IBM

IBM WebSphere Application Server V8.5

Cross-component trace

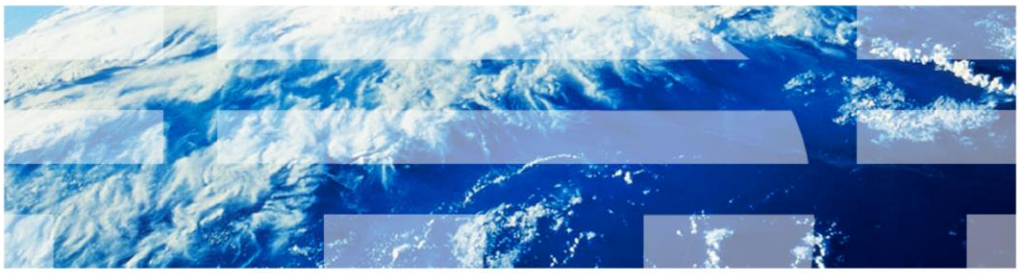

© 2012 IBM Corporation

This presentation describes support for cross-component trace (XCT) included in IBM WebSphere® Application Server V8.5

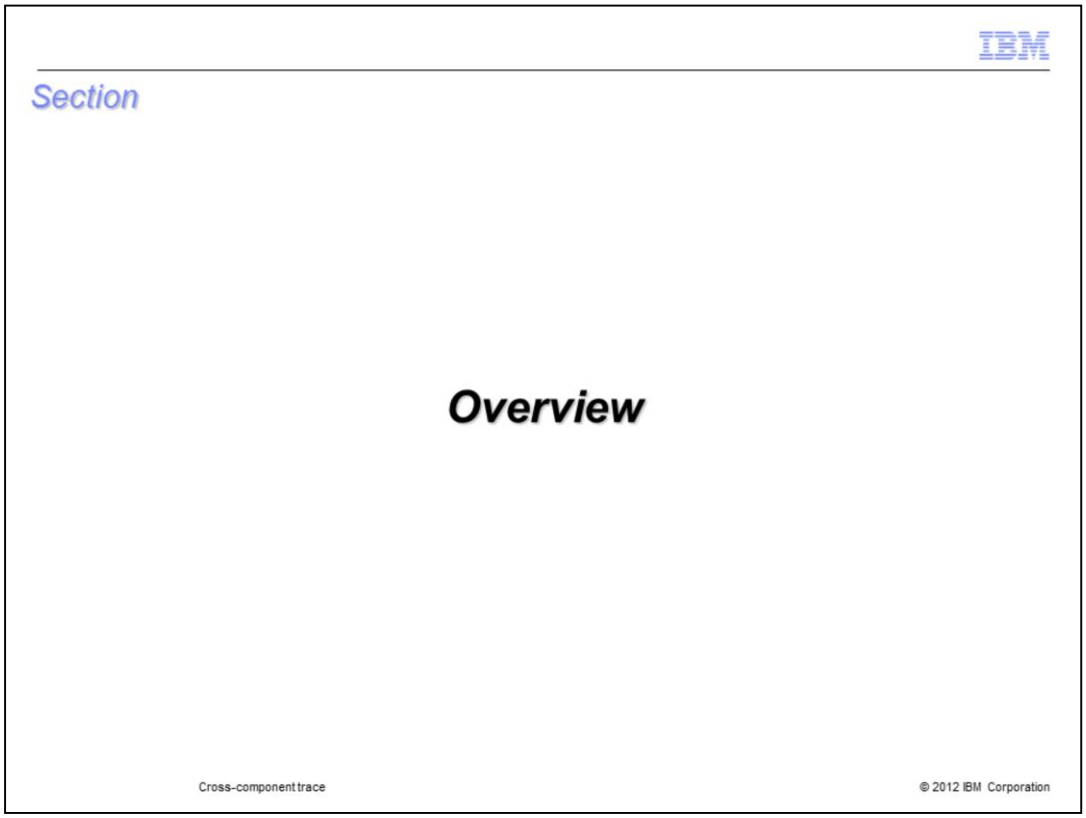

This section is an overview of cross component trace.

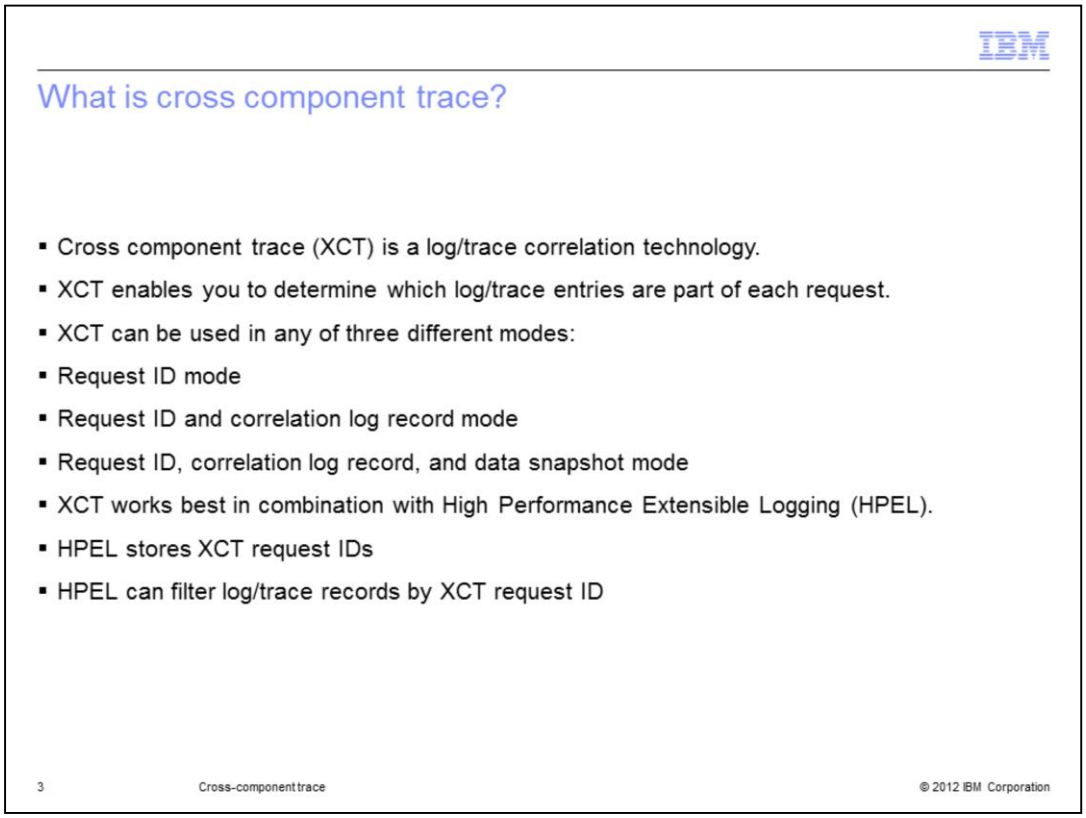

Cross component trace is a correlation technology that helps administrators see the flow of requests that span multiple threads or processes. XCT simplifies the task of determining which log or trace entries, in each application server log file, are part of each request. When enabled, XCT can be used in any of three different modes. In the first mode, XCT request IDs are added to existing log and trace records. In the second mode, XCT request IDs are added to existing log and trace records and XCT log records are added to log files. In the third mode, XCT request IDs are added to existing log and trace records, XCT log records are added to log files, and data snapshots are captured. XCT works best with the application server's High Performance Extensible Logging (HPEL) log and trace framework. XCT Request Ids are only ever added to HPEL log and trace records – they cannot be stored in SystemOut.log. HPEL also provides the ability to filter log and trace files by request ID and helps minimize the performance impact of enabling XCT log records.

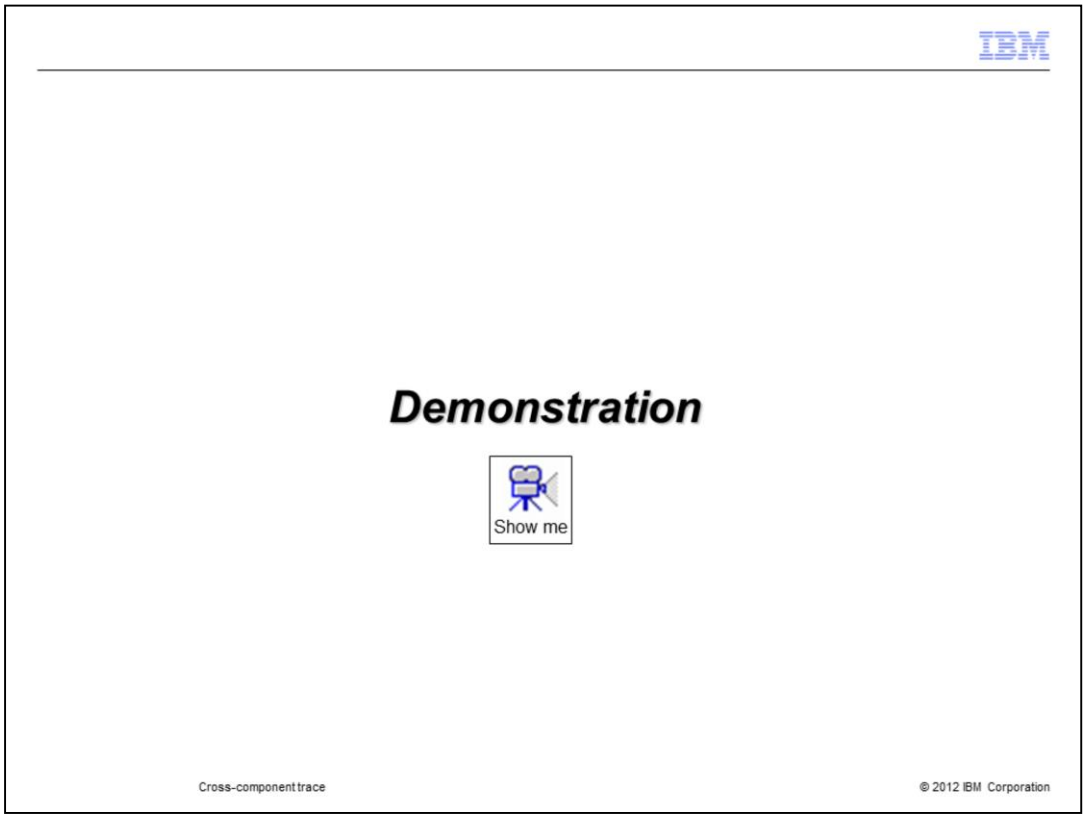

For a demonstration of how to enable high performance extensible logging, pause this presentation and click the demonstration icon.

