



## eBay Infrastructure — A Prototype for the Future

Analyst: Anne MacFarland

- *Massive growth.*
- *Constant change.*
- *Intolerance for latency in all processes.*

These characteristics of the eBay infrastructure that support a very high rate of transactions may seem like an outlandishly extreme environment. Yet, on second thought, what's not to like? As businesses are consolidating and aggregating their way to greater productivity, and looking, as always, for greater profitability, eBay is a good prototype for what they will become. As a business that designs its own infrastructure architecture, eBay does more than just shop wisely for piece parts. It builds and evolves cutting-edge architectures to support the vast scalability its businesses require.

Paul Strong, eBay Distinguished Scientist, has a pragmatic perspective not often featured in the blizzard of hype and how-to articles that are written about technology. While his infrastructure may not be typical of a smallish enterprise, his problems are. His needs for manageability and support for a high rate of change will resonate with those managing technology at enterprises of all sizes. His strategy for supporting massive IT growth concurrently with a high rate of change is of interest to any business with a growth strategy.

The eBay infrastructure is a grid of about fifteen thousand servers. It is predominantly a scale-out infrastructure of many small servers supported by some higher-processor-count servers for a federation of back-end databases. This infrastructure is focused on two custom-built applications that support the eBay world of Web commerce. **How this infrastructure is used, not what brands it comprises, is key to the success of the business.** Other tools and applications that enhance the site are integrated with eBay's application kernel and APIs. This allows a broad array of independent software vendors (ISVs) and other developers to integrate their software with the eBay environment, producing a software-as-a-service (SaaS) marketplace that benefits all parties.

The pace of development at eBay is aggressive. Builds of new functionality are planned a quarter in advance. New features are deployed every two weeks to both core and international platforms. eBay's *build Grid* development environment can build a site, starting with C++ and moving to Java, XML, and HTML, in half an hour and deploy it automatically to QA or stage it to production. Strong hopes to accelerate development to a monthly, rather than quarterly, cadence. Managing a rapidly-evolving infrastructure to minimize the latency of process that end users, both buyers and sellers, experience is an everyday reality.

The challenge is in the volume of configuration and reconfiguration that is needed. The solution is to automate the system in a way that does not limit the flexibility needed to respond to rapidly changing requirements. Strong has found grid architecture and virtualization to be the tools that make such automation possible. How eBay has used them, how this successful company combines commercial and custom software, and what Strong has to say about the standards that underlie scalable and flexible infrastructures are important to us all. For more details, please read on.

### IN THIS ISSUE

> eBay's Infrastructure Evolution .....	2
> Important Elements and Concerns of the eBay Infrastructure .....	2
> A Recipe for Using Automation Wisely .....	3
> Conclusion .....	3

## eBay's Infrastructure Evolution

### *Infrastructure Evolution*

The founders built the original eBay infrastructure over a weekend using parts that could have been bought at Fry's or Circuit City. As the business grew massively year after year, the stovepipes of the original deployment gave way to a 2-tiered architecture<sup>1</sup>, which was then made more resilient by pooling hardware assets and using a front-end load balancer. Because of the scale of the infrastructure, eBay engineers achieve resilience by  $N+M$  clustering. "The simpler the node, the fewer the special features like HBA's, the easier it is to isolate the fault. Says Strong, "you can then start to treat the computer as the FRU (field replaceable unit)."

### *Database Evolution*

The second phase of evolution involved the database areas. Splitting and federating the databases allowed eBay to support more transactions, yet still give the response time needed. Finally, in 2002, eBay separated the application's notion of the database from the physical implementation<sup>2</sup>. This virtualization supports more flexible access to database assets by the eBay application and allows more transactions to be run against the database tier.

### *Development*

At the same time, the development side of the eBay technology was also being stressed by rapid change. The code base moved to Java<sup>3</sup>, with MSXML for the presentation layer. eBay then implemented a *development kernel* on top of which all new features would be built. This standardization was the only way to be reactive to the needs of eBay users for new features.

### *Search*

eBay started with category browsing, but quickly added search. Over the years, search has become a significant tool on eBay. Indexes are now sliced in up to 16 grid-node-like pieces and each slice is replicated. When a seller posts an item to sell, he or she immediately does a search to make sure it is there. This kind of demand for real-time indexing and comprehensive recall was not met by existing search tools, so eBay built its own. Search has also driven eBay's tool strategy, due to the need to have the

telemetry to document that the search was fast enough. Supporting this granular level of telemetry with lots of agents was not attractive. eBay uses a commercial management framework's ability to gather and report, while writing their own equivalent of the framework's plug-ins to meet their specific needs. Strong hopes to move to more use off commercial tools, once their scope and focus on latency meets his needs.

