Paper - Rigorous Evaluation of COTS Middleware Technology

Authors – *Ian Gorton, Pacific Northwest National Laboratory, Anna Liu and Paul Brebner, CSIRO Mathematical and Information Sciences*

Summary of the Paper

The authors discuss the Middleware Technology Evaluation (MTE) project which aims at providing technical evaluations of commercial-off-the-shelf (COTS) middleware products.

Broadly speaking, middleware of any software application comprises any or all of the components that go together to build the framework of a business application. It includes but is not limited to the business logic, load-balancing, transaction validation, security and authentication processes. Having to build and deploy all these components by using only the conventional software development methodology can be quite tedious and difficult. It is possible to divide most business applications into logical tiers, a front end, middle tier(s) and back end. Since most applications have their business validations built into the middleware, a broad similarity can be drawn between various business application software. That gives the reason for the growth in the number of middleware building tools in the market.

The authors note that even though using commercially available off-the-shelf middleware products reduces application development time and provides for solutions to many of the problems encountered when building the middle tier, their use is not always so straightforward. Since different COTS products are built to achieve certain level of performance and reliability, it is quite possible that they do not meet the requirements of a particular business. There have been many examples of such products that have failed to provide the necessary functionality and robustness to the application that is built using them, resulting in under performance or in many cases, total failure. The aim of the MTE program is to develop a testing methodology that can be used to evaluate the available COTS products against the business and functional requirements of the system to be developed.

Since the middleware has to perform number of tasks and that those could be very complicated, it makes the middleware very large and complex. The rapid growth and development in the middleware framework arena implies stiff competition amongst vendors. The lack of uniform specification results in different capabilities being offered by different products; evolving technology makes it difficult to discern between what is promised and the actual functionality offered.

In the MTE project, middleware evaluation is done using laboratory based generic product evaluation and organization-specific product evaluation. Former technique uses testbed applications with the middleware product to perform the evaluation. The testbed application is then used to stress-test the product on the same hardware environment that will support the business application. The process is called the Middleware Architecture and Technology Evaluation (i-MATE) process. It provides a clearly defined set of steps that can help businesses select a middleware product suitable for their application requirements. The laboratory tests consider all the technical aspects while the cost issue is considered in the context of specific organization requirements. Evaluations done using the testbed application are used to populate a knowledgebase which can be used by the organizations when running their i-MATE process.

The laboratory based evaluations are done using two complementary activities – qualitative architectural analysis and quantitative performance analysis. J2EE being one of the most technologically advanced middleware technologies on the market, the MTE project used the J2EE implementation of the middleware to perform testing. Six competing COTS products were evaluated and ranked for scalability and availability, application throughput, using a test case that imitates a stock-trading application. It exercises a number of key COTS middleware components, namely, middleware infrastructure that implements the N-tier application, transaction service, name or directory service and load balancing mechanisms to handle client requests. With the test application running, the team experiments with the products architectural alternatives and deployment options. Then the test application is run through a series of demanding performance tests for all the products being tested on the same hardware and software infrastructure. Performance is measured both for normal and heavy transaction loads to determine performance and scalability characteristics. The authors specify the test setup used, with one machine each for the middleware, client application and the back end. Of the six middleware app products tested – Borland Enterprise Server, SilverStream Application Server, Weblogic Application Server, WebSphere Application Server, Interstage Application Sever and JBoss, it was found that Borland and WebSphere clearly stand out in absolute performance terms; Interstage demonstrated its resource-limitation on a single machine – an important issue in cost and capacity planning.

While performing the COTS middleware evaluations, the project encountered difficulties relating to scale, complexity, product evaluation, vendor participation and cost. It is important to understand that each

implementation of the middleware framework has its own stronghold and weakness. Hence it requires good product knowledge and experimentation to achieve near optimal results with it. Also the evaluation team needs to work closely with vendors to diagnose any unexplained performance and behavior. The level of evaluation MTE provides greatly depends on vendor involvement as the vendors understand their technology and can offer valuable insights about the product behavior. With new versions of the products appearing, the evaluation team has to keep pace with them. The MTE project has been going on for four years now and has proved to be quite a success. What is needed next is to increase its impact by broadening the technology base and reaching a wider audience. Direct user involvement, vendor investment and independent, rigorous testing are the prime factors that can help make it even more successful. Vendors have reported positive sales growth for their products after the MTE testing.

MTE project deliverables include the evaluation reports, i-MATE process. Also the project has been a source for many research papers. The MTE also gives its users a repeatable evaluation methodology. Custom “add-on” evaluations can be done for organizations based on their specific requirements. This way certain features that are more important to them can be investigated in more detail.

Critical Comments

There is a tight coupling between the application components that build the middle tier of the application and the COTS “containers” which manage their life cycle and execution. Because of such tight coupling and reliance of the application components on the middleware infrastructure, how well an application performs depends not only on how well the components have been built, but also crucially on the middleware infrastructure too. Naively built infrastructure could lead to errors in component services, poor performance, inefficiency or failure. The application is built after selecting the middleware components; hence it is imperative that the infrastructure meet all the requirements.

The MTE fundamental principles center around the idea that no particular technology is solely better than any other. Each one has its own strengths and weakness, knowledge of which can help businesses to select the right product to suit their needs. The MTE findings are based on rigorous investigations and factual observations with the evaluation teams working in close liaison with the vendors because the vendors know their products best and can help resolve anomalies quicker.

The test application code isolates the different COTS middleware dependent components like database access layer, business logic, test drivers and performance analysis and measurement components in their own modules, making them easy to replace as products are changed. This allows for code reuse without middleware dependencies across different product tests.

MTE project demonstrates a practical application of the 80-20 rule of software development. The evaluation reports are found useful mostly to medium and large IT organizations that deploy applications that use COTS middleware products.

The MTE project has been a real good aid in evaluating middleware products and helps organizations reduce their up-front cost incurred in selecting a suitable COTS middleware product that meets their requirements. This also helps in reduction of total cost of software development by eliminating the risks associated with complex technology selection. Also it allows for early procurement thereby having a further cost-cutting potential. The laboratory based activity involves building testbed applications with the middleware products under evaluation.

Critical Questions

1. The middleware components are tested using a single server setup. What changes would be needed to evaluate the products for load-balancing using multiple servers?

2. Generally, a middleware COTS product specifies a particular hardware configuration on which optimum performance is achieved. In MTE project, the same hardware configuration was used for testing all the products. In such case, it would be interesting to determine how accurately the performance of the system can be compared.

3. Can the same testbed application be used with all middleware products or is there a need to write a different one for each.

4. The performance of the middleware products for increased client load was also tested on a single server framework. In real applications, however, multi-server framework is usually the norm.