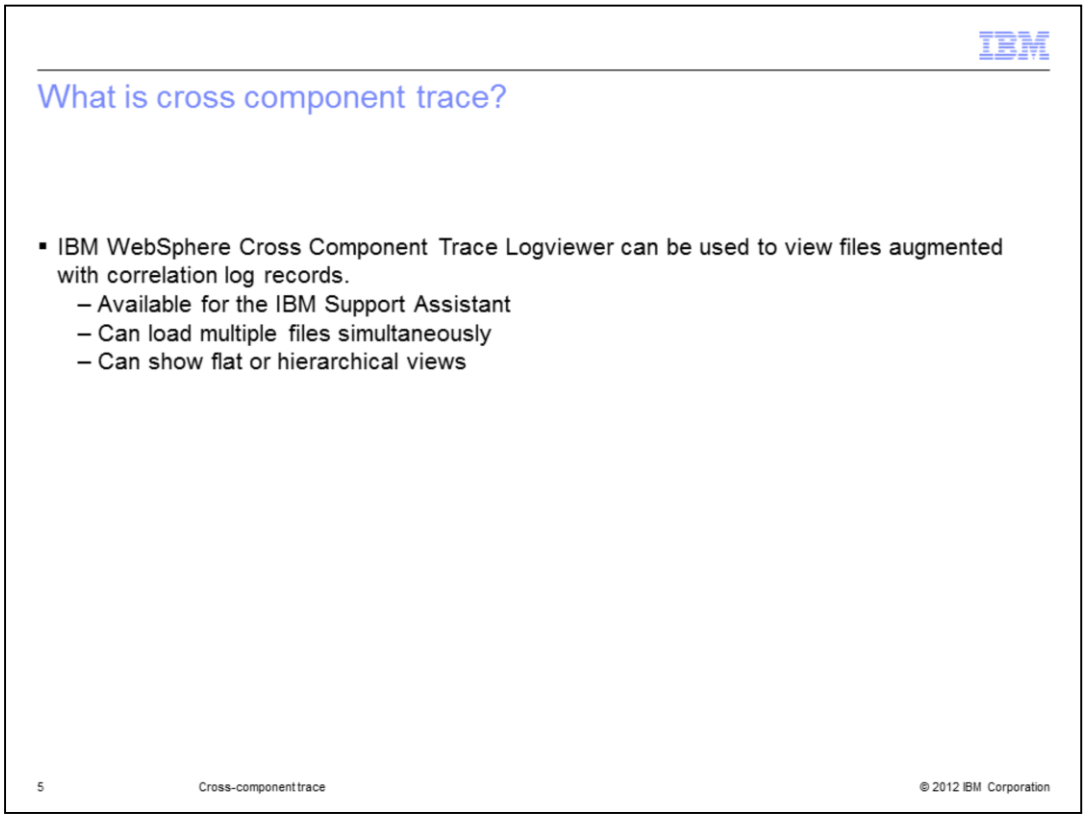

A tool called IBM WebSphere cross component trace log viewer can be used to view log or trace files augmented with correlation log records. The tool is available as a tool add-on for the IBM Support Assistant. It is able to load multiple log files at the same time, and can display log content in either a flat chronologically-ordered view, or a hierarchical request-ordered view.

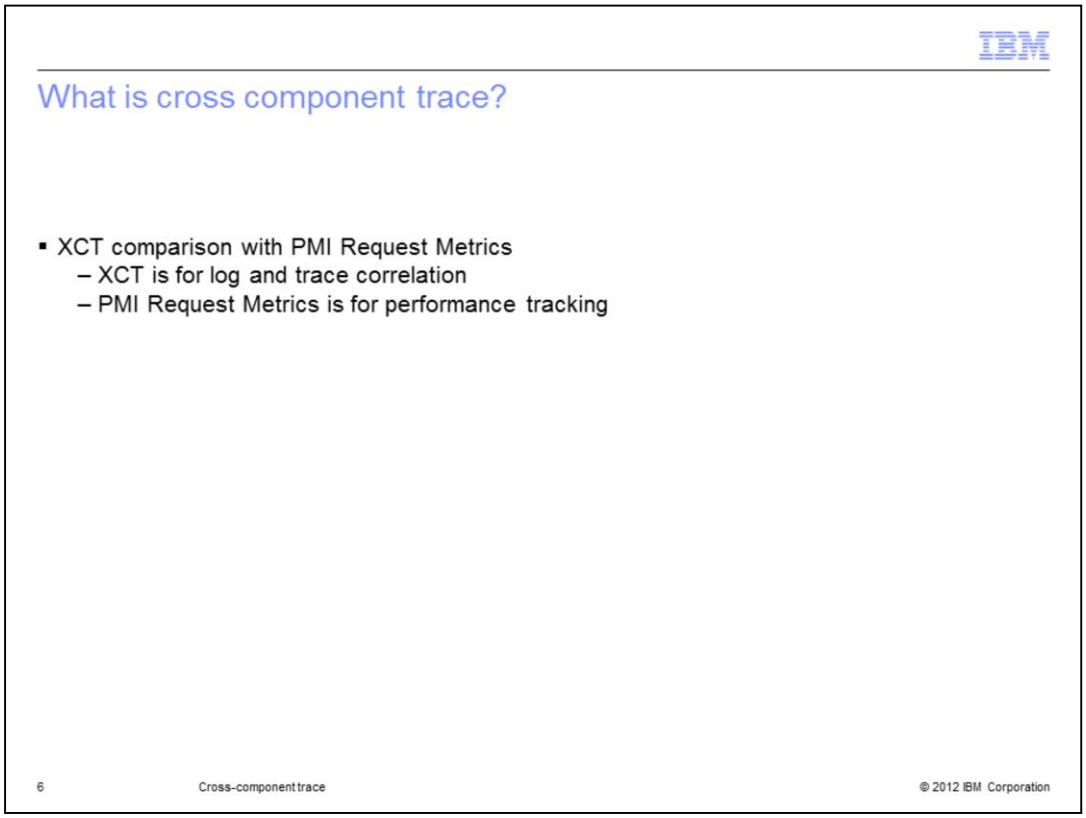

XCT and PMI Request Metrics have some overlap. Both technologies provide transaction tracking.

XCT is used for log and trace correlation, making it easy to see which log and trace entries are part of the same requests. XCT can also be used to capture request and response payload data. XCT instrumentation is sparse as it aims mostly to track where requests change threads or processes.

PMI Request Metrics is used for performance tracking. PMI Request Metrics can expose its data to Application Response Measurement (ARM) agents. PMI Request Metrics instrumentation can be verbose as it tracks the performance of individual components.

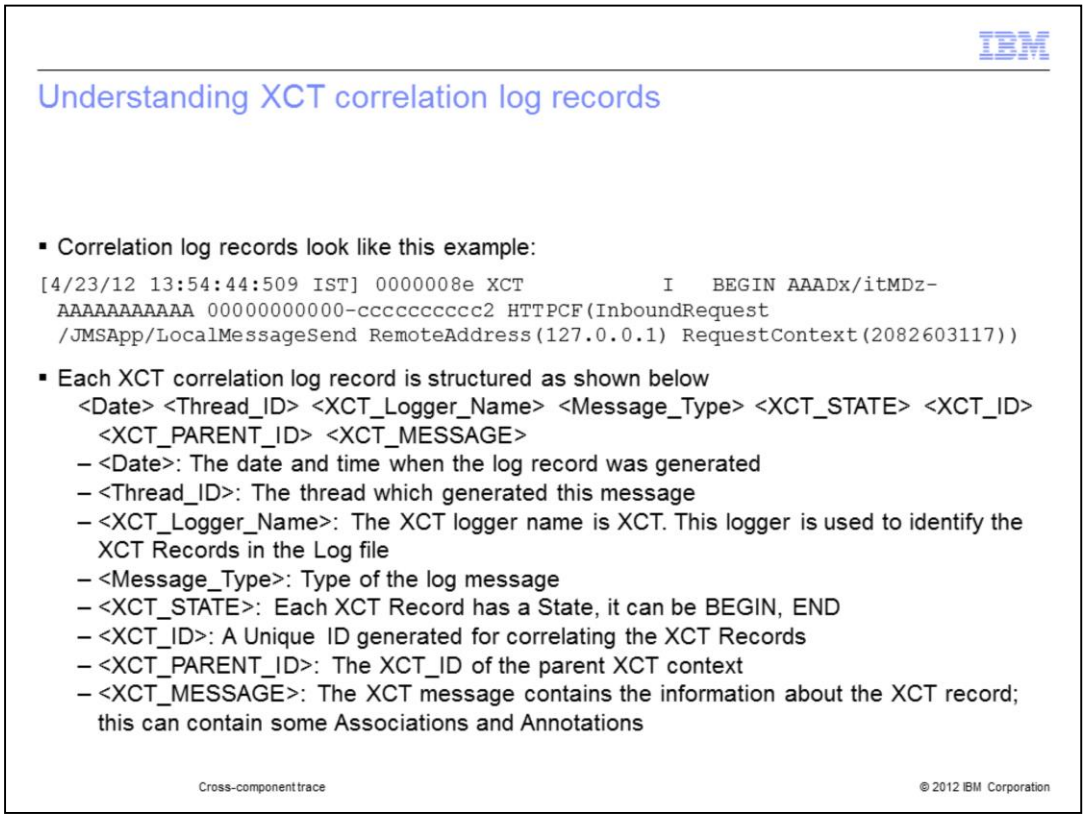

Each XCT record contains a date, thread ID, XCT logger name, message type, XCT state, XCT ID, XCT Parent ID, and message.

