



Market Roundup

October 25, 2002

**IBM Introduces the Baby Mainframe – Small, Powerful and Linux-Enabled
Back in the Old Same Place**

Intel Puts Its Money Where It Hopes Its Chips Will Be

**Google Suppresses Sites in German and French Search Results – Drawing
the Line on Information Access**

IBM Introduces the Baby Mainframe – Small, Powerful and Linux-Enabled

By Joyce Tompsett Becknell

This week, IBM announced the eServer z800 model OE1, the company's newest entry level mainframe computing solution. While the OE1 has the same architecture as the earlier z800 and z900 systems, its primary difference is size. The previous entry level model, the OA1, has an 80 MIPS engine while its new sibling comes with a 40 MIPS engine and an integrated facility for Linux (IFL). The OE1 uses the new zELC software pricing which provides aggressive pricing for the zOS designed for z800 models, to reduce the overall cost of ownership of a mainframe for low-end workloads. Additionally, the system will support the z/VM software to enable logical partitions and clusters. The z800 OE1 system is upgradeable to the OA1, and from there all the way to a z900. IBM has also decided to let customers convert or downgrade existing OA1 systems to OE1 configurations if that is a more sensible approach for their workload.

The new z800 OE1 is designed primarily for the installed base of traditional IBM mainframe customers (i.e., S/390 and RS 6000) users. IBM would dearly love to see all these folks upgraded to the newer zSeries architecture they introduced last year. While many customers have made that leap, customers with workloads well below the 80 MIPS entry level of the OA1 have been hesitant to pay for MIPS they do not need. In addition to software savings, offering a smaller, less expensive system serves as a competitive deflector to companies like Sun Microsystems which have been pitching their UNIX systems as mainframe alternatives. This new OE1 system ups the competitive ante by providing mainframe power more competitive price points. In addition, many customers who have sunk significant investments into traditional mainframe architecture are curious to explore Linux. IBM is positioning the z800 OE1 as an option that lets them have their cake and eat it too. By offering both capabilities on a standard system, IBM hopes to entice more users to experiment with minimum risk.

We have been interested to watch vendors continue to blur the lines between computing architectures. It is becoming common to find the mainframe competing with high-end UNIX offerings, and we assume that will extend to 64-bit Intel architectures as they move into the mainstream. UNIX OS flavors have been battling for years now, and Linux versus Microsoft has become a commonplace battle cry in the fight for infrastructure and low-end application servers. While Sageza believes architectural evolution is a good thing, we also believe that IT managers and their vendors are entering a period of great uncertainty, perhaps unknowingly. Not all operating systems and architectures make sense for all applications or compute environments. Managers need

to make sure they have a firm grasp of what they really need as they welcome salespeople into their offices. Vendors need to resist the urge to treat all situations as a nail in need of their unique hammer. Balancing support capabilities, the in-house knowledge base, costs of software, and ability to integrate with other applications will become even more important than hardware and OS capabilities in the near term. Sageza believes that the vendors who offer highly flexible solutions such as the z800 OE1 that meet specific client needs are doing what they need to create loyal customers and slowly win over the confused.

Back in the Old Same Place

By Jim Balderston

IBM announced this week that it has officially formed an autonomic computing unit, which will be headed by Alan Ganek, former vice president of strategy at IBM Research. The new unit is designed to integrate autonomic computing efforts across the company. Among the various efforts to be mounted within IBM around autonomic computing will be the development of a deployment model designed to guide customers through the process of developing an autonomic computing environment. IBM Global Services will develop a Resilient Business and Infrastructure Solutions Practice, and the company will develop autonomic computing design centers to help partners and customers design and test autonomic computing products. IBM also plans to move its existing technology product portfolio into the autonomic computing model. The company said WebSphere Applications Server Version 5.0 will include features that will allow it to automatically monitor and fix performance problems. It also said that its Tivoli products already contain more than two dozen autonomic computing features and that its DB2 database also already includes self-managing and self-tuning features. The company also said its Enterprise Storage Server, Code-named Shark, will have ease-of-use technologies built in. Finally, the company said it will continue developing autonomic computing features for its PCs.

IBM has been talking quite a bit lately about autonomic computing. The phrase pops up in various press releases, presentation slides, and company speeches. Before this announcement, IBM was in a good position to abandon the concept if it did not gain any traction in the Tower of Babel that is IT marketing. With this step, one thing appears clear. IBM believes that the “autonomic computing” concept has got legs. Formalizing its commitment to the concept by naming people who own the idea inside the company, and telling the world that this is a strategic direction, makes it a lot harder to walk away from now. Rarely known as a rash and hasty outfit, we suspect Big Blue is making this decision on very solid ground.

So what happens now? One could argue that naming existing products with the first iterations of autonomic computing technology already in place is proof of actual substance to the concept, or conversely, simply the repackaging of existing product in a new, shinier wrapper. While certainly a little of the latter is going on, we believe that the next twelve months will see some meaningful enhancements to the autonomic computing line, just as the past twelve have. By promising increasingly intelligent network and computing components — be they Sharks or PCs — IBM is promising a lot. Meanwhile, IBM is not alone in this latest search for the holy grail, with HP’s Adaptive Infrastructure and Sun’s N-1 efforts in high gear. However, if IBM’s delivery of truly intelligent enterprise IT environments continues on pace, it will secure for IBM a familiar position in the IT vendor space: top dog. It may seem like years since that IBM has been acknowledged as such, but those years were measured only in Internet time, after all.

