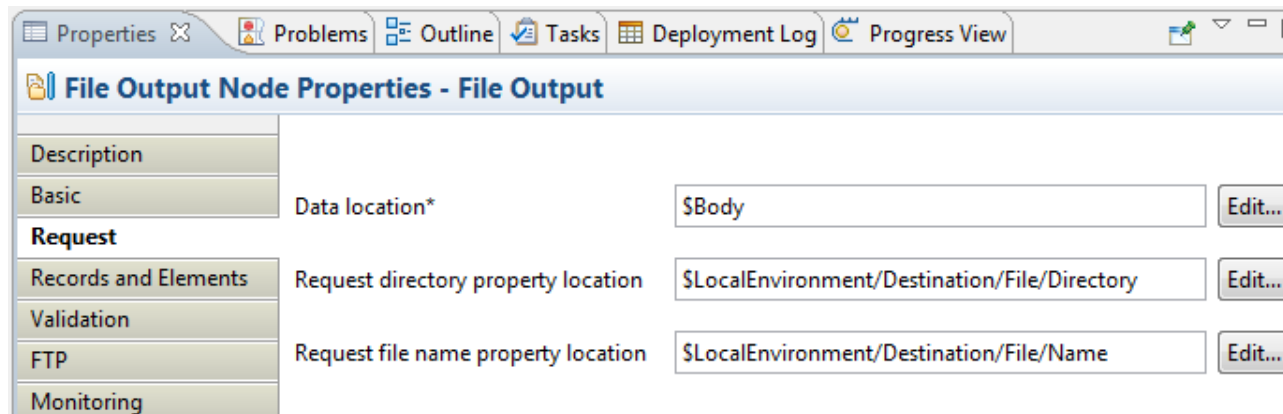


- In the simplest scenario, the received message body is written to the pre-configured file:



- When writing to the output file, the wildcard (if present) is replaced with the value of *LocalEnvironment.Wildcard.WildcardMatch*

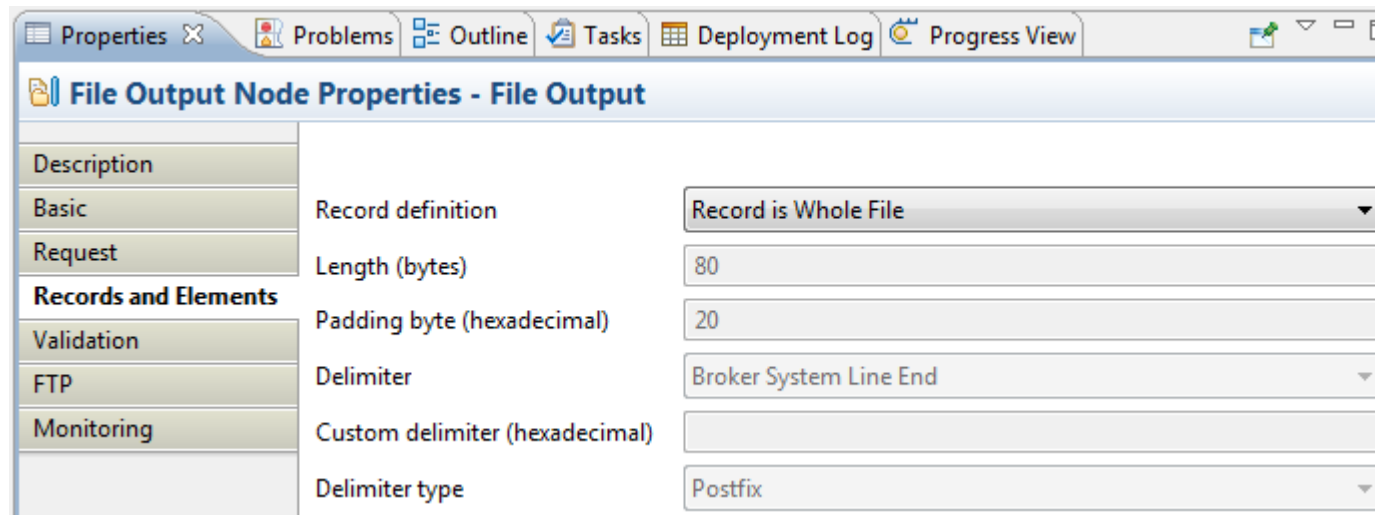
- Allows you to preserve elements of a filename during processing



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



Appending Records

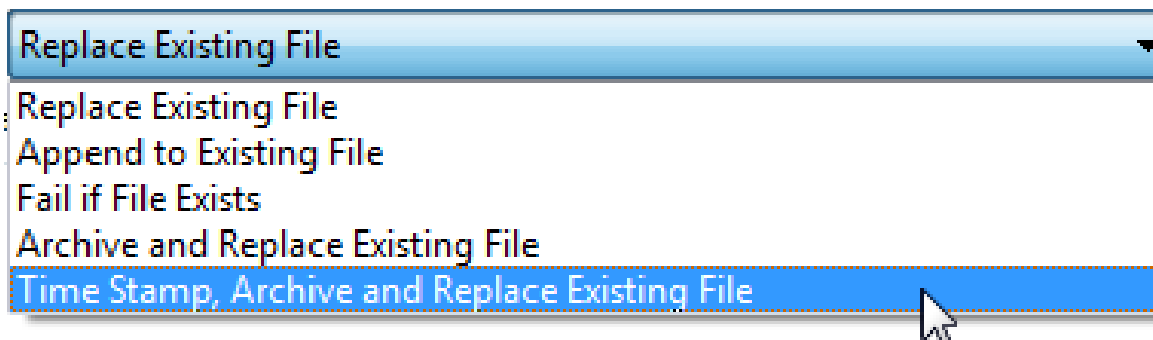


- “Records and Elements” tab defines how multiple writes to the same file are handled
- Record definition options:
 - Record is Whole File – close file automatically after first write
 - Record is Unmodified Data – the message bit-stream appended to file
 - Record is Fixed Length Data – specify length in bytes and padding character
 - Record is Delimited Data – specify delimiter and infix/postfix option
- Unless “Record is Whole File” is selected, the file will be closed when the “Finish File” terminal is triggered
 - The Finish File message is propagated on to the FileOutput’s “End Of Data” terminal
- The *mqsitransit* subdirectory holds all files that have not yet been closed

