

EMI Messaging Requirements

Summary

In this document the Messaging Requirements from the EMI Product Teams are presented.

An in-depth analysis of these requirements, along with tests of the available broker technologies, will be performed later this year and reported in another document.

However, it is already clear that these requirements will be very hard to fulfill with a single monolithic broker cluster.

From 10,000 feet, a broker simply forwards data. How much data is forwarded is roughly the product of:

- the total messages rate (including messages received **and** sent by the broker)
- the messages size

The worst case scenario from the numbers reported below would represent 100 GB/s. Even the (guessed) average represents a significant amount of data in the 100 MB/s - 1 GB/s range, with very often short-lived connections over WAN.

It is already pretty clear that a messaging service for EMI should be split into separate dedicated services that could each be scaled according to the real needs.

Types of clients

Most applications only have simple clients that act either as a producer or as a consumer.

The most frequently used model is the "one-way reliable data transfer". Usually, there is a large number of producers and a small number of consumers. In some cases, the number of producers and the number of consumers are similar (same order of magnitude).

The second main model used is the "volatile information publishing" where information is lost unless consumers exist.

Clients are usually distributed across the WAN. Few are only on the LAN.

Main characteristics

Here are the main characteristics of how messaging would be used:

- Destinations number: from few destinations (~10) to ~500 destinations per application
- Message content: monitoring data, synchronization information, metadata
- Message body format: XML, YAML, JSON, LDIF, SMIME wrapper, JMS MapMessages
- Message lifetime: some applications want few seconds lifetime, others require few days
- Acceptable latency: at most a few seconds (for some applications)
- Connection lifetimes:
 - ◆ the most typical usage is: short-lived connections for producers and permanently connected consumers
 - ◆ some applications need permanently connected producers
- Security (on the broker):
 - ◆ some applications require authentication, with username+password and/or certificates
 - ◆ some applications require authorization to protect sensitive destinations

- ◆ some applications require SSL encryption to protect the data in transit

Aggregated numbers for all destination types

If we sum all the **known** requirements, here is what the brokers should be able to support:

- Producer instances: hundreds of permanent connections, thousands of short-lived connections
- Consumers instances: hundreds of permanent connections
- Messages rate: from ~10 msg/s to ~10K msg/s peaks with an average that could be ~1K msg/s
- Message size: from ~1KB up to ~1MB with an average that could be ~100KB
- Amplification factor: highly variable, hopefully usually in the 1-10 range

N.B. The amplification factor is the number sent messages divided by the number of received messages (from the broker's point of view). For a queue, it is normally one; for a topic, it is normally the number of connected consumers. This ratio is very important because, if it is big, it could lead to I/O bottlenecks.

Client programming languages

The following programming languages should be supported:

- C/C++
- Java
- Perl
- Python

Per product team requirements

- APEL client
- ARC
- CERN Data Management - LFC
- gLite Information System
- gLite Job Management
- Logging & Bookkeeping
- Service Monitoring Task Force

This topic: EMI > EMIMessagingRequirements

Topic revision: r3 - 2011-04-19 - LionelCons



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback