



Research Report

IBM's Smarter Systems — Not Quite What We Expected

Introduction

The primary reason that *Clabby Analytics* (more specifically, the reason that I, Joe Clabby) journeyed from Maine to IBM's Almaden Lab in San Jose, Calif. was to learn a bit more about IBM's new Smart Analytics Systems. These are new systems bundles that have been designed and packaged to perform business analytics in an optimal fashion on IBM System x (x86 multi-core); Power Systems, and System z platforms. What I was expecting was a deep technology drill down into systems technologies, software components and related bundling. What I got, instead, was something even better: a view into how IBM customers are actually using IBM "smart system" to run their businesses. Further, I got a competitive overview of the business application and database markets from both the customer's, as well as IBM's perspective. (IBM rarely provides competitive commentary — so a snapshot of IBM's view of competitive markets was pretty unusual — and very insightful).

My "takeaways" from this conference are structured into the "lessons learned" subsections that constitute the rest of this article. They are organized around: 1) competitive positioning; 2) the amazing database; 3) workload optimization; and 4) summary observations.

Lesson 1: Competitive Positioning

At Almaden, IBM set up two panels with three customers on each panel. And these customers described how changing market dynamics (such as architectural consolidation and competitive positioning) have caused them to adopt IBM systems and infrastructure solutions.

From a competitive perspective, here are some of their insights (and/or conclusions I have derived):

- The Oracle acquisition of Sun has created some very interesting market dynamics. First, there is a belief that Oracle infrastructure and applications will be optimized for Oracle's Sun platforms. This leaves Hewlett-Packard (HP) in the lurch — because HP relies heavily on Oracle for both middleware and business applications. Second, the acquisition left several of these customers concerned about the future of Sun's SPARC architecture — to the point that they are moving off it. These customers confirmed a market trend I am already seeing — that a big migrate-away-from-Sun trend is continuing despite Oracle's assurances (and accordingly, these customers have indirectly added credence our position that SPARC will no longer be a leader in general purpose computing). Third, the Oracle/Sun deal leaves Sun customers who run SAP kind of out-of-step in that if you're running SAP on Sun, you're now running it on a platform owned by SAP's #1 competitor. And this makes Sun/SAP shops a little edgy...).

IBM Smarter Systems — Not Quite What We Expected

- If HP wants to sell business applications (ERP, CRM, SCM), then it can continue to sell Oracle solutions (this must warm their hearts a bit given that Oracle is now competing with HP for server business) — or they can get more aggressive about selling SAP applications to counter Oracle, thus potentially irritating their biggest application partner. But IBM's customers pointed out that they believe IBM has a much tighter relationship with SAP than HP (several customers mentioned IBM's joint development labs in Waldorf, Germany near SAP's headquarters, as well as advanced development taking place in Zurich, Switzerland). Further, IBM pointed out that SAP is certified on its systems.

Given these dynamics, here is what I think will happen in the business applications marketplace over time:

- As demand declines for Itanium and SPARC, I expect Oracle's application base will ultimately move to x86 multi-cores. Note that IBM's x86 multi-cores have very strong advantages over other vendor's implementations, so IBM will also pick-up its fair share of Oracle applications as they migrate to x86.
- Given clear superiority in RISC architecture, I expect that, over time, most of SAP's business application base that today resides on Sun or Itanium will move toward x86 multi-cores, POWER or z architectures. And, accordingly, I think SAP will have broader range than Oracle when it comes to deployment on future "workload optimized" systems.

Lesson 2: The Amazing Database

To be honest, I find database and storage technologies about as exciting as watching paint dry. So thank goodness that the other half of *Clabby Analytics* (Jane Clabby) loves covering those technologies (Jane was at Data General/EMC for 17 years). Unfortunately for me, at this Almaden event, the topic of databases came up fairly regularly. But, as it turned out, that ended up being a very good thing.

Several customers pointed out that they were in the process of moving away from Oracle's database products to IBM's DB2. There were testimonials about how unexpectedly easy the move was (using the advanced tooling provided by IBM's Migration Factory). And there were discussions about pricing advantages customers experienced by making the move. And, amidst all of the pricing advantages, one of these advantages kept coming up: the sophistication of the compression algorithms in DB2.