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

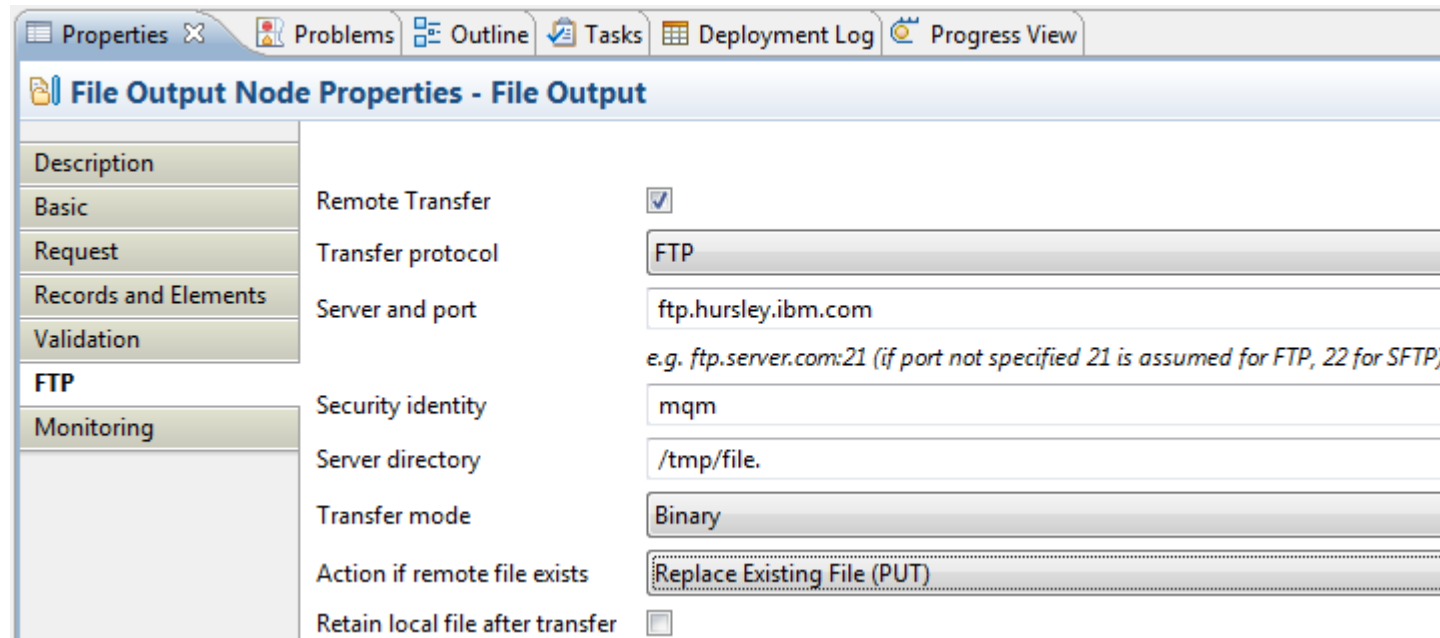
Options if the file already exists

- Output file action (basic tab)



- If the file already exists
 - Replace it
 - Append to it
 - Go down failure terminal
 - Move to mqsiarchive subdirectory and replace
 - Add timestamp, move to mqsiarchive subdirectory and replace

FileOutput - FTP Support




Properties	Value
Remote Transfer	<input checked="" type="checkbox"/>
Transfer protocol	FTP
Server and port	ftp.hursley.ibm.com <small>e.g. ftp.server.com:21 (if port not specified 21 is assumed for FTP, 22 for SFTP)</small>
Security identity	mqm
Server directory	/tmp/file.
Transfer mode	Binary
Action if remote file exists	Replace Existing File (PUT)
Retain local file after transfer	<input type="checkbox"/>

- If enabled, whenever a complete file is closed an FTP transfer of the file is attempted to the supplied FTP server
- File is optionally deleted from the local file system when the transfer completes
- Transfer is synchronous
 - Use additional instances if throughput rate is an issue
- If remote file exists, choose either replace or append

