A Conversation with Jeremiah Wilton

By Dr. Paul Dorsey

SELECT's Executive Editor Paul Dorsey asked Amazon.com's lead DBA, Jeremiah Wilton (JW) about helping administer one of the largest e-Commerce databases in the world.

SELECT: What types of hardware and software architectures should large e-Commerce sites use in order to guarantee scalability and availability?

JW: The biggest mistake many companies make in trying to achieve scalability and availability is to adopt various technologies without understanding what those technologies actually provide. For instance, 9*i* RAC (formerly OPS) is a great technology, but it is not a panacea. If the goal is disaster recovery, then the fact that 9*i* RAC consists of only one database disqualifies it as an acceptable solution.

Similarly, 9*i* DataGuard (formerly Standby Database) is a great technology. But if the goal is to load-balance between systems, DataGuard is inappropriate because the standby can only be opened for read-only, and while open, cannot continue to recover.

Oracle's recommended solutions for scalability and availability are powerful and compelling technologies, but the basic business needs must be considered ahead of considering any specific technology. It may be the case that a company can achieve its goals using separate, non-related instances in an N+1 scalable shared-nothing architecture, where availability is managed through the development of intelligent application software.

SELECT: How can e-Commerce sites properly estimate the amount of hardware needed for such architectures?

JW: The most spectacular failures in e-Commerce have been companies that failed to correctly anticipate the scale of their operations. The best tool anyone can have for estimating necessary hardware resources is an existing production system. By using trend analysis, business projections, and knowledge of the software, capacity estimates for existing systems can be fairly straightforward. If a high-availability solution such as Data Guard is part of the architecture, then DBAs can employ resource and cost-saving strategies such as standby hosts with several databases, or standbys co-hosted with other production databases.

For brand-new systems, hardware estimation can be quite difficult. If a DBA wants to guarantee reliability up to a certain level, but doesn't really know how much load to expect, it may be necessary to over-engineer the hardware solution in order to guarantee scalability. Alternately, the system can be engineered such that small hosts can be added incrementally to accommodate additional load. This type of scalability can be achieved in Oracle using 9*i* RAC or N+1 shared-nothing components.

SELECT: What are the most effective types of software for use in a Web front-end?

JW: Web server software is more or less a commodity component. Among the popular solutions, no single one holds a great deal of advantage over another. This means that paying a lot of money for Web server software is probably unnecessary.

The ability to debug and profile running server processes, and run with a very low per-process memory size is important for maintaining availability and reducing hardware costs. Choosing Web server software that runs on Unix means that free software products such as GDB (the GNU Project debugger) and strace (system call trace) can be used for debugging and profiling the software's operation.

The hardest choice that Web-based Oracle users will face is the decision to use three-tier architecture. Using the 9*i* Shared Server feature (formerly Multithreaded Server), Oracle can be scaled up to many tens of thousands of concurrent Web server connections. This means that Oracle users can put off implementing a transaction processing layer for a very long time, possibly forever. However, a TP monitor will, at the very least, reduce the number of connections to the database, offloading utilization of CPU and memory from the database host.

SELECT: How much time and what types of effort must an e-Commerce company invest in systems and software development before they open their virtual doors?

JW: The software development time that must be invested before introducing a new service depends on many factors. The number of features, extensibility and quality will have the greatest impact on the timeline. Some companies developed software for two years before introducing their Web presence, only to have the service fall over under the load on the first day and stay down for six more months. Other more successful Web sites have started with a very basic Web site, and built incrementally on that with additional features and users.

Both load testing to scale using synthetic load software and thorough usability and QA testing using experienced QA specialists are essential to guarantee the stability of a new service. How much time you spend in these efforts depends upon the scale and complexity of the application.

SELECT: What are potentially the biggest challenges facing an e-Commerce company in creating a working system?

JW: Among the greatest challenges in creating a working system are hiring top-notch people, obtaining high quality vendor support, and choosing appropriate technology.

