HP Bladed Architecture

Index

About the Author:

Richard Buckle is senior director, Business Development, GoldenGate Software, San Francisco. He spent seven years as a director on ITUG Board and served as its chairman for two years (2004 - 2005). Buckle worked in the Tandem / NonStop marketplace for more than two decades. He is known as RT Writer (for Real Time Writer and taken from the column Real Time View published in the ITUG *Connection* magazine). Richard Buckle created the ITUG-Connection Blog in 2007. He made his first post to the blog in August and he has since continued to blog his experiences with the ITUG community.

Read more at his blog – **http://itug-connection.blogspot.com/.**

Introduction

The creation of an ITUG community blog was first raised in July 2007 following an ITUG board meeting in Chicago, and the first posting to the blog was made on August 20, 2007. The blogs use a number of techniques to attract readership – the use of photos, the description of my current location, the ITUG communities I visit, the cars and bikes I am using, and song lyrics I recall all play a part in developing the story that makes up each blog posting.

This is the first grouping of blog postings around a single theme – and the theme is the impact on users of the arrival of blades, and the much-anticipated HP bladed architecture. HP has been signaling the arrival of such a package for more than a year and it is widely known through the community that Early Adoption Programs, or EAPs, are under way at user sites, and a number of Independent Software Vendors (ISVs) have been involved in platform testing..

In the September 17, 2007 posting, “Fast Forward! Transformation and Morphing!”, I introduce the topic of social networking and then transition to the main theme – HP’s imminent release of a bladed architecture spanning the major systems now shipping, including NonStop. I start out by looking at the impact of server consolidation within large corporations, where considerable effort is being taken to consolidate many servers onto a single, highly available, and scaleable platform. I then take a look at the efforts of both HP and IBM to consolidate their enormous IT infrastructures and how both vendors are pursuing a similar course, and both are making use of their biggest systems – HP with NonStop and Neoview and IBM with System z. I conclude with the observation that while the systems may begin to look very similar, the main difference will be that HP will leverage industry-standard chips and open software.

In the postings of September 26, 2007, “It’s time to leave the Duke!” and February 12, 2008, “’My Wish’ for NS Blades”, I pick up on comments made by Winston Prather about HP’s investments in the NonStop platform, most notably moving it to a bladed architecture, and how this is consistent with observations reflected in an IDC survey of CIO’s who were “seeing that the economics of yesterday’s data center (wasn’t) going to work in the future” and more and more servers simply couldn’t be thrown into the data center, nor could more staff be added to manage each new application.

In other words, adoption by HP of a industry-standard (Intel) building blocks capable of supporting any mix of operating systems suggests that NonStop may become an option on every bladed server deployed. This leads to the wish list that I provide where (1) HP delivers on the “shared infrastructure blades” product line, (2) HP delivers virtualization that supports NonStop as a guest operating system, and that (3) HP delivers the APIs and underlying support to allow any mix of

operating systems, including NonStop, to be based on the incoming transaction types with the number of processors supporting any operating system “floating” to meet the demand at the time.

In the posting of March1, 2008, I return to IBM and look at the probable impact the tenth generation mainframe (z10) will have on IT and come away a little disappointed that IBM persists in talking about mainframes, rather than servers, and continues to deploy books, not blades. However, the introduction of Infiniband, with “the ServerNet footprint all of Infinband”, having the potential to steer other vendors, including HP, to broader acceptance of Infiniband as a cluster technology. Blades with Infiniband support would provide a more consistent packaging and remove the need for NonStop to have specialized packaging with mezzanine cards in support of ServerNet.

Finally, in the posting of March 9, 2008, I look at the release of the NonStop Cluster Essentials package as a key piece in the management of the new bladed architecture where a mix of NonStop and Linux will be most likely. Without a uniform management capability any execution on Martin Fink’s shared infrastructure blades program would be flawed, given that each system sharing the package would need its own management console and appear anything but integrated.

The development of this group of blog postings around the theme of HP’s bladed architecture – what to expect, what the competition will do, and where it will find value – is the first “chapter” in a line of such compilations. Future themes will address the roles of key IT groups, the greening of our data centers, as well as specialized topics like SOA, Virtualization, Security, and Compliance to name a few.

Monday, September 17, 2007

Fast Forward! Transformation and Morphing!

I have now returned to the US – after a short stop over in Boulder, I am back in my Simi Valley office and the picture I have included here is of me just off of Presidential Drive overlooking Simi Valley. It doesn't matter how many times I make the trip to Sydney, I still come back whacked each time … there's just no fun from those long-haul over-water trips.

But now that I am back at work, with access to my system, I have gone back and added a few comments to earlier postings. I won't cover them all here, I promise! I have to admit though that I am pretty pleased with the level of response I have seen so far.