Intel Puts Its Money Where It Hopes Its Chips Will Be

By Jim Balderston

Intel has announced that it plans to invest \$150 million in companies that develop WiFi (802.11) technology. The Intel Capital Communications Fund will invest in companies that are developing hardware and software products and services that create an easier to use, more secure wireless environment, simpler billing procedures, improved network infrastructure, new ways to connect to high-speed networks outdoors, and ways to deliver services over the network. Intel said it has already invested \$25 million in more than ten

companies in this area. The company also said it would continue to invest in its Baniyas chip, mobile computing technology that has been designed from the start to be applied to mobile PC uses and which contains dual band — both 802.11a and 802.11b wireless capability. Baniyas is due out in the first half of next year.

We have had little quibble with Intel's Baniyas initiative; anything that gives us more WiFi capabilities is just fine and dandy with us. It should come as little or no surprise that Intel is investing some of its sizable nest egg in companies that will help drive a market that Intel execs put at 30 million desktops within three years. It's not the first time they have made such investments, nor will it be the last. However, the play will be incomplete if the effort stops at WiFi. For mobile computing to reach everyone and change the fundamentals of connectivity, it will be necessary to complete the circle by seamlessly integrating WiFi with other technologies so that users seamlessly transition between next generation wide area and local wireless without affecting content streams. But that's for another discussion.

There is little argument that much of the WiFi core elements — like robustness and security — still need some fine tuning. As this weekend's WorldWide WarDrive will no doubt illustrate, there are plenty of improperly secured wireless LANs out there. As these geeks roam around cities around the world looking to identify — but not use — these rather wide open wireless access points, it will have little impact on the increasing momentum of WiFi deployments in the long run, regardless of the scale of ineptitude demonstrated by the WiFi network admins in this weekend's big driveby. As P.T. Barnum once said, "It doesn't matter if you write good things or bad about me, as long as you spell my name correctly." The name of this game is WiFi and we suspect few people will have trouble spelling it in the coming years.

Google Suppresses Sites in German and French Search Results – Drawing the Line on Information Access

By Joyce Tompsett Becknell

This week, a Harvard Law School report indicated that sites were missing from search results conducted at the French and German Google sites. According to the report, most of the missing sites were those that support racism or deny the Holocaust. The reason for these missing sites is that both France and Germany have strong laws that prohibit hate speech, and these sites fall under the provisions of these laws. It is common for the various search engines to run different sites for different countries, adapted to the native language and currency. Users would generally not be aware of a missing link unless they were to perform the exact search at several locations, such as Google U.S. as well as Google France or Germany. Restricting access to these sites when in violation of a local law has become a common practice of most major search engines.

While the idea of blocking offensive material presents a popular topic for freedom of speech advocates, the problems go well beyond freedom of expression issues. Google, as well as Yahoo!, Amazon, and Alta Vista all face continuing threats of lawsuits for allowing questionable material to be procured through their services. German companies threatened to bring suit to several companies earlier this year when it was alleged that directions for how to disassemble components of Deutsche Bahn systems could be found on line. The Church of Scientology is another organization that has leveraged lawsuits to stop links to sites critical of the church based on legal reasons such as copyright violations. The Chinese government is one of several that restrict the type of information citizens have the right to access. In the wake of the September 11 terrorist attacks last year, the U.S. Government considered censoring itself regarding the types and range of information it made available publicly on the Web through its various agency sites. The argument underlying all of these examples is that search engines make it considerably easier for dangerous, valuable, or misleading information to get to the hands of large numbers of people. Before the Web, all of this information existed but was much more difficult to disseminate widely or quickly. The ongoing crawl toward broadband, the increasing sophistication of search engine technology, and the growth of Web access means that the entire spectrum of human thinking is available to us with a few clicks of the mouse. The troublesome question is how to regulate that access.

It is impossible to remove the offending sources themselves. Like mushrooms they will proliferate overnight

despite all efforts to remove them. The current compromise is to force the search engine companies — the providers of the linking technology — to police the Web. Mostly this functions on a reactive level. Someone reacts to the presence of objectionable material and the link is duly severed. However, at some point, limiting access to information becomes censorship and the definition of questionable material becomes highly subjective. Additionally, it becomes ever more difficult for search engine companies to maintain the responsibility for managing Web content through their services. Already some groups argue that the service providers who host the sites where content is stored should be responsible for this information. We believe that there are two complex issues at stake here. First is the difficult question of who should be able to determine what is acceptable and what is not. The second issue is the question of who is responsible for monitoring and enforcing these definitions, assuming anyone can agree to them at all. It is hard to draw concrete geographic boundaries on a logical structure. It is harder to enforce on country's set of rules on the citizens of another when the information is delivered to them by a company based in a third country, delivered on equipment owned by a company in a fourth country. Institutions like the United Nations or even something regional like the European Union are wholly unequipped to tackle these issues effectively. Sageza expects this to be an ongoing battle pitting localized political power against the forces of human expression and technology that will make today's ongoing World Trade Organization disputes look like child's play in comparison.