### *Pervasive Virtualization*

eBay has adopted global billing, centralized application logging, and business event streams as objects to be managed. A multi-cast infrastructure allows workloads to be parallelized and search to stay fast. The significant change, for eBay, in moving to a more holistic grid organization was to abstract the tools by which the infrastructure was managed to move the focus of management away from the IT objects to business process objects. Strong has been working to insert the layers of virtualization needed to accomplish this. Strong's strategy at eBay for the near future is to work on tools integration, develop infrastructure tools based on the EGA<sup>4</sup> framework, and make management of both services and components a more navigable space.

## Important Elements and Concerns of the eBay Infrastructure

Throughout eBay's history and aggressive roadmap, certain issues surface repeatedly.

**Network-centric computing (Grid)** - It is not the hardware assets themselves, but the way they are connected that optimizes how business uses technology. Strong, speaking in Washington, D. C., last month, at *GridWorld*, emphasized strongly that the network-based computing paradigm of grid is the way to design scalable and flexible infrastructures, applications, and the tools that manage them.

The wealth of telemetry that can be exposed in grid architecture makes predictive capacity planning more accurate. This accuracy allows elements to be quiesced, reducing the environmental overhead of the "slop at the top" which is a big cost in most IT environments.

It is not just the resilience of grid architecture that provides benefits, but, more importantly, the management discipline of common practices for configuration and deployment, and other lifecycle

<sup>1</sup> There is no separate application tier. The application and the Web presentation are a unified functionality.

<sup>2</sup> A classic software separation of concerns that underlies, and is the basis of, all virtualization.

<sup>3</sup> Java allowed the coding to be "more vanilla." J2EE is used mainly as a servlet container. Keeping dependencies to a minimum is key to supporting both a high degree of growth and a rapid rate of change.

<sup>4</sup> For some time, Strong has worked actively with the Enterprise Grid Alliance<sup>4</sup> (EGA), chairing the technical steering committee. The EGA is the Grid organization that has taken the Open Grid Service Architecture (OGSA) and defined the management elements that map to it.

events, that allows the infrastructure to *both scale and change easily*. With that discipline, mapping the value of business flows to the cost of the physical infrastructure, in a time-sensitive way, can begin to happen.<sup>5</sup>

**Standards** - Adherence to industry standards is a key way to grow technology environments to be both large and inclusive. Standards develop, Strong notes, both explicitly via standards bodies and implicitly through predominant implementations and, increasingly, via open source. What is important, in any case, is to get a broad input from end users to guide the direction and pace of standardization. At this time, there are few end users on standards bodies. Most enterprises are second-hand users of standards, consuming those their vendors have chosen to implement.

Vendors can and do bring the requirements of their customers to the table, but such feedback is focused on particular assets. The richness of this feedback route depends on the patience of the end users in responding to requests of all the vendors of their infrastructure. Strong feels that more direct end user engagements with standards bodies would give a broader perspective and produce more timely and relevant standards.

**Modeling** – IT modeling, Strong feels, ranks right up there with inventories and accounting as things that have to be done well for a business to succeed. A good model is at the core of any explicitly-shared environment. If each new element is added to the IT model when provisioning is initiated, the model represents the desired state, against which monitoring and discovery can be compared. It is more effective than discovery alone,<sup>6</sup> because it can tell you what is not working, while discovery only discovers functional elements.

When you have a scale-out environment of thousands of low-cost servers, discovering dysfunctional elements is an important daily chore. One must be able to query a model, Strong continues, rather than having to continually query the individual components of it. Updates to the model must be precipitated promptly to all processes that use it. Strong's goal is to model all the layers of functionality, not just the physical infrastructure.

<sup>5</sup> For more on Grid and SOA as time-oriented architectures, see **The Clipper Group Navigator** entitled *IBM Multiplexes the Data Center with SOA and Grid*, dated September 30, 2006, and available at <http://www.clipper.com/research/TCG2006086.pdf>.

<sup>6</sup> Discovery will find elements that shouldn't be there – but in an enterprise environment, there should be nothing there that you have not deployed. This is a matter of deployment discipline.