The date is the date and time when the message was generated.

The thread ID is the thread which generated the message.

The XCT logger name is the logger name in XCT, the logger is used to identify the XCT records in the log file.

The message type is the type of the log message.

The XCT state identifies whether an XCT context is beginning or ending. Think of an XCT context as a request running on a particular thread.

The XCT ID is the unique ID generated to correlate the XCT records.

The XCT parent ID is the XCT ID of the parent request.

The message contains the information about the XCT record; this can contain Associations and Annotations.

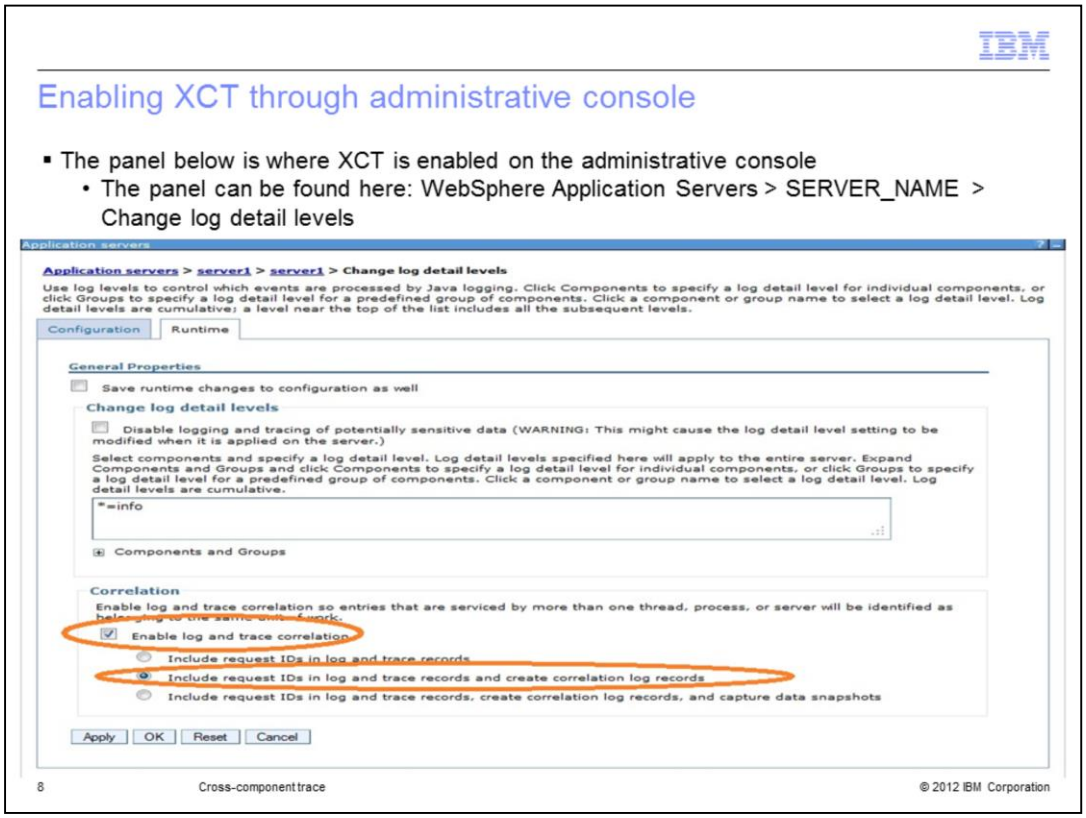

Cross Component Trace (XCT) can be enabled using the administrative console or WSADMIN scripts. In this illustration XCT is enabled using administrative console.

To enable XCT for the server , server1, navigate as follows: Servers > Server Types > WebSphere Application Servers > server1 > Change log detail levels > Runtime Tab

Select the check box that says Enable log and trace correlation and select the radio button labeled Include request IDs in log and trace records and create correlation log records

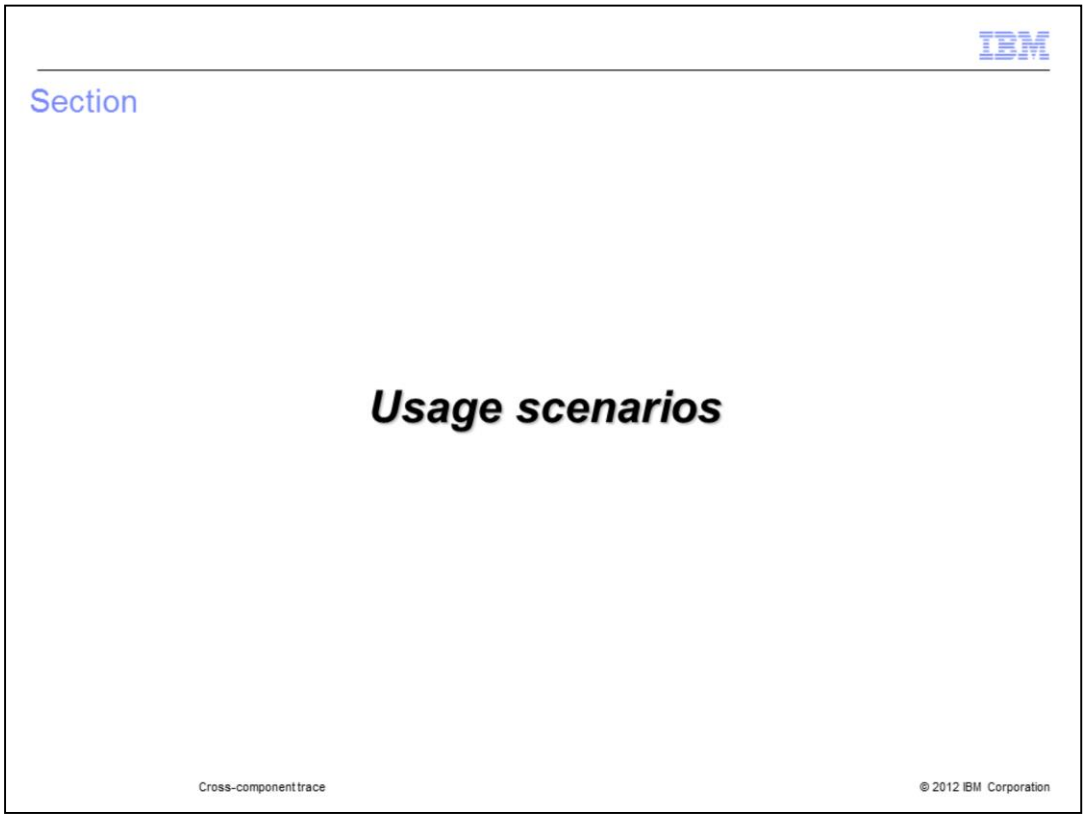

Cross component trace is used in these scenarios.

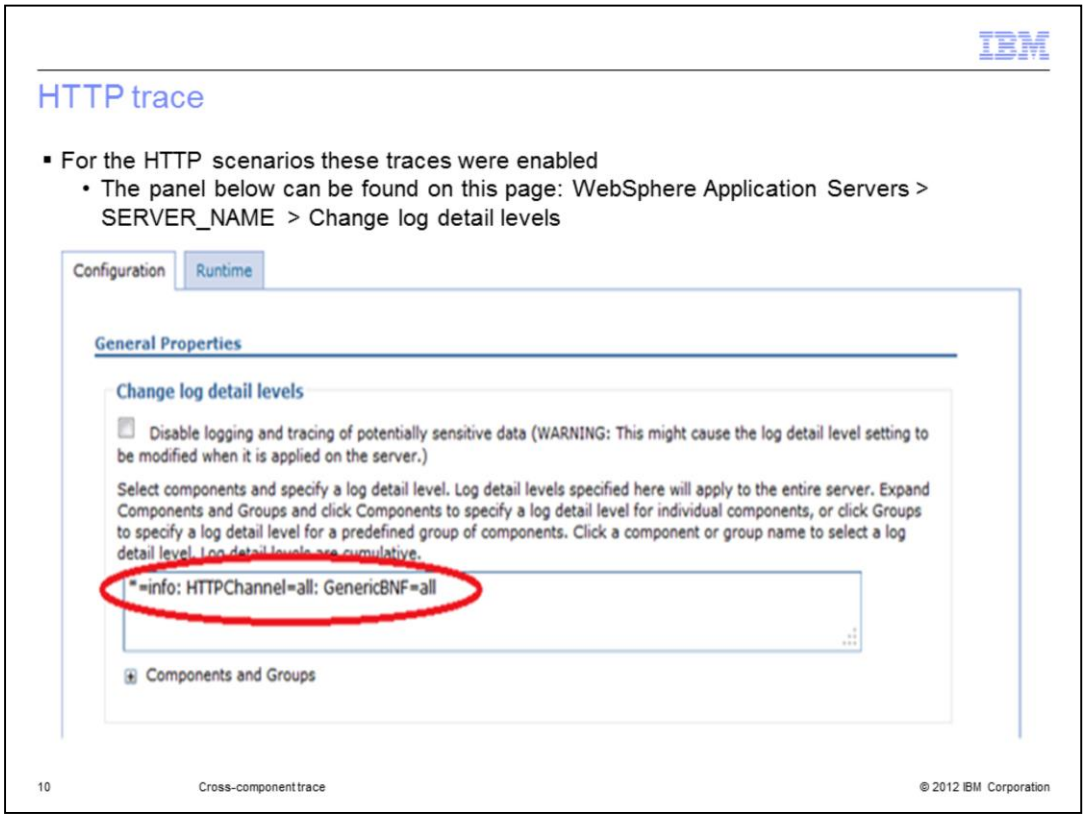

For the HTTP scenarios to follow this trace was enabled:
*=info:HTTPChannel=all:GenericBNF=all

