



Contribution ID: 47

Type: oral presentation

Building a robust distributed system: some lessons from R-GMA

Wednesday, September 5, 2007 3:40 PM (20 minutes)

R-GMA, as deployed by LCG, is a large distributed system. We are currently addressing some design issues to make it highly reliable, and fault tolerant.

In validating the new design, there were two classes of problems to consider: one related to the flow of data and the other to the loss of control messages. R-GMA streams data from one place to another; there is a need to consider the behaviour when data is being inserted more rapidly into the system than taken out and more generally how to deal with bottlenecks. In the original R-GMA design the system tried hard to deliver all control messages; those messages that were not delivered quickly were queued for retry later. In the case of badly configured firewalls, network problems or very slow machines this led to long queues of messages, some of which were superseded by later messages that were also queued. In the new design no individual control message is critical; the system just needs to know if each message was received successfully.

The system should also avoid single points of failure. However this can require complex code resulting in a system that is actually less reliable.

We describe how we have dealt with bottlenecks in the flow of data, loss of control messages and the elimination of single points of failure to produce a robust R-GMA design. The work presented, though in the context of R-GMA, is applicable to any large distributed system.

Primary authors: Mr PAVENTHAN, A (RAL); Mr KUSEJU, Adebisi (RAL); Mr DUNCAN, Alastair (RAL); Dr WILSON, Antony (RAL); Mr JIANG, Ming (RAL); Ms BHATTI, Parminder (RAL); Dr FISHER, Steve (RAL)

Presenter: Dr FISHER, Steve (RAL)

Session Classification: Grid middleware and tools

Track Classification: Grid middleware and tools