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

IBM Integration Bus : FileRead Node



File Read

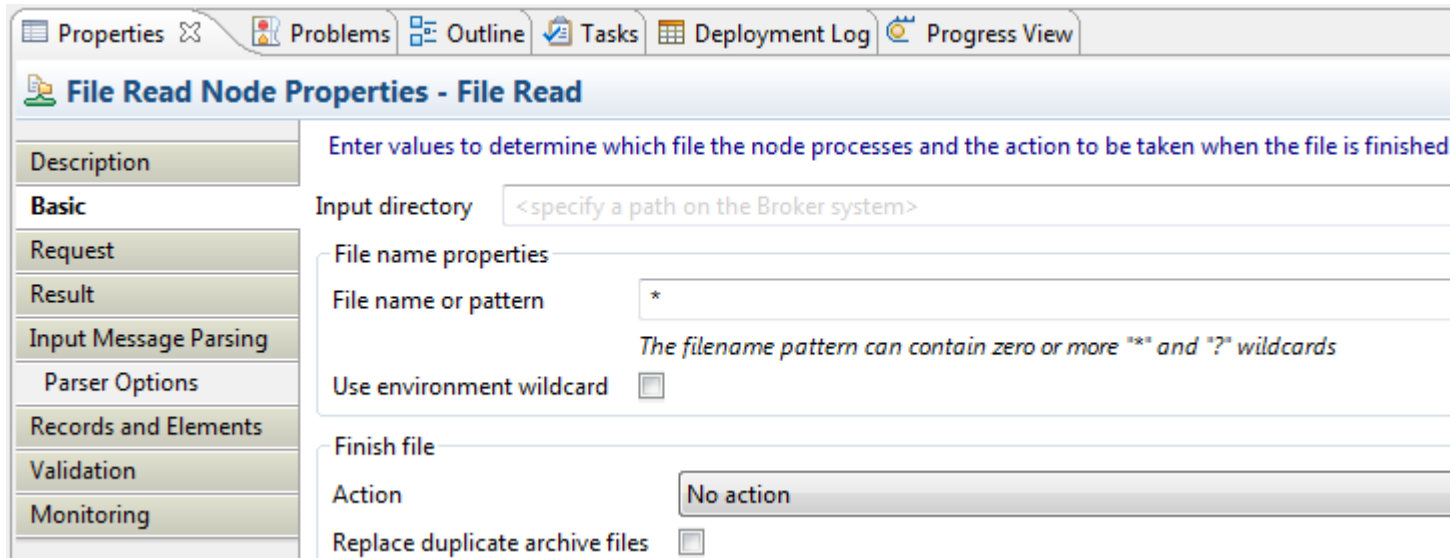
Complete your session evaluations online at www.SHARE.org/Seattle-Eval



FileRead node



- File Read Node
 - Reads data in middle of flow like an MQGet or a TCPIPReceive node
 - Reads either whole file contents or one record from the file
 - Allows user to override the file to be read and the offset within the file to start reading from
- Has a *No match* terminal which the message is sent to if it can not find a file or a record in the file



The screenshot shows the 'File Read Node Properties - File Read' dialog box. It has a tabbed interface with 'Properties', 'Problems', 'Outline', 'Tasks', 'Deployment Log', and 'Progress View'. The 'Properties' tab is active, showing a description: 'Enter values to determine which file the node processes and the action to be taken when the file is finished.' Below this is a 'Basic' section with the following fields:

- Input directory:** A text box containing '<specify a path on the Broker system>'
- File name properties:**
 - File name or pattern:** A text box containing '*'. Below it is a note: 'The filename pattern can contain zero or more "*" and "?" wildcards'
 - Use environment wildcard:** A checkbox that is currently unchecked.
- Finish file:**
 - Action:** A dropdown menu currently set to 'No action'.
 - Replace duplicate archive files:** A checkbox that is currently unchecked.

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





Notes on: FileRead Node



NOTES

- FileRead node behaves like a MQGet or TCPIPReceive node in the sense that it reads in data within a flow without first sending data out. For example: MQGet reads a message from a queue, TCPIPReceive reads data from a TCPIP input stream and the FileRead node reads data from a file.
- Can either read the whole contents of a file or a single record from the file and then parses and constructs a message to propagate down the flow.
- The node is very configurable both at design time and during runtime where most properties can be overridden based on the message content.
- The basic properties are similar to a File input node where the details of the file to process are given.

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



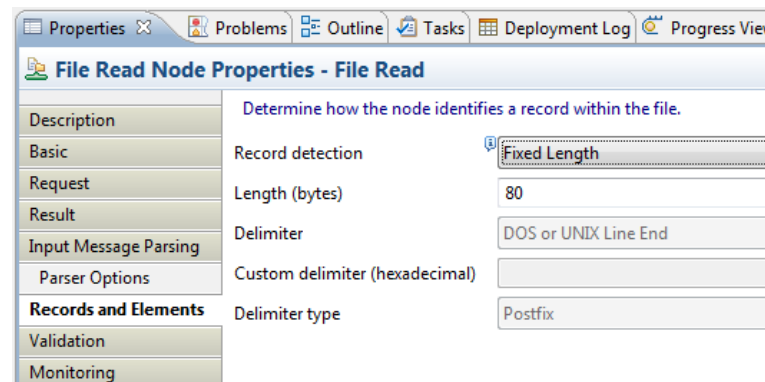
Defining which record to read and propagate



- **Where the record starts**
 - Defaults to the beginning of the file
 - Give the offset into the file to start record from based on contents of the message

- **Where the record ends**
 - Define the record detection mechanism:

- Whole file
- Fixed size
- Delimited
- Parser



- **Which record to propagate**
 - Define an expression to specify which record to propagate
 - Node iterates through all records from the start offset until one matches
 - Only propagates the first match

Record selection expression*

example: \$InputRoot/XMLNSC/Invoice/Name=\$ResultRoot/MRM/KeyField





Notes on: Defining which record to read and propagate

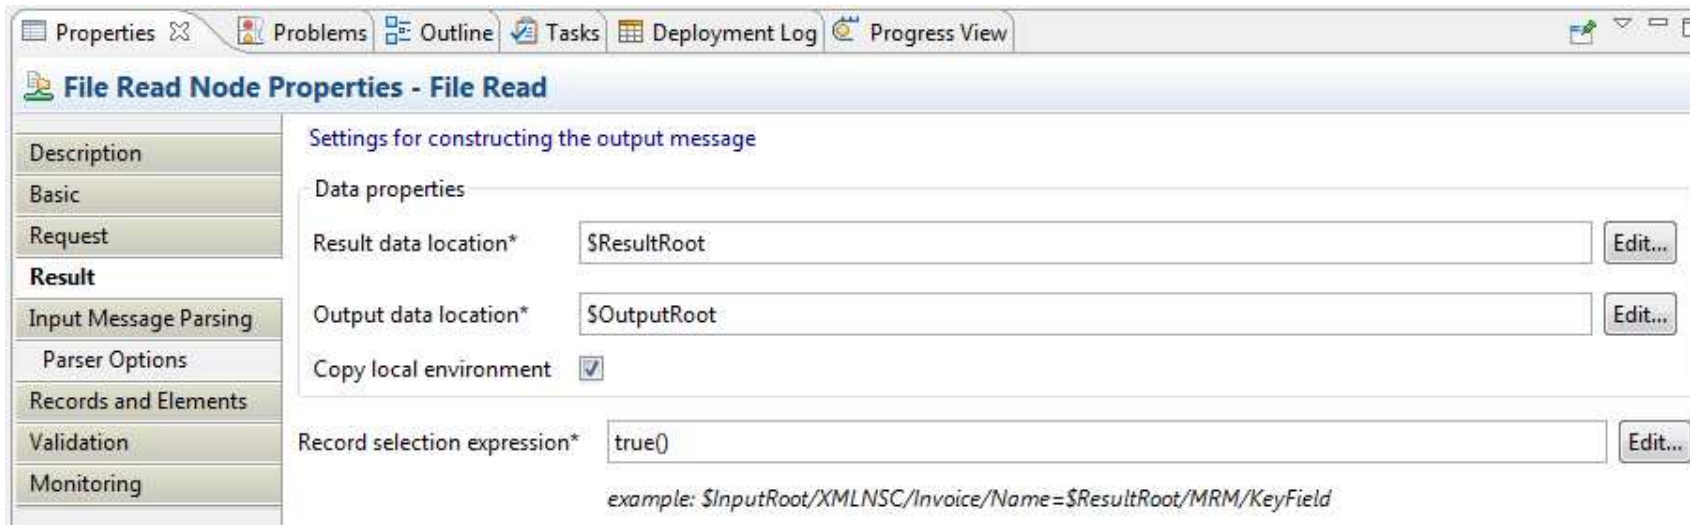


NOTES

- Three key pieces of information need to be defined to create and propagate a record from the file.
- Where the record starts. By default the file read node always starts the reading of a record from the beginning of the file. It does not remember or store the result of the last record read so the next time through the node will start from the same place. Unless the user specifies in the local environment where to start from. It is possible to configure the node to use the end record offset from a previous fileread node as the start of the current record.
- Where the record ends. The normal record detection mechanisms are used to find the end of the record. Fixed size just reads that many bytes, delimited scans for a delimiter and parser uses an integration node parser to determine the end (like XMLNSC if the record is an XML document). When using fixed size it is possible to override the length being used using the message content or local environment.
- Which record to propagate. After finding the record the *Record selection expression* is evaluated. If it is true then the record is propagated otherwise the next record is found using the end of the last record as the start of the new one. The process is repeated until either the expression is true or the end of file is found. Only the first matching record is propagated. If no record matches then the file gets sent to the no match terminal.

Constructing the message using the record read location

- Which part of the record read to propagate
 - Result data location
- Where to put the record in the outgoing message
 - Output data location



The screenshot shows the 'File Read Node Properties - File Read' dialog box in a software application. The dialog has a sidebar on the left with tabs for Description, Basic, Request, Result, Input Message Parsing, Parser Options, Records and Elements, Validation, and Monitoring. The main area is titled 'Settings for constructing the output message' and contains the following fields:

- Data properties**
- Result data location***:
- Output data location***:
- Copy local environment**:
- Record selection expression***:

example: `$InputRoot/XMLNSC/Invoice/Name=$ResultRoot/MRM/KeyField`



Notes on: Constructing the message using the record read location



NOTES

The *Result* panel has properties which specify how the outgoing message is constructed based on the contents of the file and the incoming message.


By default, the whole incoming message is replaced with the contents of the record retrieve from the file.

The Result data location is used to extract a piece of information from the file to insert into the outgoing message.

The Output data location is used to find the location to write the extracted information to.

Changing the file disposition after the read

- By default file is left unchanged after read
- Disposition change is always done when the FileRead node executes *Finish File*:
 - The end of the file is reached
 - A message is sent to the finish file terminal
- Following actions are available:



- Archived files are moved to the mqsiarchive directory
- Can override the archive directory and archive name using local environment or data in the message



Notes on: Changing the file description after the read




NOTES

- By default, after any read, the file read node will leave the file unchanged and will close any connections to it.
- It is possible to configure the node to modify the file disposition when the nodes *finish file* action is triggered. Finish file is defined as either when the file read node reads to the end of a file (whole file mode or when the last record is read) or when a message arrives at the *finish file* terminal.
- It is possible to delete or archive the file. Archiving moves the file to the mqsiarchive directory but it is possible to override this based on the contents of the message to move the file to any new directory and any new file name.

IBM Integration Bus : Connect:Direct nodes



CD Input



CD Output


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



IBM Sterling Connect:Direct (no IIB)