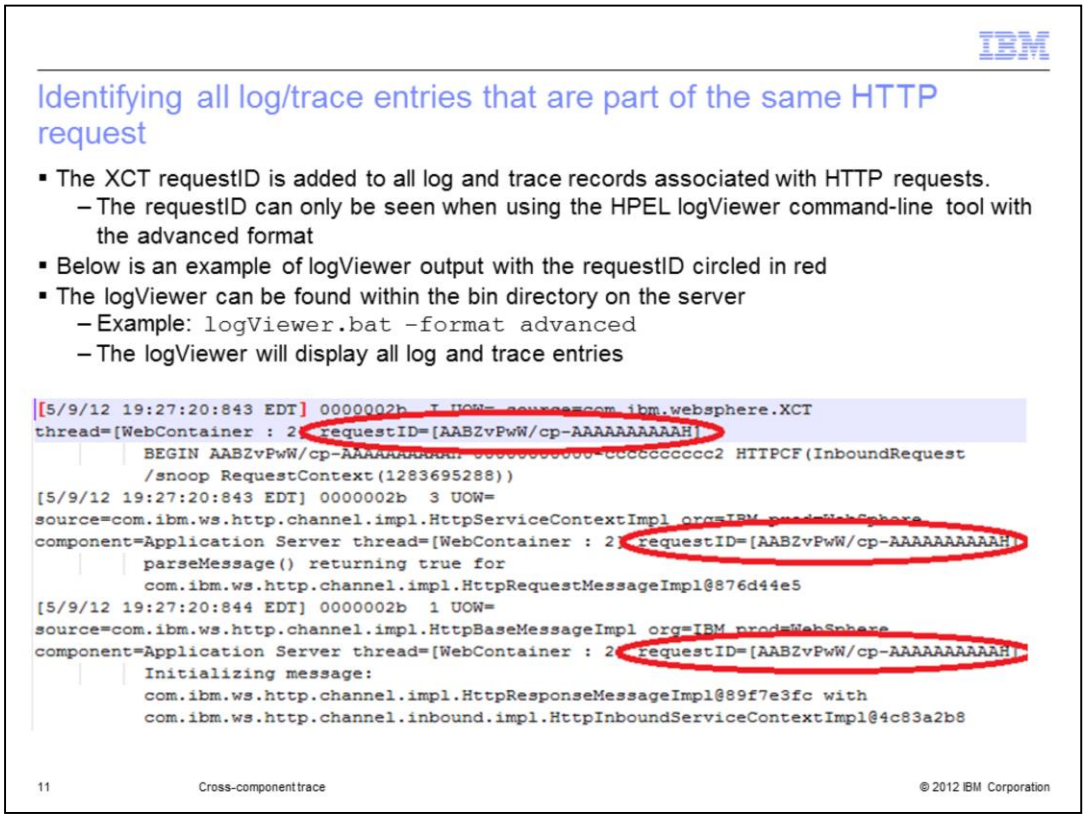

An administrator might want to use XCT to identify what trace entries are part of an HTTP request. To accomplish this the HPEL logViewer command-line tool can be used. This tool is found within the bin directory of the server. When the “logviewer” is run with the advanced format option, the requestID can be seen on each trace entry. In order to have the requestID present XCT must be enabled.

To search for log and trace records that match a particular requestID, use the command logViewer -includeExtensions requestID=<some Id>.

For example:

logViewer -includeExtensions requestID=AABZvPwW/cp-AAAAAAAAAAH

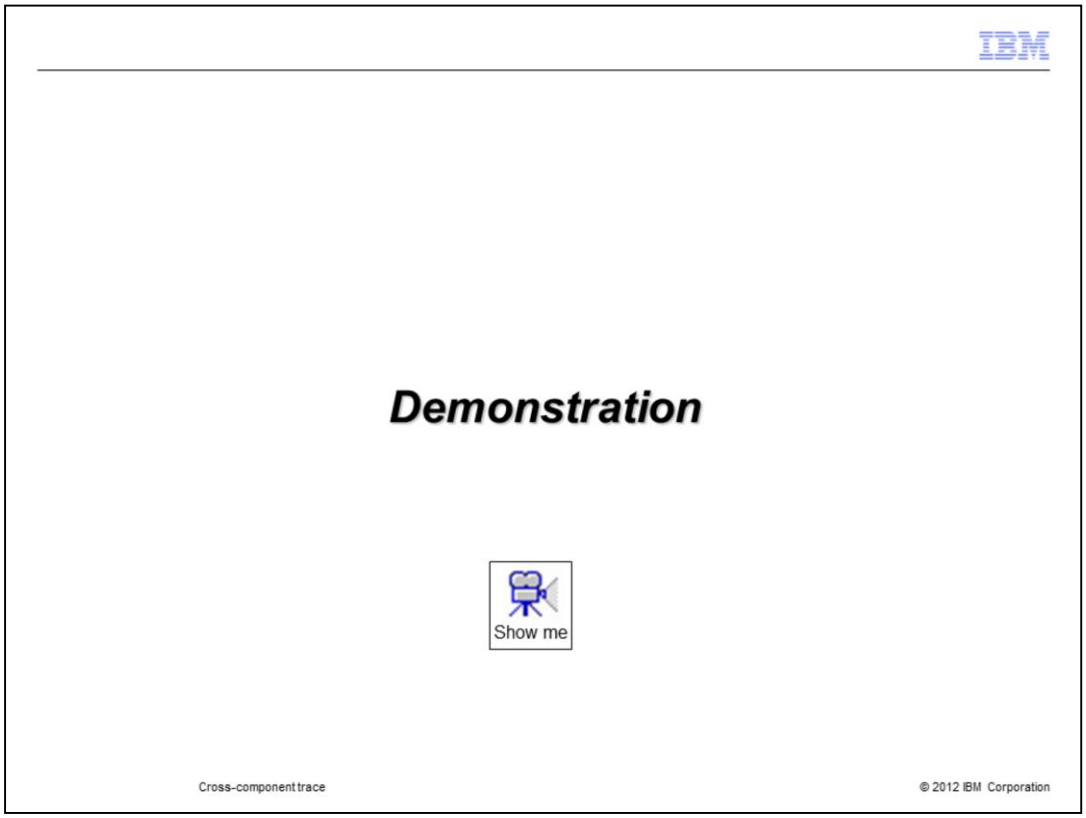

For a demonstration of how to Use request IDs to see log and trace entries related to a particular request, pause this presentation and click the demonstration icon.

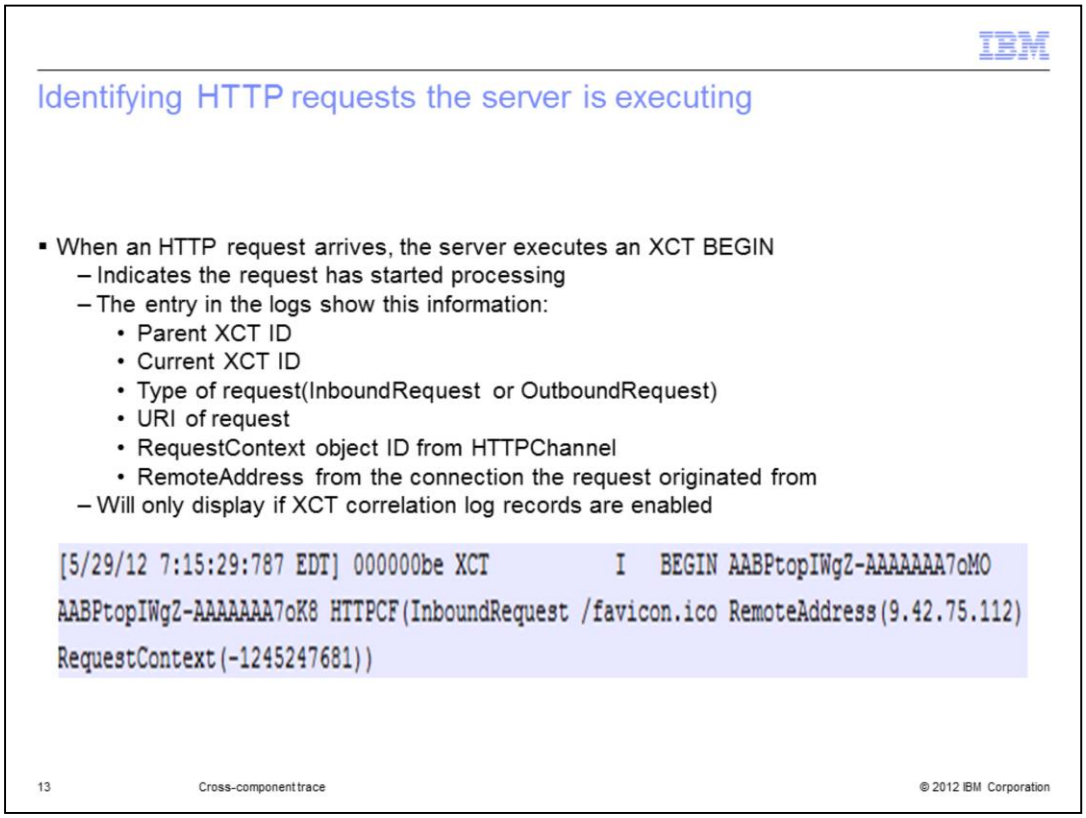

An administrator can see what HTTP requests the server is running. When the request arrives, the server does an XCT BEGIN. This marks the request as having started processing. In the logs an XCT BEGIN for an HTTP request can commonly be seen with attributes showing the XCT Parent ID, XCT Current ID, the type of request, which includes InboundRequest and OutboundRequest, the URI of the request, the HTTP Channel RequestContext object ID, and the Remote IP Address from the connection the request originated from.

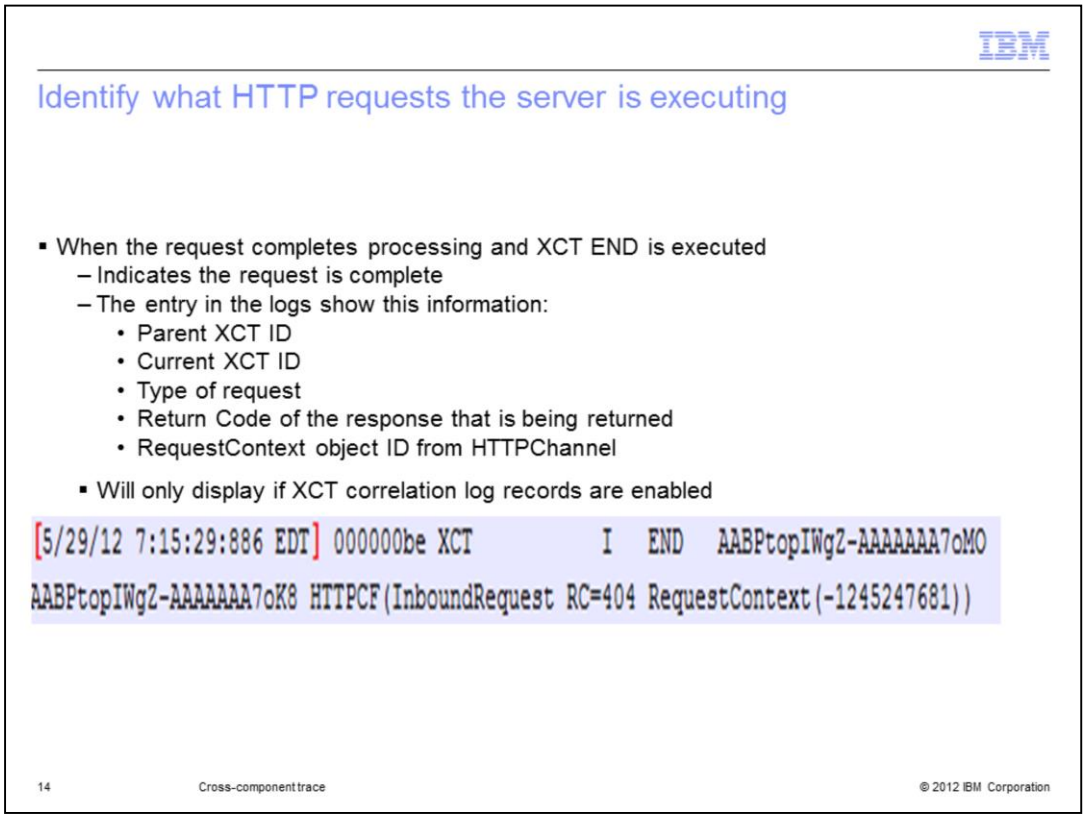

When the request completes processing the server does an XCT END, this marks the request as finished. In the logs an XCT END for a HTTP request can commonly be seen with attributes showing the XCT Parent ID, XCT Current ID, the type of request, which includes InboundRequest and OutboundRequest, the return code of the response, and the HTTP Channel RequestContext object ID.

IBM WebSphere cross component trace log viewer

- The scenarios following this slide use the IBM WebSphere Cross Component Trace Logviewer – available as a tool add-on for the IBM Support Assistant
- Tool used to examine XCT entries in a log
- Logs can be loaded from multiple servers and they are stitched together

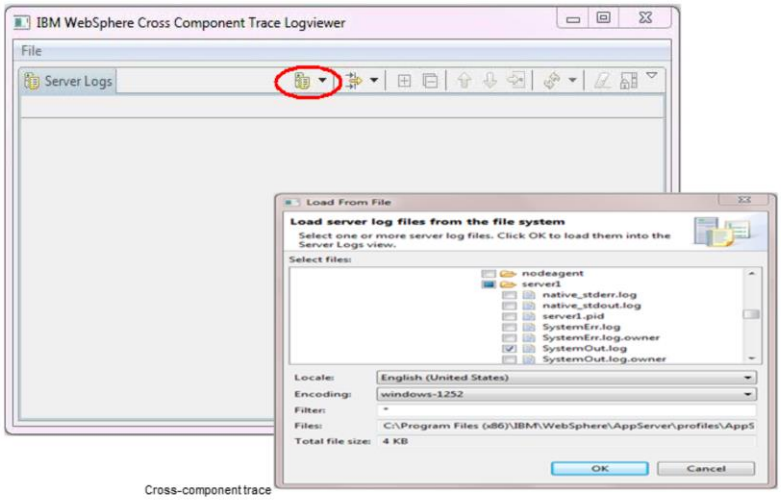

© 2012 IBM Corporation

The scenario that follows uses the IBM WebSphere cross component trace log viewer. The tool is used to examine XCT entries from a server log. Log files from multiple servers can be loaded and they are stitched together for a combined view.

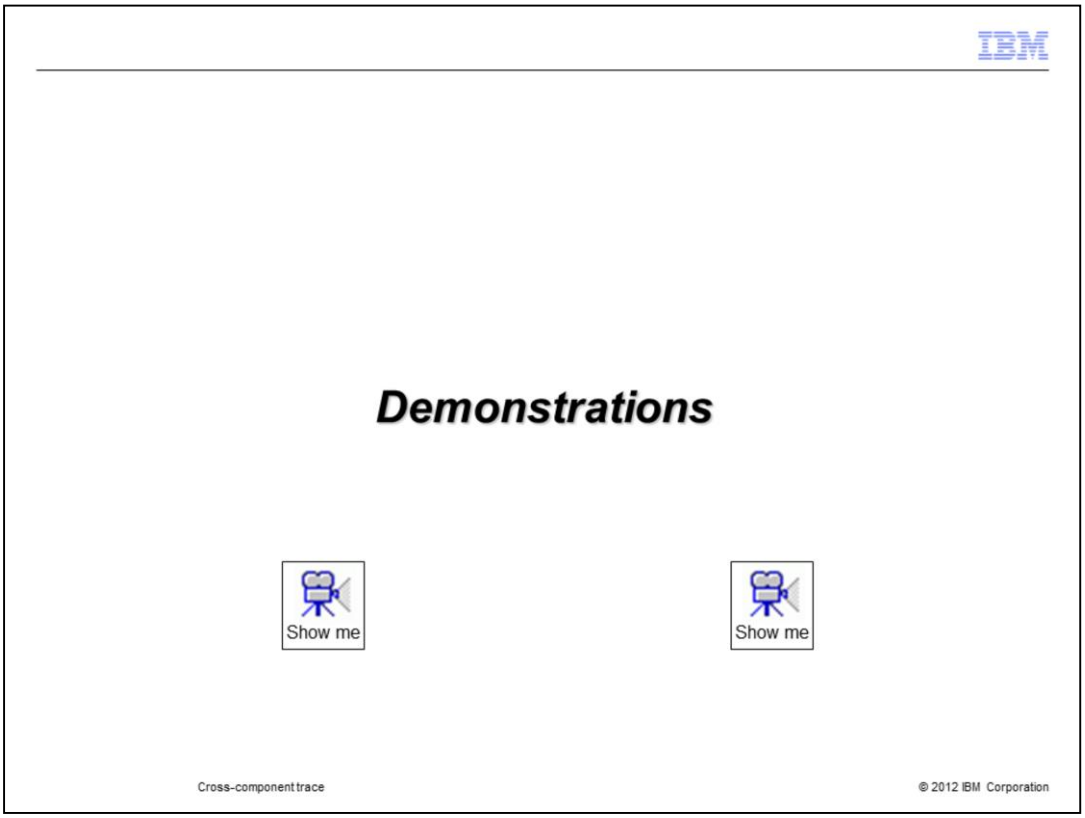

For demonstrations on how to get the cross-component trace log viewer and on how to see request hierarchy in the log view, pause this presentation and click each demonstration icon.

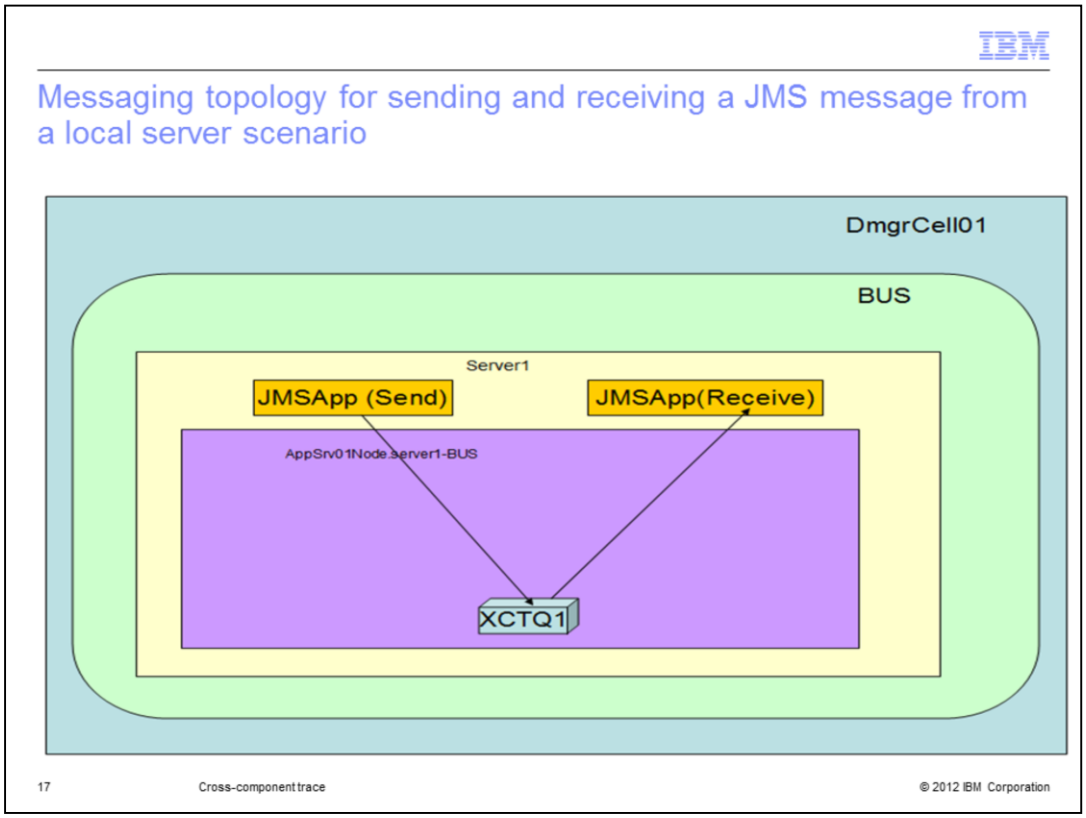

In this scenario, the JMS application and messaging engine are running in the same server process.

