All Hands Conference, Nottingham 2005

Abstract

The requirements for collaborative services, especially pertaining to order and delivery, are quite different compared to traditional distributed applications. The NaradaBrokering (NB) messaging substrate enables scalable, fault-tolerant, distributed interactions between entities, and is based on the publish/subscribe paradigm. The substrate also incorporates support for Grid and Web Services. We have performed benchmark tests on machines in the USA, UK and Australia to gauge the performance and scalability of the NB substrate over worldwide networks using benchmarking software developed alongside NB and NB using GridFTP with varying message payloads and large attachments over Parallel TCP. In this paper, we demonstrate how our results show that the substrate can indeed be utilised to achieve maximum performance in a collaborative messaging environment over worldwide networks with the ability to maintain performance notwithstanding scalability constraints. The results will show some measurements for the Web Service implementations in NB and highlight the performance of SOAP/HTTP in these settings.

Extended Abstract

Message-Oriented Middleware (MOM) can be viewed as a middleware layer lying between the application and transport layers at both ends of a communication path. MOM virtualizes communication in distributed, heterogeneous environments, thereby increasing interoperability, portability, and flexibility. MOM supports asynchronous communication, typically through the use of message queues, which help insulate distributed applications from the impact of destinations that are busy due to high message traffic, or that are temporarily disconnected. MOM is well-suited for applications requiring event or notification services, and for service-based systems in general since it provides a conceptual model for communication between services. The MOM model facilitates loosely coupled, asynchronous communications which in turn enables complex interactions between participating entities. The notification service being developed at Southampton University for the MyGrid project and the iJob environment developed at San Jose State University by Fatoohi [1] may both be regarded as examples of MOM.

As the Open Grid Services Architecture (OGSA) gains increasing acceptance in the e-Science community, it is becoming apparent that MOM is an important aspect of the service-oriented architecture (SOA) that has not received sufficient attention. The main motivation for the research detailed in this paper is to demonstrate that the NaradaBrokering (NB) [2,3,4] MOM substrate developed at Indiana University could be incorporated into a SOA for grid computing and illustrate how it performs as a messaging infrastructure for the grid environment. A table that compares the main features of some leading messaging systems with NB is available at http://www.cs.cf.ac.uk/User/David.W.Walker/NB.pdf. (Updated Version??).

NB is a distributed messaging infrastructure, which provides support for centralized, P2P and distributed interactions. NB efficiently routes any given message between the originators and registered consumers of the message in question. Messages could be used to encapsulate information pertaining to transactions, data interchange, system conditions and finally the search, discovery and subsequent sharing of resources. As far as support for standards based messaging is concerned NB incorporates support for the Java Message Service (JMS) and WS-Eventing; support for the WS-Notification specification is currently being incorporated into the substrate. A list of services available within the substrate, supported subscription formats, and supported standards can be found at the project website[1].

The objective of the NB substrate is to provide a recognised open-source standard for a messaging infrastructure to build SOAs and Grids upon. The settings within which NB has been deployed successfully include audio/video conferencing, distance education, remote visualizations, mobile computing and Web Service settings. The software has been in the open source domain for the past 3 years. The 1.0 release of the substrate has more than 1100 classes with approximately a quarter million lines of code. It is hoped that these results will influence others to use NB as a messaging protocol for their own applications and develop its credibility within the distributed computing workgroups and consortia.

[1] http://www.naradabrokering.org

To test the performance issues of NB on a worldwide scale, NB brokers were set up at Cardiff in the UK, Indiana in the USA and Melbourne in Australia. This facilitated global testing over international networks. Port connection speeds of 100mbs were used as standard with extra testing over 1 gigabit networks between Cardiff and Indiana. Benchmarking software developed alongside the NB substrate was used to send message packets of varying sizes between messaging clients at all sites using exhaustive permutations of brokers and clients. An NB-GridFTP test bed application has also been developed to collate results of global network tests with varying message payloads and large attachments over Parallel TCP. Please note that the NB-GridFTP application leverages several services currently available within the NB substrate: the fragmentation service takes large payloads and generates smaller fragments; these fragments are delivered reliably by the reliable delivery service [5]; finally, the coalescing service merges these fragments into the original payload. This scheme enables very precise recoveries from failures while ensuring minimal retransmissions during these recoveries. Additional details regarding this scheme can be found in a recent volume of the Journal of Future Generation Computer Systems [6].

More recently NB has incorporated support for Web Service specifications such as WS-Eventing and WS-ReliableMessaging. Here we report some preliminary measurements related to incorporating support for Web Service specifications.

References

[1] R. Fatoohi and N. Gokhale, A Distributed Job Execution Environment using Asynchronous Messaging and Web Technologies, Proceedings of the 5th International Conference on Enterprise Information Systems (Angers, April 2003), pp. 246-251, April 2003.

[2] The NaradaBrokering Project at the Community Grid Labs: http://www.naradabrokering.org

[3] Shrideep Pallickara and Geoffrey Fox. NaradaBrokering: A Middleware Framework and Architecture for Enabling Durable Peer-to-Peer Grids. Proceedings of ACM/IFIP/USENIX International Middleware Conference Middleware-2003.

[4] Geoffrey Fox and Shrideep Pallickara. An Event Service to Support Grid Computational Environments Journal of Concurrency and Computation: Practice & Experience. Special Issue on Grid Computing Environments. Volume 14(13-15) pp 1097-1129.

[5] Shrideep Pallickara and Geoffrey Fox. A Scheme for Reliable Delivery of Events in Distributed Middleware Systems. Proceedings of the IEEE International Conference on Autonomic Computing. New York, NY. pp 328-329.

[6] Message-Based Cellular Peer-to-Peer Grids: Foundations for Secure Federation and Autonomic Services. Journal Of Future Generation Computer Systems. Volume 21, Issue 3, pp 401-415 (1 March 2005). Published by Elseiver.