**Control Mechanisms** – Control mechanisms depend on having the right information model, and determining the right way to persist and access it. Strong feels configuration management databases – a real hot system management topic today – are a part of the solution, but only one part of the management amalgam.

One must also remember, Strong cautions, that stability is always in contention with agility. Overly rigid constraint in the name of discipline is not good for business. As he explains, “manual processes that enable stability cause ossification. Once you automate those processes, you get repeatability (stability) and speed (agility) where required.”

Eventually, Strong feels, all layers of infrastructure, from the physical hardware through the logical partitions, operating system(s), and applications must be part of a federated wholeness that can be queried. Within that federation, a variety of data constructs, from relational databases to flat files to triple stores<sup>7</sup> or relational databases. Only this approach will allow process flow management to be automated in a sensible and safe manner.

## A Recipe for Using Automation Wisely

Strong feels that there are four prerequisites for automating massive, responsive, and rapidly-evolving systems.

**1. Make processes reliable.** If administrative processes are reliable and disciplined, then the components administered will be, as well. *Only then* can you safely automate provisioning, installation of operating systems, and applications. This is a bottom-up initiative.

**2. Improve your monitoring infrastructure at many levels.** You will want to focus on your business process goals, and collect all the information that bears on achieving these goals, be they quick response, identification of new sources of revenue, or simple support for rampant growth. It is not one magic tool or even one technology, but probably a bunch of them. You must be disciplined in your choice of criteria. This is a top down process. With well-considered metrics, you can tie business value to IT cost.

**3. Find ways to transform the data collected into useful and actionable information, both for the IT environment and for the business strategist.** Work on links – ways to foster understanding of each other's perspective. Strong notes this may be hard,

<sup>7</sup> Triple-stores are three-column tables of relationships that can be manipulated to bring the desired results into a quick, sequential read. In situations where they are appropriate, they give a faster response than other alternative structures.

because goals are different, vocabulary does not map cleanly from one domain to the other. This consensus between business and IT must also find a collaborative mode different from that of squabbling over budget expenditures. If you tie business value to IT cost, as in item 2 above, you will develop an environment where the vocabulary to talk about alternative strategies can thrive.

**4. Develop a knowledge base of patterns and trends of your business' use of technology.** This is different from documenting the tweaks of reassigning resources on the fly that are typical of Grid environments. It is a more long-term characterization of how your business uses technology. This is an analytic process. Its use is not just feedback to inform business strategy, but as a preliminary benchmark against which to assess how that business might use technology more effectively. As an example of this, Strong feels it is necessary to understand the long-term goals of the business, and what kinds of IT change are a consequence of those goals. Only then can enterprise IT do capacity planning and performance management effectively.

## Conclusion

The eBay environment embodies the infrastructure requirements of a new breed of aggregator enterprises. These new service providers must capitalize on the economies of scale to deliver a myriad of services, but must also compete with others who can do much the same thing by delivering more value, or delivering sufficient value at a lower cost. The one big eBay grid is not just an operational silo. The catalogs are part of it, as are the auctions, the real-time searches of indexes in cache, the billing, and even some batch jobs. The inclusiveness of this holistic system allows trade-offs to be made and sharing and change to be a matter of routine.

The eBay infrastructure is not just a platform; it is a way of doing business. Take time to understand this multi-dimensional bazaar, for it is the way of the future.



### ***About The Clipper Group, Inc.***

***The Clipper Group, Inc.***, is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at [www.clipper.com](http://www.clipper.com).***

### ***About the Author***

***Anne MacFarland*** is Director of Data Strategies and Information Solutions for The Clipper Group. Ms. MacFarland specializes in strategic business solutions offered by enterprise systems, software, and storage vendors, in trends in enterprise systems and networks, and in explaining these trends and the underlying technologies in simple business terms. She joined The Clipper Group after a long career in library systems, business archives, consulting, research, and freelance writing. Ms. MacFarland earned a Bachelor of Arts degree from Cornell University, where she was a College Scholar, and a Masters of Library Science from Southern Connecticut State University.

- ***Reach Anne MacFarland via e-mail at [Anne.MacFarland@clipper.com](mailto:Anne.MacFarland@clipper.com) or at 781-235-0085 Ext. 128. (Please dial "128" when you hear the automated attendant.)***

### ***Regarding Trademarks and Service Marks***

**The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager,** and "***clipper.com***" are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "***Navigating Information Technology Horizons***", and "***teraproductivity***" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

### ***Disclosure***

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

### ***Regarding the Information in this Issue***

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.