The blog postings "Is 30 Minutes too Long?", "Got Security?", and "Whereto CTO?" all generated comments that I felt are worth revisiting. A steady trickle of postings is beginning to appear. Clearly, this isn't the only dialogue taking place across the ITUG Community – a number of online discussion forums are supported by ITUG and can be easily accessed from the ITUG web site – **http://www.itug.org/**

There are newsletters as well, and Bill Highleyman's Availability Digest is a particularly good source to learn how users are deploying highly available

solutions. Bill has told me that “the Digest really is not a newsletter. It does not talk about current events. It is a Digest in the purest form of the word, summarizing various topics in high availability.”

What’s a little different about this blog is that it is not so much logging questions as they arise but rather, attempting to anticipate the questions that are likely to arise soon; covering topics and trends I pick up on as I talk to clients and partners – and to HP itself. So far, I have discussed the user group itself and where it could be headed, looked at virtualization and Linux, thrown in some thoughts on availability and even given consideration to the roles of architects and technologists. For the next few weeks I will continue to introduce different subjects and see which of them generate traction with the community - I am looking at the comments posted as one way to determine interest levels, and to validate my belief that the issues are about to arise.

What users of HP NonStop servers should be glad to see is that there is a bit of an ecosystem developing around the platform – and more information is becoming available. Whether you pick up information from this blog, a discussion forum on the ITUG web site, or from newsletters or digests, such as the one Bill is producing – the good news is that people are now talking in a public forum about NonStop. In the past, these exchanges have gone on but usually captured within emails so that most of us never had the opportunity to read about, and to leverage other users experience is an incredibly important element of growth.

I don’t think there are any of us that do not want to see growth in NonStop, see more users deploy HP NonStop servers. Every time we hear of a new application on NonStop we get excited and every time we hear of a new customer – someone trying NonStop for the first time - we are pretty impressed. Throughout the corporate world, consolidation is rampant and when discussions focus on the consolidation of many servers back onto a single platform – then the question always comes back to “is the platform reliable? Is it highly available?” and “can the platform scale?” and still, there’s no platform the equal to NonStop on both counts.

Now, talking about Boulder and Sydney, have you read the recent announcements from IBM? While most of us have been made aware of HP’s CIO,

Randy Mott's mission to collapse thousands of HP internal servers back to a few large complexes based on HP NonStop, and to deploy them across six data centers – some of which are still under construction – did you pick up on IBM adopting pretty much the same plan?

In case you missed it, check out:

http://www-03.ibm.com/press/us/en/pressrelease/21945.wss

ARMONK, NY - 01 Aug 2007:

IBM data centers in Poughkeepsie, New York; Southbury, Connecticut; Boulder, Colorado; Portsmouth, UK; Osaka, Japan; and Sydney, Australia, will participate in the initiative. IBM has established world-class teams to migrate, test and deploy the applications, which include: WebSphere® process, portal and application servers; SAP applications; and DB2®.

And for even more details, check out:

http://www.ciol.com/content/980798930.aspx

Thursday, August 09, 2007

In a significant transformation of its worldwide data centers in a generation, IBM has announced that it will consolidate thousands of computer servers onto about 30 System z mainframes running the Linux operating system. "The mainframe is the single most powerful instrument to drive better economics and energy conservation at the data center today," said James Stallings, general manager, IBM System z mainframe. "By moving globally onto the mainframe platform, IBM is creating a technology platform that saves energy while positioning our IT assets for flexibility and growth."

When you take a good look at the configurations – each of these System z mainframes will be a hybrid of sorts – with z/OS and DB2 being at the center, and the application code (mostly Java) running within Application Servers on Linux. Sound familiar? Among the locations will be Boulder, Colorado as well as

Sydney, Australia. Go figure … and I had nothing to do with it!

The point here is that the HP NonStop server was not only the logical choice to anchor HP's consolidation, but HP's main rival also picked it's own largest system in pretty much a parallel move to HP's. It's kind of reassuring, at the end of the day, to see both HP and IBM electing to go with systems most of the customers have been depending on for key transactional and data base applications for several decades.

Over time, you will see both of these mainframes, or "mainframe class of servers" morph to where they will be very similar in many ways – the main difference will always be that the HP servers have NonStop highly integrated for better availability and scalability. HP will leverage industry-standard chips – and as long as IBM continues to own its own chip technology, IBM will depend on higher-cost proprietary technology. I kind of like where this is headed – I have worked with both systems over the years – and they are both suited to the consolidation tasks awaiting them.

I think most of us understand that we live with a many-tiered technology deployment model with Microsoft on the desktop, some Microsoft and even Linux on the web servers, mostly Unix on the application servers, and significant populations of NonStop and zOS at the core. So what we are seeing today from HP and IBM plays into this many-tiered model and potentially reflects where many of us are headed. Perhaps more as Blades becomes better defined.

With the advent of newsletters, digests, discussion forums, and blogs – an information ecosystem focused on providing independent news to the NonStop community, it will become much harder to ignore the success that HP has with NonStop. And so the question needs to be asked – are we taping into these services? Are we forwarding information to our management? Or do we think that the moves being made by the two biggest vendors don't really apply to our situation?

Wednesday, September 26, 2007

It's time to leave the Duke!

