Fiorano
Enabling change at the speed of thought

www.fiorano.com

Message Driven SOA --- Enterprise Service Oriented Architecture

How well does your messaging infrastructure scale?

Flow Control, Latency, Scalability and Persistence

--- Atul Saini

AMERICA'S
Fiorano Software, Inc.
718 University Avenue Suite
212, Los Gatos,
CA 95032 USA
Tel: +1 408 354 3210
Fax: +1 408 354 0846
Toll-Free: +1 800 663 3621
Email: info@fiorano.com

EMEA
Fiorano Software Ltd.
3000 Hillswood Drive Hillswood
Business Park Chertsey Surrey
KT16 0RS UK
Tel: +44 (0) 1932 895005
Fax: +44 (0) 1932 325413
Email: info_uk@fiorano.com

APAC
Fiorano Software Pte. Ltd.
Level 42, Suntec Tower Three 8
Temasek Boulevard 038988
Singapore
Tel: +65 68292234
Fax: +65 68292235
Email: info_asiapac@fiorano.com

Entire contents © Fiorano Software and Affiliates. All rights reserved. Reproduction of this document in any form without prior written permission is forbidden. The information contained herein has been obtained from sources believed to be reliable. Fiorano disclaims all warranties as to the accuracy, completeness or adequacy of such information. Fiorano shall have no liability for errors, omissions or inadequacies in the information contained herein or for interpretations thereof. The opinions expressed herein are subject to change without notice.

HOW WELL DOES YOUR MESSAGING INFRASTRUCTURE SCALE?

Flow Control, Latency, Scalability and Persistence

Executive Summary

Messaging technologies have evolved in response to constantly changing corporate IT environments, which have shifted over the past three decades from mainframe systems to client/server applications to Web-based Internet applications. Most JMS-compliant message queuing (MQ) servers support multiple clients and address heterogeneous operating systems, databases and applications. But they fall short of fulfilling some basic *scalability* requirements imposed by a business's growing needs.

Good products support scalability to address growth requirements. Outstanding products support scalability without compromising the performance and reliability of the overall system. As business conditions necessitate the support of more concurrent users and/or higher throughput (measured in messages/second), the easiest recourse is to invest in more powerful hardware to drive your MQ performance. But while processors and memory costs trend downwards, net performance gains from adding more CPUs and memory do not scale linearly because of well-known bus and memory access latency limitations in multi-CPU-based hardware.

A software solution to address higher performance requirements involves distributing the workload across a *cluster* of MQ servers mounted as software on different machines. However, clustering MQ servers to meet performance and scalability requirements necessitates a closer inspection of the following issues.

Guaranteed Messaging

Guaranteed messaging ensures that messages are reliably delivered once, and only once, to their intended customers. Guaranteed messaging, traditionally a key requirement for financial and B2B supply chain markets, is increasingly a "must have" for most customers. Additionally, mobile clients need to be able to retrieve their messages on demand--as opposed to having to stay logged on and subscribed to a particular topic all the time. To support an on-demand delivery of messages, they need to be marked as *persistent.*

Persistent messages must be recovered in the case of an MQ or client failure, and the MQ server must provide the retrieve-on-demand flexibility discussed above. JMS-compliant MQ servers that support guaranteed delivery of persistent messages need to implement an offline storage mechanism for persisting messages to local disk or databases or across storage devices attached to a storage area network (SAN). This storage ensures message recoverability in the event of an MQ or client failure. Storage implementations range from flat files to relational databases to object databases.

On the other hand, messages such as stock quotes, which are refreshed periodically, are inherently transient in nature. However, they may also be guaranteed and, hence, necessarily stored for later retrieval. Relational/object databases designed to store huge blocks of data for long periods of time are suboptimal in their support of these smaller, more dynamic transient messages. File-based data stores designed specifically to address a broad spectrum of messaging requirements are a superior alternative to relational/object databases.

After ensuring that your MQ server meets guaranteed messaging requirements, the next two questions are:

- How reliably can the MQ implementation support guaranteed?
- Does supporting guaranteed messaging cause a significant performance?