- IBM Sterling Connect:Direct is a managed file transfer solution



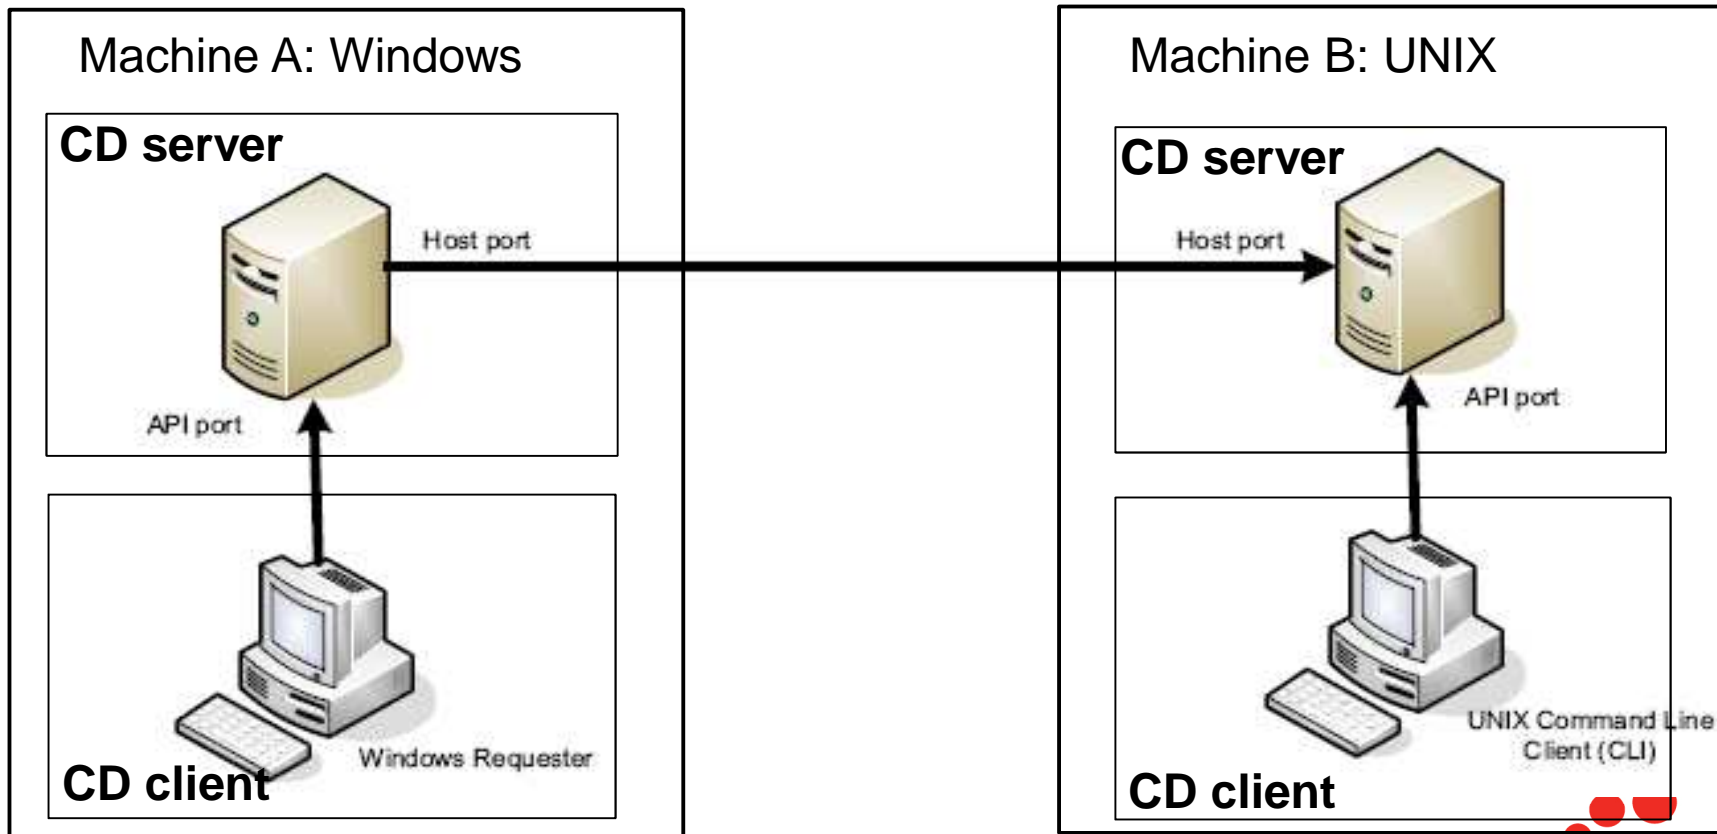
configure



track