According to these customers, when moving from Oracle's database to IBM's DB2, they are finding that they are able to recover 40% of their database disk space. 40 PERCENT! Dwell on this number. This means that enterprises that have spent \$20 million for database storage could have saved \$8 million in storage costs if they used IBM DB2. Amazingly compelling if true, and according to these customers, IBM's assertions of major storage savings due to DB2 compression are true.

After the customer sessions, I tracked down an IBM database expert to shed a little light on why this is the case. According to this specialist, the reason for IBM's marked advantage in compression has to do with how DB2 handles database dictionary look-ups. Apparently,

information that is stored again and again (like the word “encyclopedia” which takes up a lot of space) can be simplified to a few characters and placed in an index.

Those characters can be used in the database rather than the whole word — and accordingly, a lot of space can be saved. IBM appears to have a distinct advantage over Oracle in compression — and the company’s customers are finding this to be a huge cost savings advantage. Further, the IBM DB2 customers I talked to expect that, with the projected growth in information capture in coming years, IBM’s DB2 compression will have a huge effect in helping them contain costs.

Lesson 3: A Workload Optimization Matrix?

Wouldn’t it be great if IBM created a matrix that portrayed which workloads belong on which systems? Then, IT buyers could refer to this matrix to achieve maximum business efficiency — right? Well, I’ve concluded that this just isn’t going to happen — and I suspect that the reason why it isn’t likely to happen is because of one of IBM’s big selling mantras: we offer our customers “choice.” So IBM isn’t going to say — “run this here on this systems architecture — and run that there.” Instead, it is going to let its customers choose the platform/platforms best suited to run their application mixes.

So, instead of producing a workload guidance matrix, IBM’s customers described how they are optimizing workloads within their information systems environments. And these customers presented a gold mine of information and shared some brilliant insights.

My takeaway was this: IBM’s customers are figuring out for themselves which of their applications run best on which systems. IBM is providing guidance — but customers are weighing the differences between x86, POWER, and z — then choosing the platforms best suited to serve their needs. Accordingly, I no longer expect a massive matrix that describes which applications belong where in its systems portfolio.

Summary Observations

To date, IBM’s latest high-level messaging (smarter planet) has largely focused on describing how vast amounts of information are being generated, and how vast numbers of devices feed information capture (usually described as instrumentation), and/or on how new devices require access to captured information (often concentrating on a plethora of new and evolving mobile devices). Smarter planet messaging has also concentrated on the vast amounts of data that need to be stored — and on processing vast amounts of data in real-time (or near real time). Add in some green messaging — and then show how new data processing capabilities can serve certain industries better (such as smart traffic systems, smart water management, smart energy grids, smart healthcare, smart food systems, intelligent oil field technologies, smart countries, smart cities, etc) and voila — you end up with a smarter planet.

But how do you build smarter planet solutions? Last year IBM’s dynamic infrastructure group showed the market how to glue together and manage smarter planet infrastructures. This messaging focused on lowering costs (through virtualization, energy efficiency, standardization, and automation) and optimizing business (by aligning business and IT, improving service and flexibility, and basing information systems on industry standards). And it described products that can be used to build a smarter planet dynamic infrastructure

IBM Smarter Systems — Not Quite What We Expected

(including service, asset, and virtualization management — as well as green management products). Further, IBM's information infrastructure group did a brilliant job showing the market how to lower storage costs while managing vast amounts of information better. But, the systems side of the equation was still missing.

What I had expected in Almaden was a more system-focused “how to go about building smarter solutions” perspectives. But instead, I may have gotten something better. The presentations at Almaden Labs gave me new insights into certain market competitive dynamics; a new view of workload optimization; and insights into significant database differentiators. And, through customer examples, it showed me how smarter planet solutions are being developed in the real world in traditional commercial markets. Accordingly, I walked away from Almaden not with a systems overview (which IBM has since supplied in an in-depth separate briefing), but with some very valuable competitive and market dynamics information that will definitely influence my future research as I continue to articulate how IT executives should build smarter systems environments.

Clabby Analytics
<http://www.clabbyanalytics.com>
Telephone: 001 (207) 846-6662

© 2010 Clabby Analytics
All rights reserved
June, 2010

Clabby Analytics is an independent technology research and analysis organization that specializes in information infrastructure and business process integration/management. Other research and analysis conducted by Clabby Analytics can be found at: www.ClabbyAnalytics.com.