The JMS message is sent to a local queue destination and the message is received from the local queue destination synchronously.

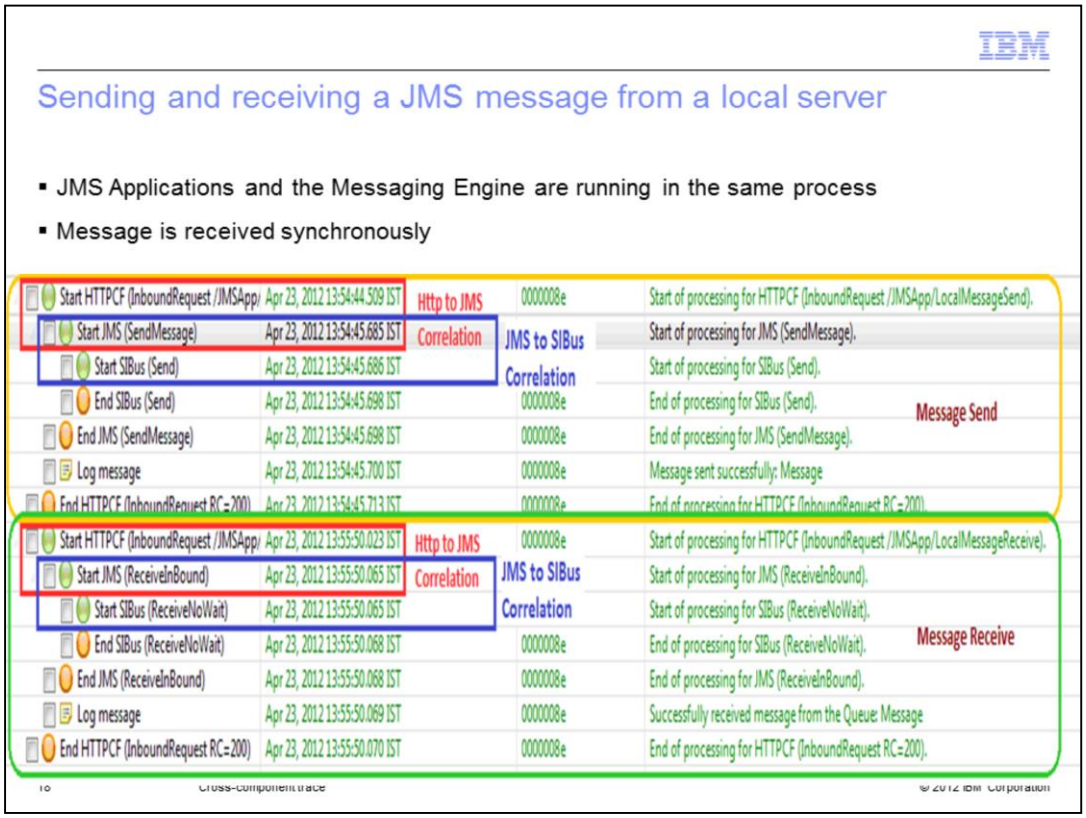

Since the JMS application and messaging engine are running in the same server process, the SystemOut.log from that server is loaded

HTTP to JMS Correlation and JMS to systems integration bus correlation can be clearly seen in the IBM WebSphere Cross Component Trace Logviewer

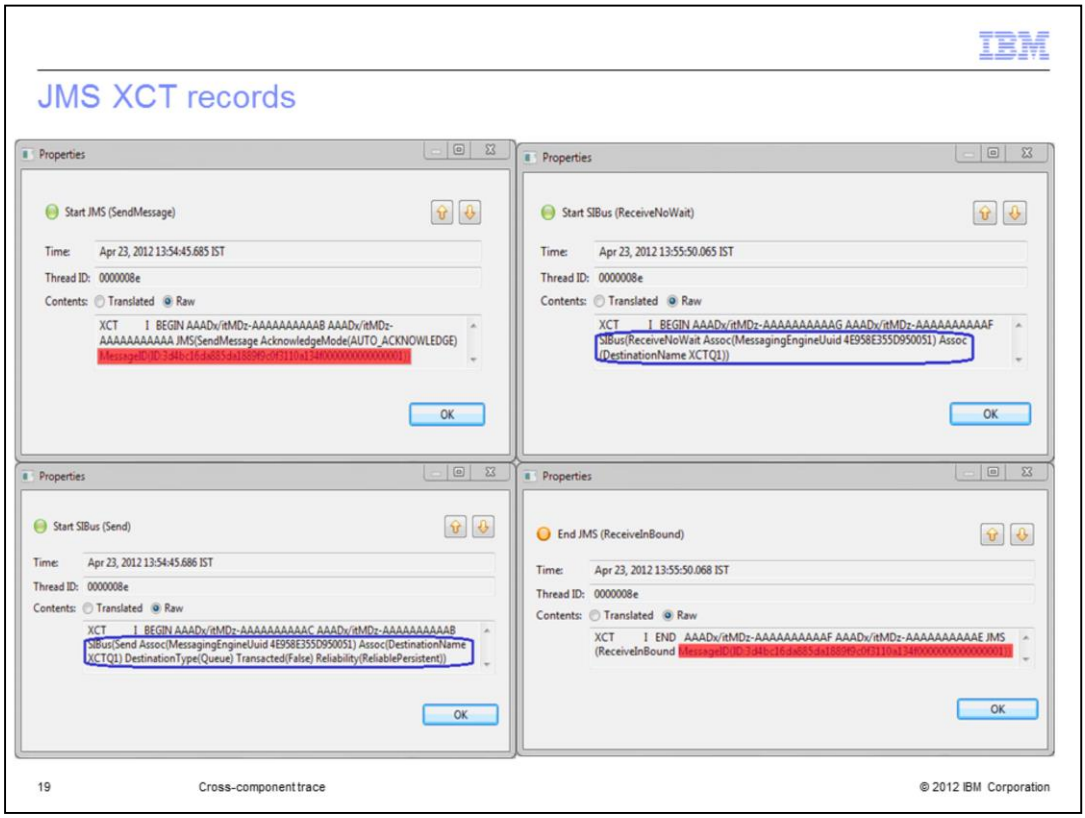

By double clicking the entry in the record list in the IBM WebSphere Cross Component Trace Logviewer the XCT records with annotations can be seen. In the JMS layer, the JMS Message ID is captured , which helps in correlating the message sent with the message received. In the systems integration bus layer, information related to the destination where the message is sent and from where the message is received is captured, such as messaging engine UUID , destination name etc.

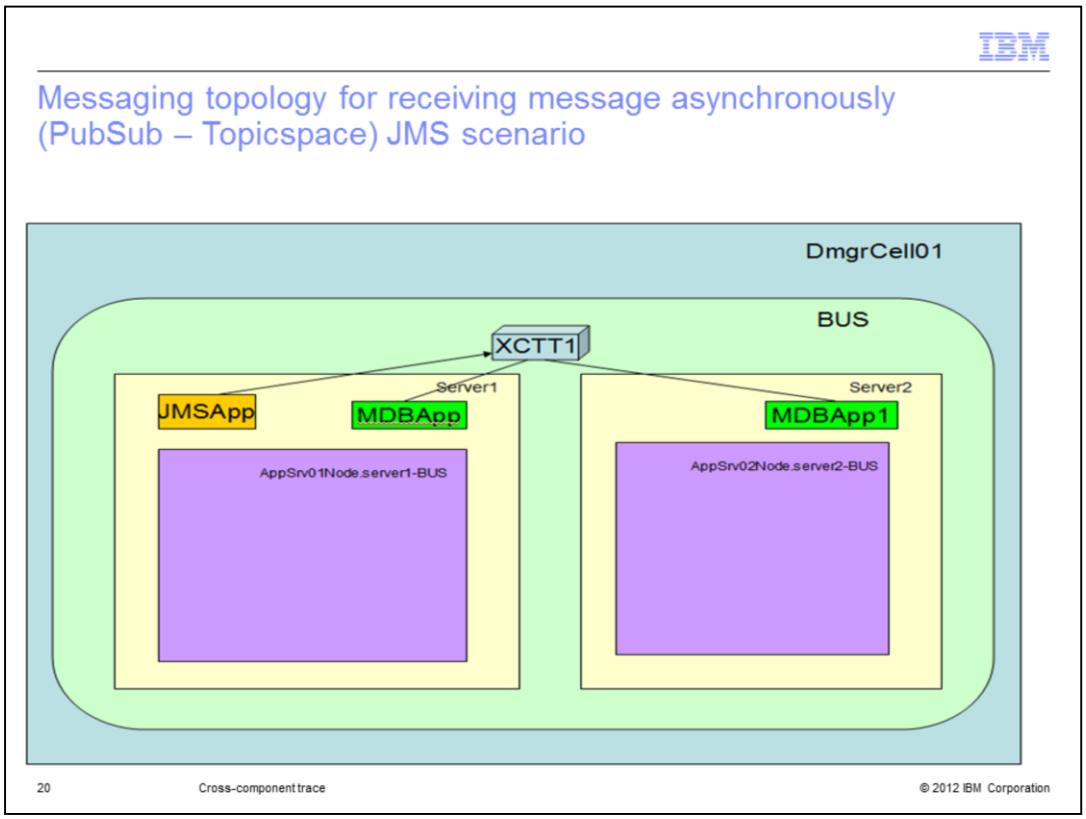

In this scenario, two servers are involved. Two MDB applications are deployed -- one in server1 and another in server2.