In the event of an MQ failure within a cluster, all clients connected to the failed MQ server need to fail over to another live MQ instance with minimal downtime and without the loss of any messages. If each MQ instance has its own associated storage mechanism, then in addition to the transparent redirection of the clients connected to it, the associated storage needs to instantaneously fail over as well. Any delayed failover directly affects message availability. Real-time replication of stored messages to support failover implies process communication overheads between MQ servers in a cluster, which can degrade overall cluster performance. Shared data stores (based on standardized RAID disks or SANs supporting multiple MQ servers) minimize these communication overheads and are better than individual silos of MQ storage implementations.

In summary, look for MQ implementations that support guaranteed messaging without overall system performance degradation in a clustered environment. File-based data stores designed specifically for messaging requirements can offer significant scalability and performance advantages over database implementations.

Reducing Management Complexities

As multiple MQ servers become part of your cluster, management complexities must be well scoped out and controlled. The configuration of server-to-server communication routes between servers and global destinations based on server clustering implementations needs to be simple. Within each MQ server instance, configuring global/local queues and MQ cluster index can become a beast in itself unless managed well. So finding an MQ server with a rich set of management functions is imperative.

It should also offer an easy process for performing MQ software upgrades.
The chosen MQ server should support common industry security standards and basic security management features. As an integral component of your overall enterprise security system, the MQ server can become a dangerous vulnerability from a security perspective. To that end, configuring and managing all the access control lists (ACLs) to MQ cluster access, as well as managing security certificates, becomes paramount. MQ solutions ought to support standards-based security at the transport level as well (using 40-bit and 128-bit SSL encryption standards, for example).

Restrictive Control Flow Support

Most MS solutions have an ugly limitation that can become your worst nightmare. Assume that a publisher is sending stock quote messages
(at 400 messages/second) to an MQ server connected to 100 clients/subscribers. Most MQ servers utilize the machines' internal buffers to store the incoming messages. Subscribers, on the other hand, consume these messages at varying rates.

Each subscriber is limited in its consumption rate of the messages, depending on the implementation of the subscription application, hardware resources, network access rates and other factors. Typically, a slow subscriber can force the internal buffer in the MQ server to fill up quickly. Most MQ servers throttle back the publisher in order for the slowest subscriber to drain its internal memory buffers. In theory, the publisher could be brought down to a grinding halt which is unacceptable for most real life applications. Look for MQ products that provide flow control for managing internal buffers and external disks intelligently so that any slow subscriber still results in an acceptable level of drop in overall performance. The key is to ensure that all messages are delivered reliably, albeit at a slower rate.

Interoperability Minefields

It is a given that any MQ product has to integrate seamlessly with the existing ecosystem. Standards based support across multiple languages (C/C++, Java, Perl), operating systems (Windows NT, Solaris, AIX) and platforms (Java and .NET) is a fundamental requirement. Integration with existing application and LDAP servers and support for XML data can simplify your scalability goals.

Interoperability also implies a proven, seamless ability to work with other vendors' MQ products at connectivity, management and--most importantly--debugging levels. Standards-based MQ solutions can be built on and hence leveraged to solve much larger and more complex problems, such as enterprise application integration. For example, as Web services begin to get more standardized and available, your MQ product ought to integrate SOAP (Simple Object Access Protocol), a standard way to access Web services APIs so that SOAP messages between Web services can leverage the MQ transport services.

Look for ways to protect and extend your investment in MQ solutions by insisting that vendors integrate these interoperability features as part of their standard MQ clustering products.

Conclusion

Messaging technology began with a simple objective of transporting messages across the enterprise, and it followed the trend in computing from centralized mainframes to client/server to distributed computing. However, many MQ servers today do not stretch well to cover enterprise requirements of guaranteed message delivery, management ease of use, flow control, security and interoperability. To ensure that your messaging technology can scale with your business, pay careful attention to how your MQ server fulfills these critical needs.

About Fiorano Software

Fiorano Software (www.fiorano.com) is a leading provider of enterprise class business process integration and messaging infrastructure technology. Fiorano's network-centric solutions set a new paradigm in ROI, performance, interoperability and scalability. Global leaders including Fortune 500 companies such as Boeing, British Telecom, Credit Agricole Titres, Lockheed Martin, NASA, POSCO, Qwest Communications, Schlumberger and Vodafone among others have used Fiorano technology to deploy their enterprise nervous systems.