Scenario 1: "The missing message problem"

A large clearinghouse uses a request-reply computing model for a major application. This is a three-tier client/server system with a queue manager involved at every tier. An alarm is generated when a reply does not arrive within the expected time. This happens every now and then for a multitude of possible reasons. Internally this is known as "the missing message problem". The standard procedure when the alarm goes of is to investigate what happened with the request-message, or its corresponding reply-message. The reply-message's correlation-id is the message-id of the request-message. The systems administrator can see if the message is stuck in one of the transmission queues, system dead letter queues, application dead letter queues, intermediate queues, or target queues. He does this manually using WebSphere MQ-browser utilities. It is a labor-intensive task, made more difficult because the message might still be on the move during the investigation.

The workload of the application is growing significantly and the IT department decides to implement multiple instances of the second- and third-tier queue managers. They decide that the easiest way to do this is to use WebSphere MQ clustering, which has the added benefit of workload balancing and high availability. But now the way the missing message problem has been handled becomes impractical, because it is impossible to know which of the many available routes the message took.

With ReQuest™'s multi-machine message propagation report, the current whereabouts of the message can be found immediately, plus the complete history of the journeys of the request-message and the reply-message.

Cressida Technology Ltd.
1 Lammas Gate, 84a Meadrow,
Godalming, Surrey GU7 3HT UK
Tel: +44 14 83 23 93 00
Fax : +44 14 83 23 93 83
Email: info@cressida.info

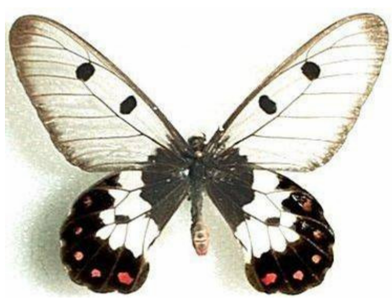

Scenario 2: "Decrease the overhead of the API-wrapper"

Three years ago a leading European bank started to use WebSphere MQ as the strategic asynchronous messaging middleware platform on all their operating systems (z/OS, AIX and Windows/NT). They quickly realized that information on every MQPUT operation needed to be logged to facilitate application restart, accounting and reporting. The system programmer who was asked to investigate soon discovered that the information they needed was not readily available; the log format is not documented, is dependent on the platform, and is not easy to read. So the bank decided that every MQPUT call had to go via one subroutine - an API-wrapper - and that this would log all the necessary information. Writing and maintaining the API-wrapper on all platforms was very costly and slowed the performance of WebSphere MQ, but in the end it all worked. With the exponential growth of WebSphere MQ applications, the bank is now faced with a tremendous overhead caused by the API-wrapper and looking for solutions.

The bank evaluated ReQuest™ and discovered that, for the persistent messages, it could produce all the information the API-wrapper gave them. About 90% of the messages are persistent. ReQuest™ runs can be scheduled at low-activity hours, and, more importantly, ReQuest™ can run on another machine, all it needs is the logs, or a copy of them. After some study the bank decided to make all messages persistent, and get rid of the logging from within the API wrapper altogether.

Cressida Technology Ltd.
1 Lammas Gate, 84a Meadrow,
Godalming, Surrey GU7 3HT UK
Tel: +44 14 83 23 93 00
Fax : +44 14 83 23 93 83
Email: info@cressida.info

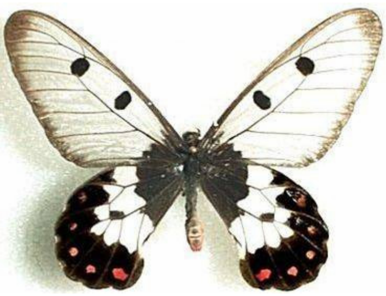

Scenario 3: "Use of timestamp recoveries in Service Level Agreements"

A large data center maintains the data processing environment of its customers, most of which are independent companies within the same group. The data center has agreed on a number of Service Level Agreements with the customers. One of these is related to disaster recovery, where the complete data processing environment is rebuilt on an alternative site.

This SLA states "We guarantee that, in the case of a disaster recovery situation, all data in the databases we manage will be recovered within 12 hours to the state they were in at 23:00 the previous day."

To achieve this SLA in the high-availability environment where off-line backups can not be easily considered and scheduled, the data center currently uses DBMS-specific techniques or third party tools to establish valid recovery points, and to do the time-stamp recovery. However, it turns out that the recovery of WebSphere MQ queues is also critical.