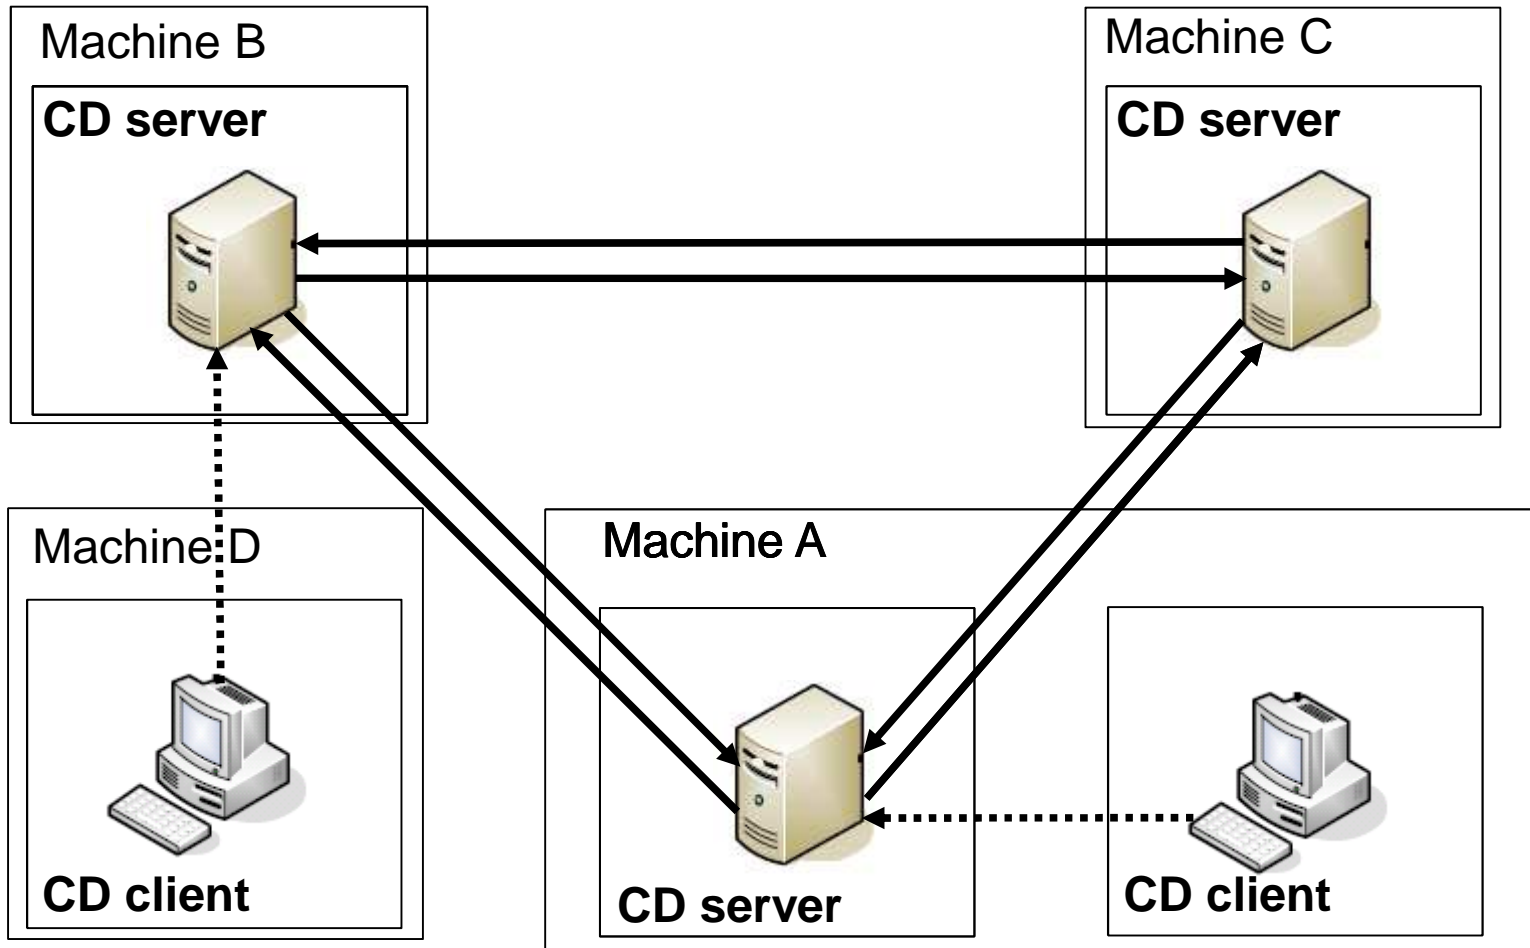
audit



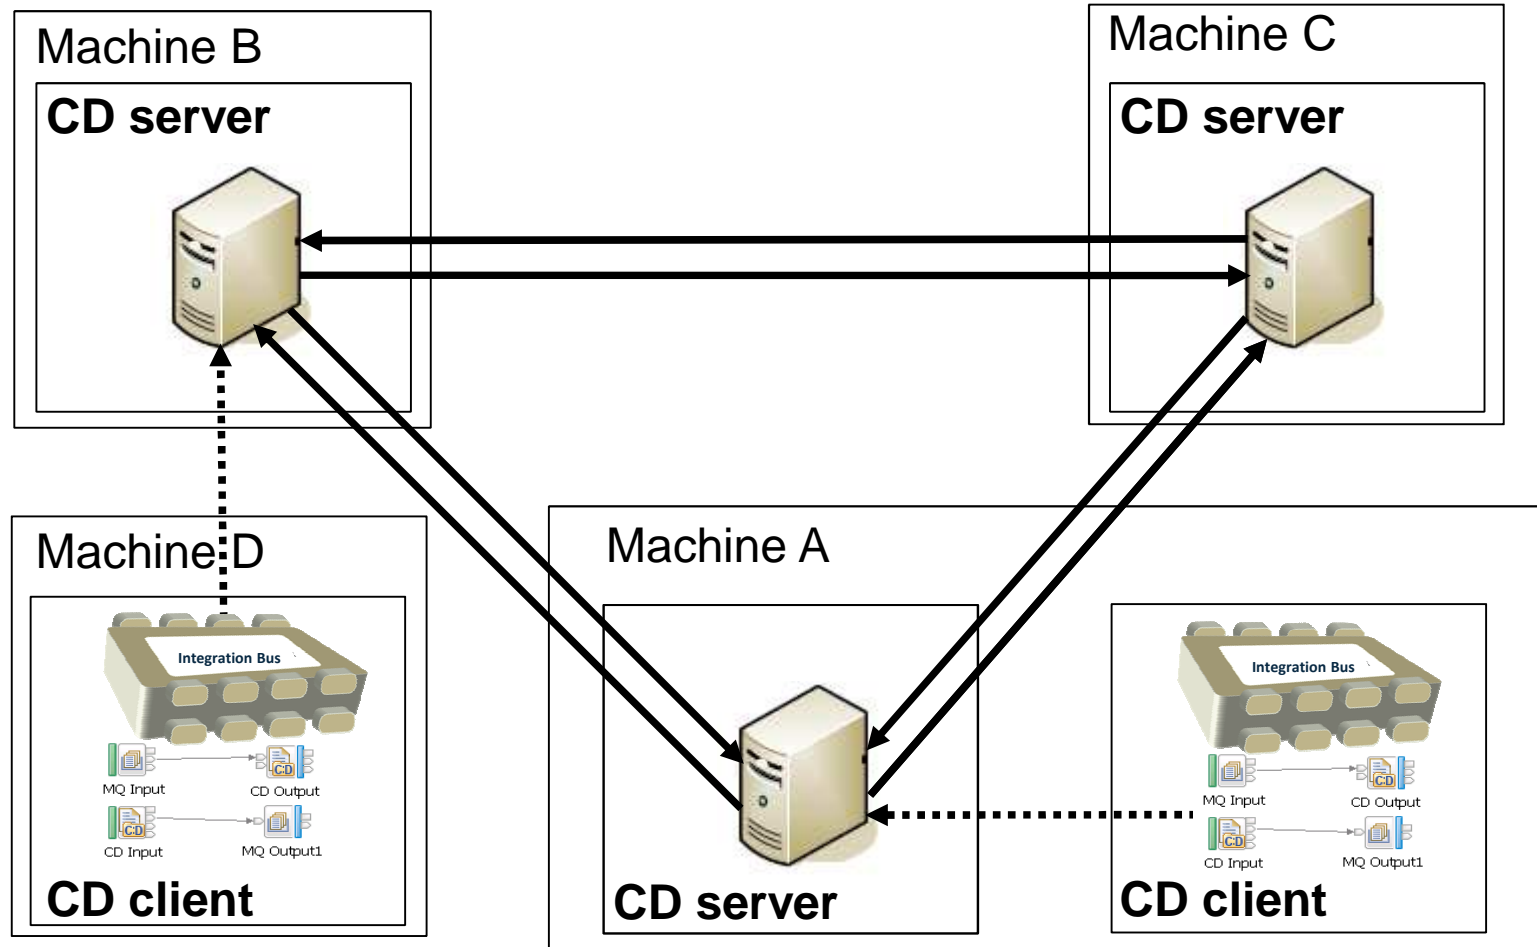
Complete your session evaluations online at www.SHARE.org/Seattle-Eval