The JMS application running in server1 publishes a message to a topic that is subscribed to by the MDB applications running in server1 and server2

The message is asynchronously consumed by the MDB application.

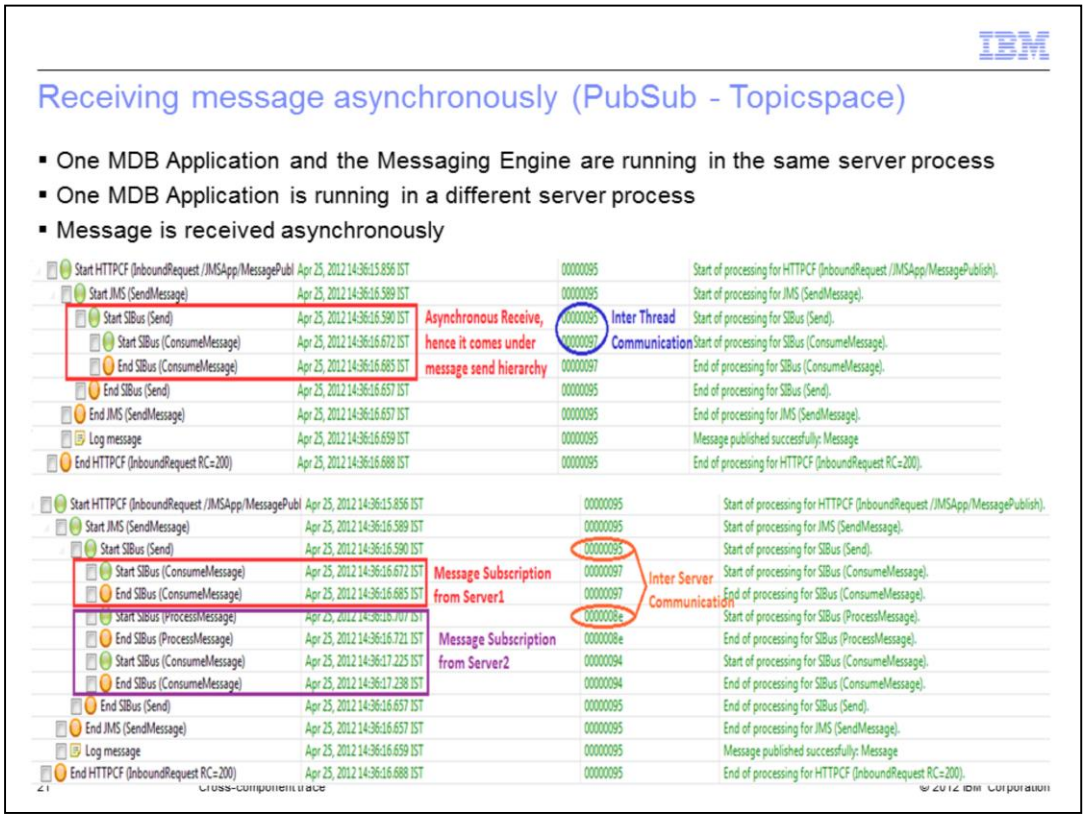

The SystemOut.log from server1 is loaded to show the inter-thread communication. Inter-thread communication is captured by XCT and displayed in IBM WebSphere cross component trace log viewer, where the JMS application and MDB application are running in the same server process.

The SystemOut.log from server1 and server2 are loaded to show the Inter server/process communication

Inter server/process communication is captured by XCT and displayed in the IBM WebSphere cross component trace log viewer, where the JMS application and the MDB application are running in different server processes.

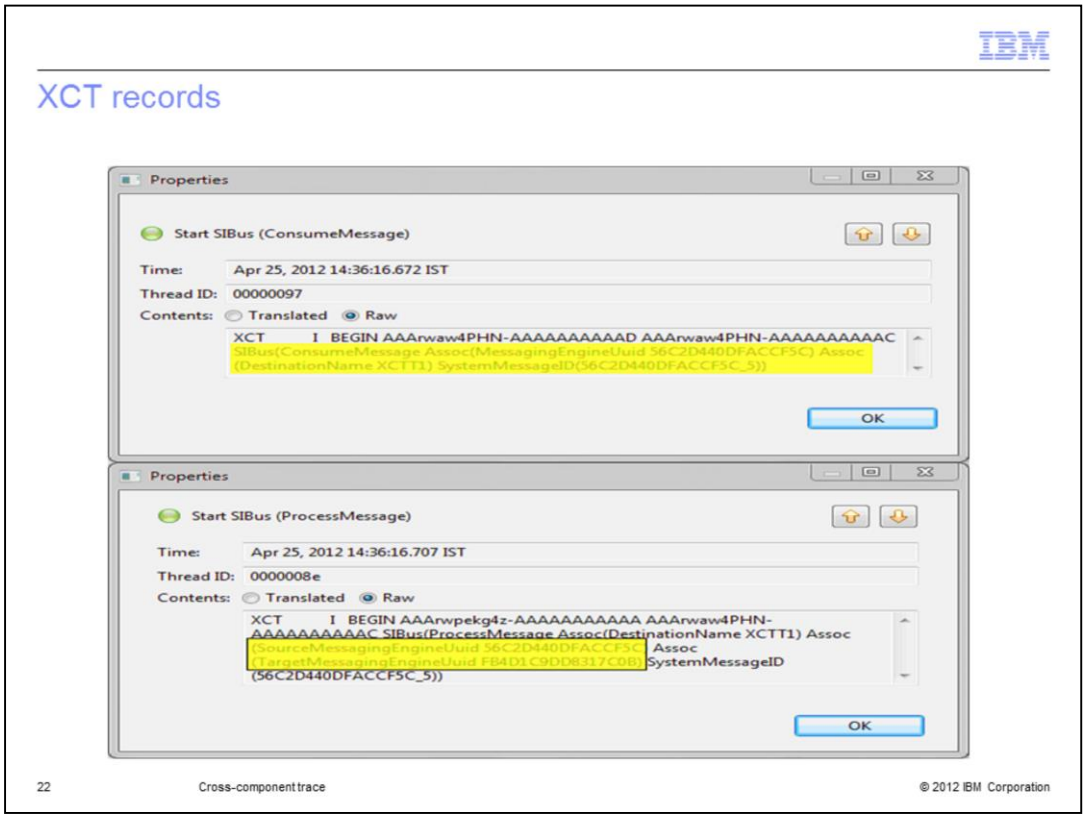

The top image shows the message being consumed by the first server, where the JMS application and the MDB application are running in the same server process with the messaging engine. The messaging engine UUID, the destination name from where the message is consumed and the system message ID are captured.

The bottom image shows the message being processed by the second server, where the MDB application runs remotely from the JMS application. The messaging engine UUID of both the source and target messaging engines, the destination name where the message is consumed, and the system message ID are captured.

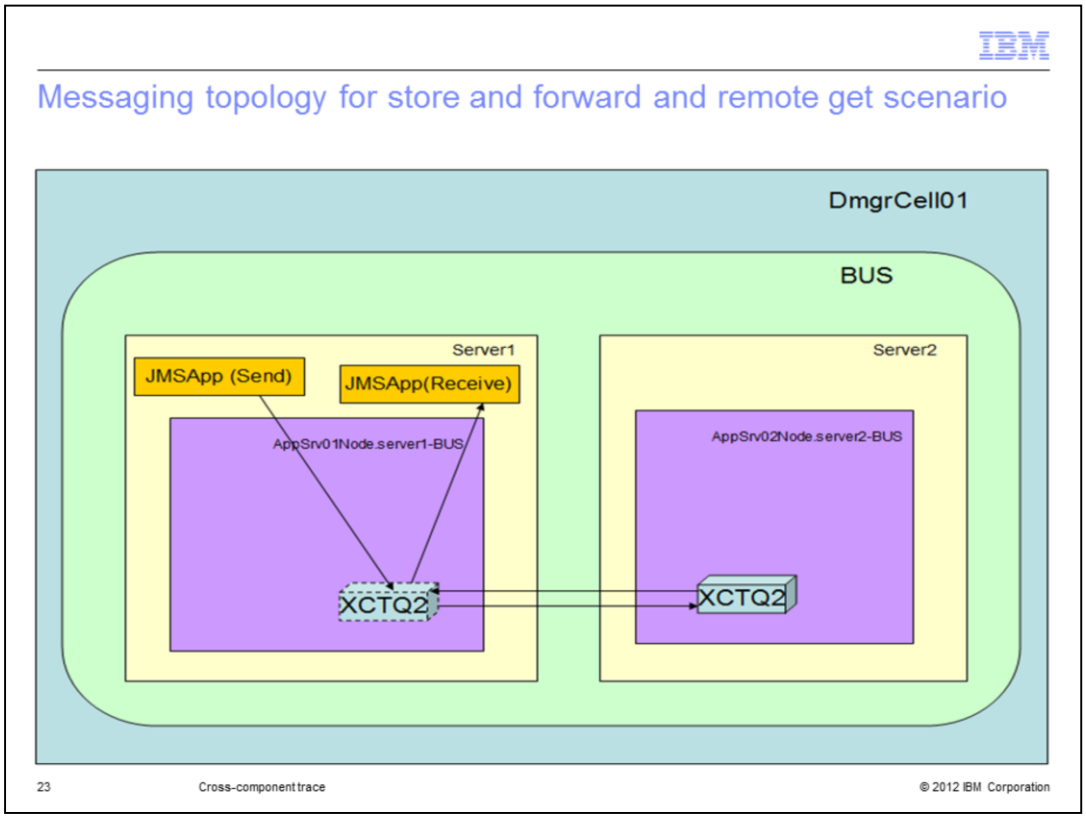

In this scenario, the JMS application deployed in server1 sends a message to a destination in server2 and receives a response message from that remote destination

