Frequently Asked Questions (FAQs) on Sizing Ariba Buyer in the IBM Environment

How is the value "concurrent" users defined and calculated ?

In general, Ariba documentation states that for light users, use 1% ratio, for moderate users, use a 2% ratio, and for heavy users, use a 3%ratio. This ratio set assumes that the using population is homogeneous in that no particular individual is more or less likely to be requisitioning overhead items. So if Merrill Lynch has 3% of it's population buying at any point in time, the ratio would be 3 in 100. Thus, 10,000 employees would generate 300 concurrent users. If the 10,000 were spread out over 3 non intersecting time segments, then the using population at any point in time would be 3333 and there would be 100 concurrent users.

One needs to develop the requirements with the customer to figure out the ratios depending upon their usage plan. Not only that, Ariba publishes these numbers as guidelines only. Each environment must be planned and sized carefully. The idea of using a test center, such as at Gaithersburg or Dallas, to build the exact customer environment is a very good one The number of variables is high and the customer environment is not one that IBM can control.

What if there is a specific department that handles all requisitioning?

There are cases where only a few of the total population are considered users and are registered and they have the task of requisitioning materials for an entire organization as their primary job. In this case, as much as 10% of the registered population might be concurrent. Of course, the population is, therefore, far from the total population. If, for example, a procurement department were 10% of the total population and their usage was 10% concurrent, then the overall concurrent usage would be 1%.

If I add a location what happens to the concurrent load?
In adding a location in, say, Japan or elsewhere, take 1,2, or 3%, as appropriate, of the using population and add it to the population on that particular time zone. Then evaluate the concurrency.

How are response times developed and what do they mean?

The test standard used by Ariba, and to which we test, is that all elements of a 10 step requisition process must be completed in less than 6 seconds each and 90% of transactions must be completed in less than 5 seconds. The first of these elements is "get URL". The last is "submit requisition. So, when we publish a number of, say, 300 concurrent users as a maximum, this means that above 300 users there would be at least 1 transaction over 6 seconds or that caused the 90th percentile to go above 5 seconds. Also, each requester generates an average of 1.2 requisitions(random) and each approver approves 4.8 requisitions(again random). Also, each requisition is approved so there is a 1:1 ratio of requisitions and approvals.

Also, this test must run for at least 5 hours with NO exceptions to the performance listed. If any transaction causes an exception to the standard, the test fails and Ariba does not certify the performance result.

Ariba certifies the Buyer application's performance on each partner's platforms based on the

above criteria. Ariba has developed a load module that simulates the requisition process generated by users. One simply establishes the environment, ratchets up the number of users until “failure” occurs, backs off until it doesn’t and then submits data for certification.

Performance data published for the Ariba Buyer sizing and scaling brochures and tools are based on this process.

What is the role of requisition count in establishing requirements?

As an example, for 600 concurrent users, approximately 2750 requisitions will be approved in one hour. This is another metric that is used in sizing. With the ratio of 1.2 items per requisition, about 3000 items are purchased per hour. The customer needs to evaluate this kind throughput against their requirements.

If SSL(Internet – Secure Socket Layer – a network layer with security encoding and decoding) is enabled, the throughput is much lower, say 50%. Most people use an SSL card rather than the main processor for this reason, or they utilize a separate SSL box to protect the system.

What is the load generated by the ERP connector and how is it simulated in test?

The load by the ERP connector and the delays by the ERP system is impossible to generalize about because it is not possible for us to know the size and configuration of the ERP system and it's network. If the users are all local on a high speed(100Mb or greater) network, the performance will, obviously, be better than if they are dispersed around the globe and sharing a common database. The Ariba certification uses two methods and simulated loads. One for the Buyer application itself and another with the ERP connector added in. We have found that the ERP connector (TIBCO) adds about 6% to the processor load, thus reducing the maximum throughput by that amount. So, all of our test data that is published is with that 6% factored in as almost no one runs without some kind of an ERP system and associated adaptor.

What is the role of architecture in performance?

Some customers optimize the Buyer application performance by spreading function out on a number of boxes. The application is a three tier application. Some customers utilize the 3 tiers on two boxes and others on 3 boxes. There are cases that the customer wants to run with a single box. It depends on the IT team implementing the application. What we can say is that in absense of SSL, an ERP load, and with a high speed (100 Mb) network, all local users, and with a combination web/application server and a separate database server, the published values are appropriate.

What do I do about the customer’s request for an SLA(Service Level Agreement)?

If one is going to sign up for an SLA, be very conservative because the environment, the definition of the "using" population, the architecture, all play into the performance. Therefore throughput has to do with response time standards that eminate from customer behavior and satisfaction. The customer's environment must be put together in such a way to support that customer satisfaction.

Regarding scaling up as usage or business grows

At this time, the application is limited to single node form. This limits the concurrent users to about 1000 concurrent users regardless of processor count or speed. In releases next year, Ariba plans to offer multi-node application capability. This is not available as yet but shows promise to over 5000 concurrent with the right machinery and environment...

IBM’s scaling tools provide solutions for processors from 2 up through 6, though we have tested successfully up to 24 processor machines. Going past 6 processors, as of Buyer 7.0.6, doesn’t yield sufficient performance improvements to justify such a decision.