IBM Sterling Connect:Direct (no IIB)



IBM Sterling Connect:Direct (with IIB)




Adding IBM Integration Bus to a CD network

- Install Integration Bus
- Create and start an integration node on the machine
- Set up a security identity to be used to connect to CD server:

```
– mqsisetdbparms BROKER –n cd::default –u gorman–p *****
```


- Create and deploy flows which interact with CD via built in CD nodes



- That is all that is needed if the broker is on the same machine

Adding IBM Integration Bus to a CD network

- IBM Integration Bus does not have to be on the same machine as CD server.
- They must have access to a shared file system where files are transfer to.
- A configurable service can be created with details of CD Server:
 - Hostname
 - API port to connect to
 - Security identity
 - Filepath of the shared file systems



Accessing a CD server from a message flow

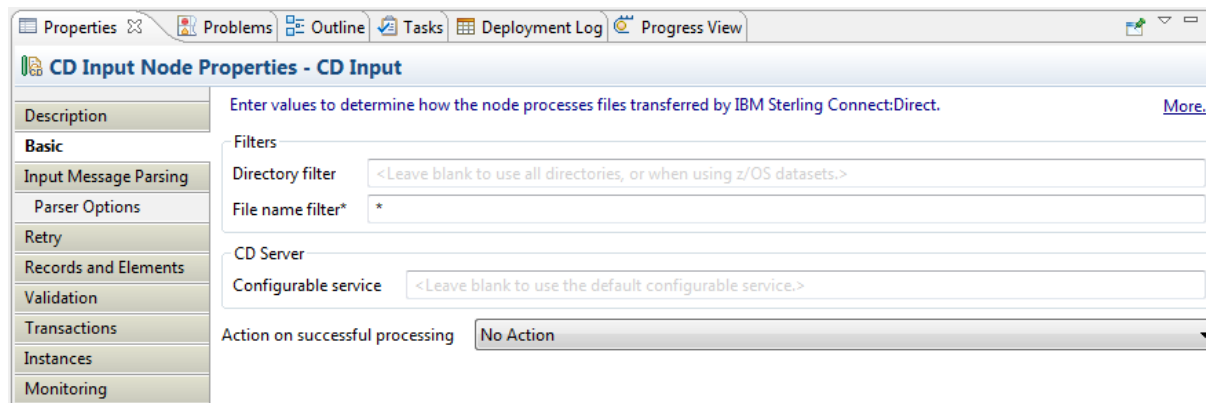


- Messages sent and received from the CD server using Integration Bus nodes:
 - CDInput node: receives any file transferred to the CD server
 - CDOOutput node: constructs a file and sends a request to the CD server to transfer the file to a remote system.
- CD nodes based on the current file nodes allowing:
 - Record based processing
 - Stream parsing for large file support
 - Can be used with all other message flow nodes
- CD server is not stopped, started or administered by Integration Bus. Both products are decoupled:
 - Integration Bus outages have no effect on CD server
 - CD server outages only effect Integration Bus when a transfer is requested by a message flow

CD Input Node



- Consistent with file input node but makes full use of CD function
- Timely
 - The CD node monitors the CD servers stats for transfers that have completed
 - The node processes the files in the transfer immediately.
 - Each copied file is processed independently
 - Can leave the file unchanged after processing and just delete notification message
- Metadata
 - Metadata associated with the transfer is sent with the notification
 - Includes user defined application data field
- Filter
 - The node by default receives all transfers
 - Can specify which files to receive using a filter



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



CD Input Node – metadata

[-] ◆ LocalEnvironment	
[-] ◆ CD	
[-] ◆ Transfer	
◆ ProcessName	BRTOHQ
◆ Step	CPBRHQ
◆ ProcessNumber	20536
◆ Submitter	mqbroker
◆ Accounting	
◆ SourcePath	C:\\temp\\CD_demo\\to_HQ\\Account_transfers.req
◆ DestinationPath	C:\\Program Files\\Sterling Commerce\\Connect Direct v
◆ PNODE	JREEVE6
◆ SNODE	JREEVE6
◆ TimeStamp	20110303_174443_333773
◆ Offset	0
◆ Record	1
◆ Delimiter	
◆ IsEmpty	false

CD Output Node



- Consistent with file output node but makes full use of CD function
- Transfer details
 - The destination CD server, directory etc are defined on the node

The screenshot shows the 'CD Output Node Properties - CD Output' dialog box. The 'Basic' tab is selected. The 'Destination' section includes fields for SNode (defaulting to the PNODE), Directory (defaulting to the default directory), and File name* (with an example: test.txt or *.txt). The 'CD Server' section has a 'Configurable service' field (defaulting to the default service). The 'Options' section includes 'Process name' (defaulting to a broker-generated name), 'Disposition' (set to 'RPL - Replace or create a new file'), and 'Transfer mode' (set to 'Binary transfer (no conversion)').