The message is sent from server1 to server2 and a response is sent back from server2 to server1

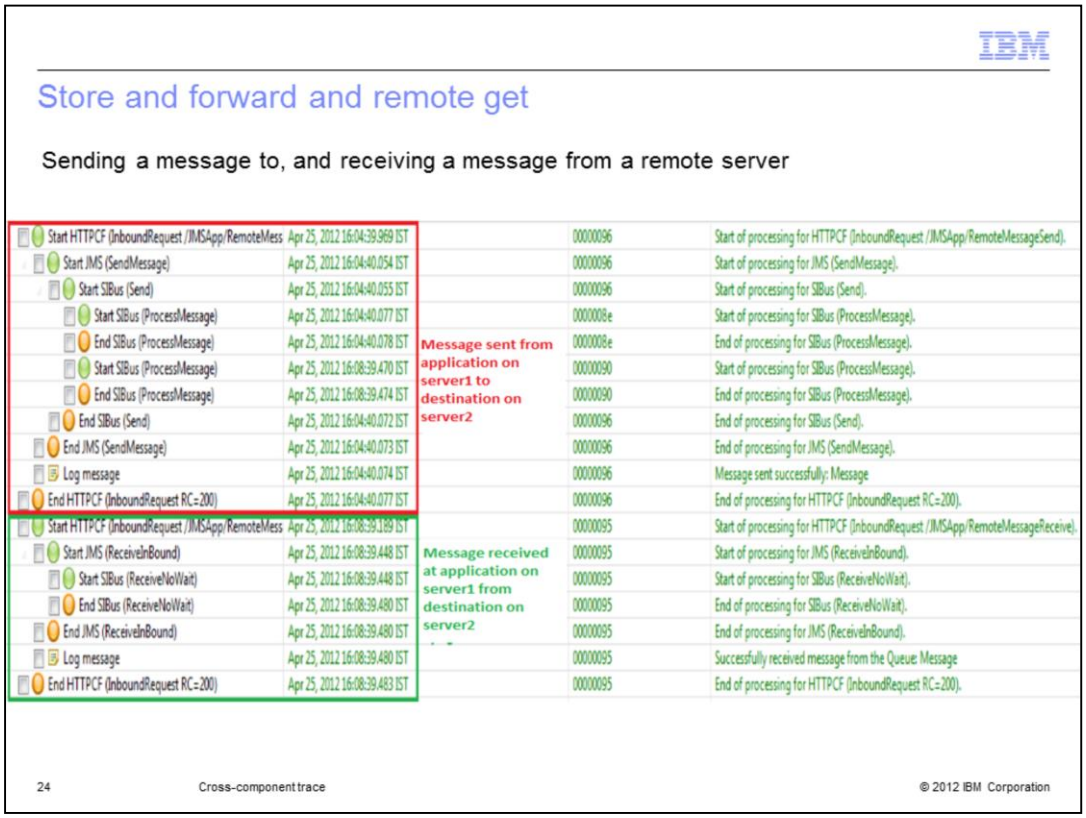

In this case, two servers are involved.

The SystemOut.log from server1 and server2 are loaded to show the message flow from server1 to server2 and vice versa

The source and target messaging engines UUID's are captured by XCT

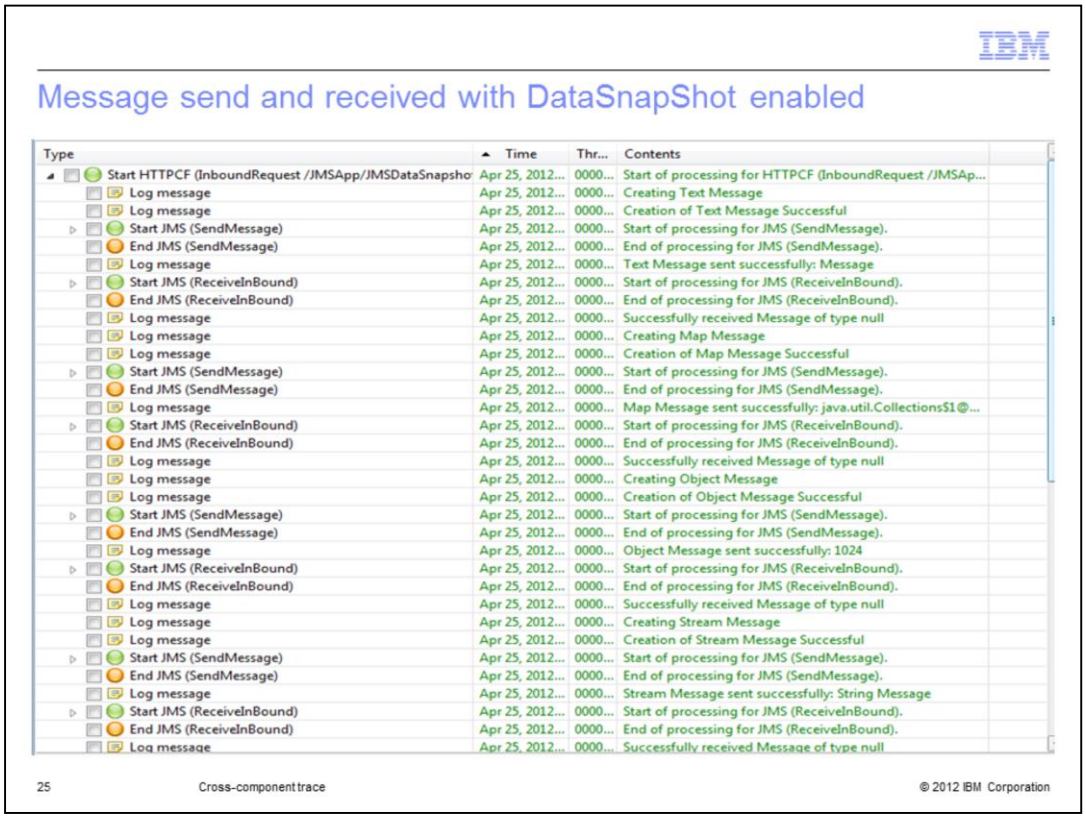

In this scenario, a message is sent to and received from a local queue destination with the XCT Data Snapshot option enabled.

When the message is sent and received the message data is stored in a file under the snapdata directory which is typically found under the server log root.

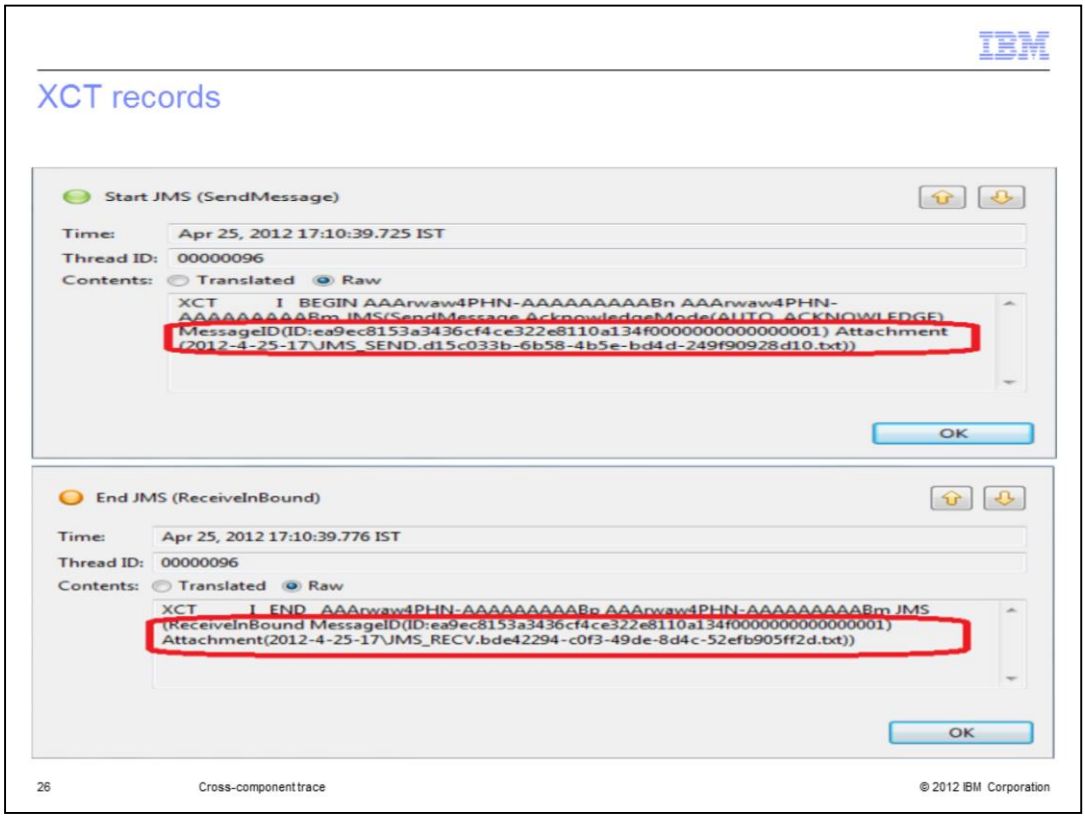

The XCT log record has the information on the attachment created

Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, Current, and WebSphere are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE 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, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2012. All rights reserved.

© 2012 IBM Corporation