Hiring great people can be the absolute key to success for any company. This is particularly true of e-Commerce companies. Having a sharp, well-trained technical and business staff will translate into fast problem-solving and a high-quality product. Conversely, careless hiring will result in a staff with employees who know they should be doing something, but aren't sure what. This scenario can be very destructive, because these employees will spend lots of time producing something, but an ill-conceived product may require additional effort to engineer your way out of at a future date.

Good vendor support is also difficult to come by. Many large software and hardware companies will not realize that you are destined for greatness, and will treat you like all of their other small customers — poorly. Poor technical support can be a major obstacle. One way to overcome it is to join discussion lists pertaining to the product, and (after thoroughly consulting the documentation for the solution) ask the experts about any problems you might have.

Finally, choosing appropriate technology can be a major obstacle in engaging successfully in e-Commerce. Engineers should try never to succumb to pressure from marketing relationships or political justifications to adopt a particular technology. Engineers should choose components such as hardware, storage, O/S, TP and database software based on technical requirements, cost and staff skill level.

SELECT: What types of security measures can you recommend post 9/11?

JW: When the subject of security comes up, people tend to jabber on about encryption, single sign-on and strong authentication. This misses one of the main points we learned from 9/11, which is that people will exploit simple, obvious security holes long before they go to the trouble of breaking a 64-bit encryption key with a supercomputer.

In the case of Oracle, the most common, gaping security holes are default passwords, misconfigured listeners, and lack of a firewall. On many Oracle installations, nobody ever changes the default SYS, SYSTEM, DBSNMP, and OUTLN passwords. This amounts to laying out the welcome mat for people wishing to exploit your systems.

Because Oracle listeners can be polled by "lsnrctl" from any host, companies who do not have a firewall blocking traffic from the public Internet to their Oracle ports advertise not only the fact that they use Oracle, but also the names of services that the listeners can connect them to. The listener configuration that can be misconfigured to the greatest detriment is the External Procedure Listener. Anyone deciding to use exproc listeners should obtain guidance with respect to proper security from the documentation and from Metalink.

Having a security audit done by a reputable firm is also money well spent. These consultants will attempt to crack their way into your system using a variety of well-known exploits and methods, exposing your most critical vulnerabilities.

SELECT: What advice would you give smaller companies who want to build e-Commerce sites?

JW: I would advise smaller companies that want to become web-based merchants to consider carefully whether it is most advantageous to build their own platform or use somebody else's. A number of established e-commerce companies have made their platforms available for a fee. The decision to take one or the other of these routes centers mostly on the cost of building the considerable infrastructure needed for a robust and professional service.

SELECT: What can a B2C retail site do around Christmas and other holidays to help support increased volume?

JW: My favorite answer to this question is "send the developers to pack boxes in the distribution centers (so that they can't run expensive ad-hoc SQL statements and write lots of new poorly tested code.)"

A major problem for many US-based Web merchants is handling the increased load leading up to Christmas. It would be unfortunate to have to maintain a huge infrastructure for 11 months out of the year running 75% idle just to survive the load for one month.

In many environments, a significant proportion of daytime peak load comes from people in the back offices running less critical business analyses, decision support programs, or general poking around. During the Christmas season, these activities should be relegated to the off-peak hours so that all available capacity during peak can be devoted to serving customers.

In addition, few sites can consider themselves perfectly tuned. All sites believing that they have tuned their system to perfection should revisit Web server data caching, SQL tuning and application profiling. In the case of Oracle tuning, DBAs should adopt a wait event-based tuning model.

About the Author

Jeremiah Wilton became Amazon.com's first DBA in May of 1997. At that time, he did not even know how to do a hot backup. Since then he has presented three times at Oracle OpenWorld, delivering the OpenWorld IOUG keynote address in 2000 and receiving an honorary Oracle Master certification in 2001. You can reach Jeremiah via email at **jwilton@speakeasy.net.**