- Support also for wild card file names
- Staging
 - Uses a local staging directory to build up a file record by record for transfer
 - Once a file is finished a request is sent to the CD Server to transfer the file
- Metadata
 - User data if provided in the Local Environment is sent with the transfer
 - Other metadata is generated by CD server or by Broker

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



CD Output Node – Metadata



[-] ◆ LocalEnvironment	
[-] ◆ WrittenDestination	
[-] ◆ CD	
◆ ProcessName	HQTOBR
◆ ProcessNumber	20542
◆ Directory	C:\\temp\\CD_demo\\from_HQ
◆ Name	branch.rply
◆ PrimaryNodeName	JREEVE6
◆ PrimaryNodeOS	Windows
◆ SecondaryNodeName	JREEVE6
◆ SecondaryNodeOS	Windows

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



File types supported

- Flat files on windows, UNIX and z/OS
 - Treated and processed the same as in the normal file nodes
- z/OS Sequential datasets
 - File name is the dataset name: JREEVE.TEST1.TEST2
 - Wildcard values allowed anywhere: JREEVE.*.*
 - Staged to HFS and then treated the same as in the normal file nodes
- z/OS Partitioned datasets
 - For the input node file name is a pattern like:
 - JREEVE.TEST1.TEST3(MEM1)
 - JREEVE.TEST1.TEST3(MEM*)
 - JREEVE.TEST1.TEST3(*)
 - JREEVE.*.TEST3(*)
 - Each member is staged to HFS and then treated the same as in the normal file nodes
 - For the output node the file name must be a single member:
 - JREEVE.TEST1.TEST3(MEM1)

IBM Integration Bus : VSAM and QSAM Support Pacs



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

SupportPac IA13 – VSAM nodes



- A suite of nodes allowing users to perform record oriented processing on VSAM files.
 - VSAMDelete
 - VSAMInput
 - VSAMRead
 - VSAMUpdate
 - VSAMWrite
- The broker can process VSAM records in the same way as it processes messages from or to other data sources.
- A suite of 5 nodes allowing users to perform record oriented processing on VSAM files: Input, Read, Write, Update and Delete operations.
- Users can combine these nodes for VSAM access with other nodes to integrate VSAM processing into message flows, or use VSAM files to drive message flow processing
- VSAM file support
 - KSDS, ESDS, RRDS, KSDS_PATH, ESDS_PATH
- **Fully supported Cat 3 SupportPac for V7**





NOTES

Notes on: VSAM nodes



- **Suite of Five nodes**
 - A comprehensive set of nodes allowing you to perform record oriented processing on VSAM files.
 - You can start a message flow with the VSAM input node, and there are corresponding nodes to perform subsequent read, write, update and delete operations on VSAM files.
 - These nodes can be mixed and matched with all other supplied nodes to create sophisticated VSAM processing capabilities.
- **Input Node**
 - The VSAM Input node is targeted at batch oriented processing. This node can read a set of records from a VSAM file (identified by an input control message) and propagate them down connected nodes.
 - Unlike message oriented processing with MQ, VSAM does not have an obvious ****trigger**** to identify when a file is to be processed.
 - The “control” terminal provides the trigger. This is usually wired to a different message flow, so that VSAM processing can be driven from another flow’s processing – it provides the trigger.
 - This node also requests additional instances, so that multiple threads can be processing a VSAM file in parallel, if processing allows.
- **Read Node**
 - Whereas the input node can read a set of records, the read node reads a single record, and places it at a user specified location in the message tree.
 - Properties identify and parse the record to be read
 - Basic, Default, Request groups (Key, RBA, RRN, ...)
 - Additional Properties
 - Advanced, Result, Status groups
 - VSAM Request message control message overrides some of these properties to allow very flexible and programmatic processing.
- **Write, Update, Delete Nodes**
 - These all follow the same pattern as read node, performing write, update or delete operations respectively.
- **Transactionality** The VSAM operations that are performed by these nodes do not participate in the transactions that exist within the message flow. Successful VSAM operations are not automatically backed out if a failure occurs in the message flow. As such you will need to bear this behaviour in mind when designing applications using the VSAM nodes.
- **SHAREOPTIONS.** When you consider the effect of using VSAM SHAREOPTIONS on a VSAM data set, note that an integration server rather than a message flow should be regarded as the equivalent of a user task or application program. **Note:** The services that are used by the VSAM nodes to access VSAM data sets do not use SHAREOPTIONS to enforce which intents the data set can be opened with. Therefore, any VSAM data set can be opened for read or update intent by the VSAM nodes, regardless of the setting of SHAREOPTIONS.

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



SupportPac IA11 – QSAM nodes

- These are similar in concept and usage to the VSAM nodes, but oriented around sequential files, rather than record oriented files.
 - FileDelete
 - FileRead
 - FileRename
 - FileWrite
- XML action control messages used to indicate when to open, close, delete rename files.
- Facilities provided
 - File to file
 - File to queue
 - Queue to file
- **Fully supported Cat 3 SupportPac for V8**



Summary

- MQ MFT
- MQ MFT(FTE) nodes in IIB
- File nodes in IIB
- IBM Sterling Connect:Direct nodes in IIB
- VSAM and QSAM support pacs

This was Session 17054. 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]	
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]		
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!	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]	



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