WebSphere MQ was first used for simple asynchronous communication between applications. Today's systems administrators realise that WebSphere MQ is not just a high-level communication protocol, but must be thought of as a database manager in its own right with business critical information stored in local queues. Applications reading and updating the databases also put messages on and get messages from queues, often in the same unit of work. Some database updates are driven by information on queues, whilst in other applications the content of the messages being put on the queues depends on information in the databases. In other words, if the databases are recovered to 23:00 yesterday, the local queues must be recovered to the same time, or you have an inconsistent situation with unpredictable results.

Since WebSphere MQ does not provide a timestamp recovery feature itself, the data center uses the only tool that does: ReQuest™ for WebSphere MQ allows the data center to detect valid recovery points and to recover the local queues to a state that is consistent with the DBMS.

Cressida Technology Ltd.
1 Lammas Gate, 84a Meadrow,
Godalming, Surrey GU7 3HT UK
Tel: +44 14 83 23 93 00
Fax : +44 14 83 23 93 83
Email: info@cressida.info

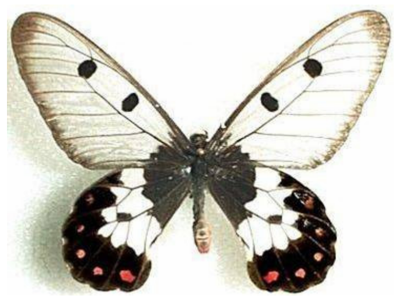

Scenario 4: "Auditing"

A UK based mail-order company has erroneously sent out 10, 000 tooth brushes to one of its customers. This shipping is driven by a WebSphere MQ application that reads orders from a queue. The application owner suspects either human error by one of their personnel, or a bug in an application. He asks the WebSphere MQ administrator to find out which applications and - if possible - which userid, generated the order. This information is needed as quickly as possible, in case the problem occurs again. Unfortunately, there are many different applications, on different machines and different platforms that can generate orders. Most queue managers run continuously and some use circular logging, which makes the available tools for dumping log information useless. Furthermore, for the few queue managers he *can* produce a log dump for, the system administrator is faced with the very difficult task of finding the requested information amongst the unformatted data.

Fortunately the company now have the option to use ReQuest™'s message propagation report to search system wide for messages containing the string "ARTICLE=toothbrush, AMOUNT=10000". In a matter of minutes he could identify the machine, application, and userid that instigated the erroneous order.

Cressida Technology Ltd.
1 Lammas Gate, 84a Meadrow,
Godalming, Surrey GU7 3HT UK
Tel: +44 14 83 23 93 00
Fax : +44 14 83 23 93 83
Email: info@cressida.info

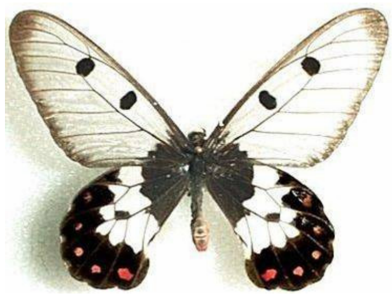

Scenario 5: "Use of timestamp recoveries for diagnostic purposes"

A major Insurance company runs z/OS with IMS/DB, DB2 and WebSphere MQ for its business critical applications. Since the introduction of IMS applications the data center team has received questions from the application owners such as "my batch application ran with the wrong input data; can you restore the database to just before the run?", or "yesterday around 2 a.m. an application generated some invalid data because of a bug and several accounts have been influenced. To correct the problem I need a copy of the databases as they where at 1:55 yesterday". In other words, requests for time-stamp recoveries. With IMS/DB this was relatively easy. The most difficult task was finding a valid recovery point, because the 'official' valid recovery timestamps were scarce. The 'let's-try-and-see' approach often worked; pick a time, do the timestamp recovery, run a database check, if that gave a return code of 0 or 4, everything was ok. Later, people found ways to detect valid recovery points from scanning the IMS log, which is easy to read and interpret. When DB2 came into the picture, the principle remained the same. True, logging of DB2 is not timestamp based, but RBA- based. This makes it a bit more difficult to handle the application owner's questions, but with some ingenuity - and possibly a third party software package - they could be addressed. For WebSphere MQ however, there are no such tools, and its recovery log is not easy to read.

Using ReQuest(tm) for WebSphere MQ the application owners will get the answers they look for. Using ReQuest™ the data center management team is able to agree on the following SLA with their departments: "... the data center guarantees that it can recover any application data in IMS/DB, DB2 or WebSphere MQ to a valid recovery timestamp not older than 90 days within a four hour period."

Cressida Technology Ltd.
1 Lammas Gate, 84a Meadrow,
Godalming, Surrey GU7 3HT UK
Tel: +44 14 83 23 93 00
Fax : +44 14 83 23 93 83
Email: info@cressida.info

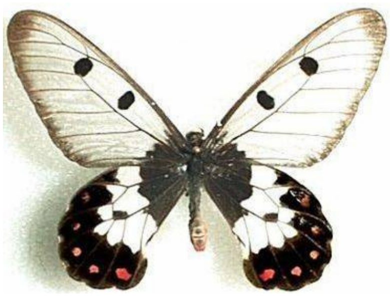

Scenario 6: "Replay for capacity planning."

The country's largest shoe distributor has recently made some major acquisitions. The IT department is proud of their inventory application that can be used by all offices and shops via a WebSphere MQ client connection. The IT director fears that the current hardware and software can not handle the expected workload growth. Furthermore, to support the new shops, some modifications to the WebSphere MQ applications are necessary. Although every application is easy to test in itself, they are all linked via WebSphere MQ and it is difficult to anticipate how they will interact in a production situation.

Both problems are addressed by ReQuest™'s replay functionality. The complete workload of messages coming in via the WebSphere MQ clients is read from the log, and replayed in a test environment. First this simulation run is used to test the new versions of the applications. Then, the rate at which messages come in is tripled to simulate the expected extra workload. Based on these tests the IT director may decide that the upgrade suggested by his hardware manufacturer could be postponed by at least a year.

Cressida Technology Ltd.
1 Lammas Gate, 84a Meadrow,
Godalming, Surrey GU7 3HT UK
Tel: +44 14 83 23 93 00
Fax : +44 14 83 23 93 83
Email: info@cressida.info

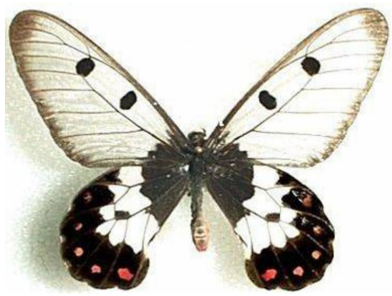

Scenario 7: "Performance of messaging applications"

The Ministry of Information is working on an e-government application that will allow every person to query on-line the status of various administrative dossiers. The application is WebSphere MQ based, and generally a query will involve multiple programs running on different platforms. The IT department is very committed to SLA management and has a tradition of detailed and segmented response-time reporting for the older, monolithic applications. They are now faced with the challenge of doing the same for multi-platform applications based on asynchronous messaging. The problem with trying to segment response time is that you need a way to correlate the occurrence of a message on one system, with the corresponding occurrence of a message on another system. For example, when message A is put on a transmission queue in QMGR1, and later on a target queue on QMGR2, how do you correlate these two events? The problem becomes more complex if one application, in order to satisfy one request, spawns multiple request messages, and waits for all reply messages to come back. Every request message can take a different route to a server application. Every server application generates a reply message, which takes some route to the reply queue.

Continued, see next page

Cressida Technology Ltd.
1 Lammas Gate, 84a Meadrow,
Godalming, Surrey GU7 3HT UK
Tel: +44 14 83 23 93 00
Fax : +44 14 83 23 93 83
Email: info@cressida.info

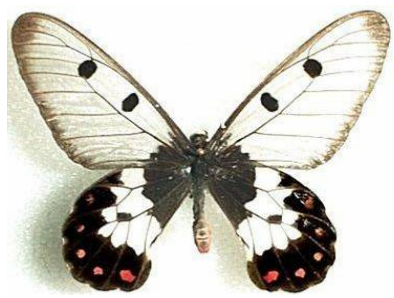

Continued: Scenario 7: "Performance of messaging applications"

They rejected solutions based on instrumenting the applications, because of the overhead, cost of implementation and maintenance, and lack of flexibility. They looked at ReQuest™ for WebSphere MQ, as it is the only tool that allows for cross queue manager, cross platform, message level reporting. Unfortunately ReQuest™ could not provide the reporting they needed out-of-the box. Since the message-id and correlation-id field in the message header was used for a specific purpose, and not unique, the tool had difficulty correlating the messages. The ministry claimed that messages could be correlated based on message contents. With this information Cressida Technology personnel showed that the tool was " open and user-extensible" and were able to produce the required, fully documented raw data. They created a segmented response time report that fulfilled the requirements of the ministry. The data is now also used in their MICS data warehouse and in the Candle Command Center.

Cressida Technology Ltd.
1 Lammas Gate, 84a Meadrow,
Godalming, Surrey GU7 3HT UK
Tel: +44 14 83 23 93 00
Fax : +44 14 83 23 93 83
Email: info@cressida.info