Back in Boulder again for a few days, and then it's back to Simi Valley next week. Later next week, it's off to the Euro ITUG event in Brighton. So what I would like to do tonight is to pick up on one theme I have been covering – the Integrity NonStop platform itself.

It is clear to me that in the months to come, HP will be distancing itself from its competitors based on its roll-out, top to bottom, of blades packages. And nowhere will the impact on users be more visible than in the NonStop community. IBM has already made it clear that, for the time being, its mainframes will not be based on blades – the basic building block of "books" will continue for new product introductions for some time to come. This is not to say that IBM's decision is wrong, and they are certainly enjoying an up-tick in mainframe usage of late, just that it will be pursuing a different path to HP.

As I flew to Denver I grabbed a couple of magazines and among them were a couple of recent ITUG Connection issues. I was flipping through the pages of the July / August 2007 issue when I ran across the Winston Prather piece "News from HP's NonStop Enterprise Division". Winston is the VP and General Manager of NonStop Enterprise Division & High Performance Computing within HP. With overall responsibility for the Integrity NonStop platform – it just makes sense to always read his column. So it was with interest that I caught the statement "HP is

making significant investments in the platform, most notably in moving it to a bladed architecture.

What further caught my attention was the follow-on remark “we will also be moving storage and communications to Linux front-end servers in the future. Greater leverage of volume economics through increased use of industry-standard components will translate into lower TCO for our customers.”

With IBM and HP going in slightly different directions, the user community will face a number of decisions – but I am always comfortable when there are choices. Blades versus books! Itanium versus Power! And different views on virtualization! These are all topics I plan to cover in future blog’s postings.

What Winston is foreshadowing, to me, is a new paradigm where a common blade building block – even within a single system, such as the integrity NonStop - may be the same, but may be running different OS’s. Underneath the covers, of a future system may be a combination of OS’s each selected to support a specific function - be it storage, communications, etc. As most of you know, I have worked in communications and networking for more than three decades, and today I take no issue with electing to run communications stacks on top of a Linux distribution. It makes all the sense in the world and I can only see more complete protocol offerings and better industry support as a result.

Servers, based on standardized hardware, have a shot at being more energy efficient and we will see their “greening” become a priority. By this I mean placement of the blades within the total package of racks can be scrutinized and then organized for optimal heat management. I just have to believe mapping the heat signature of a common building block will then lead to better blade placements and a more energy efficient overall package.

We are going to see a lot of changes, and have to adapt to new ways of doing things. We may be moving in a totally new direction with different manageability priorities and a whole range of new interfaces and tools.

Last year, in a magazine insert called Next-Gen IT (July, 2006) that was put together by the editors of Computerworld and CIO magazine, Michelle Bailey, an

IDC research director said “CIOs are seeing that the economics of yesterday’s data center isn’t going to work in the future – you can’t have more people, more I/O, and more servers with every new application.”

This week I caught up with Wil Marshman and we met in the Duke of Edinburgh for a drink. For all of you not familiar with the Cupertino campus, for the past 25 years or so many of us thought the pub was part of the campus. Back in Tandem days, on special occasions, it was even graced with a regular Tandem building “tombstone” - although I forget the specific location number it was assigned. Being back in the pub with Wil brought back a lot of old memories and as I looked around, there was Jack Trice – in his usual corner! I walked outside and there was Andy Hall on his cell phone, talking to family – it was Andy who first took me into the Duke in 1987 and where I first met folks like Roger Mathews and Steve Saltwick.

And it just stopped me dead in my tracks – for 20 years, we had kicked back in this place and brainstormed all sorts of wild and crazy ideas. It was a comfortable place, and we all knew where to head at the end of the day to catch up with folks we neede to see. We were kind of reluctant to try any place else! And it reminded me of what I had written only days earlier in my August 29th blog posting (Back Home … To NonStop) “what we have considered as our safe and trusted turf may be moving underneath us”!

I guess in some ways, we all would like to stay with what we have. We feel most comfortable working with tools and utilities that we have depended on for many years – where we can quickly comprehend, and react to, the information returned. But I sense we are headed in an entirely new direction and so much of it will be different, unfamiliar, and perhaps a little frightening.

As I look ahead, I really don’t know all the details about what’s coming and so I need to be cautious. But we do know that a bladed architecture is coming, and we will have an industry-standard building block from which any number of configurations will be built. We will see multiple operating systems in these packages, and we will have a much more efficient “green” product. But among the packages, there will be NonStop and it will be a key part of the HP server strategy.

What we can't predict with any certainty is whether this new NonStop will find universal acceptance and whether they usher in a new era of growth – in selected niches or across a broad mixture of industries. The potential is definitely there however, and I am looking forward to their arrival.

My favorite baseball coach, Tony LaRussa, talks about how he prefers his players to get themselves into positions to "manufacture runs". Tony prefers to manage the game, one innings at a time, so that a number of runs can be scored over the duration of the game. He doesn't build a team that relies on individuals scoring home runs to win a game and becoming dependent on a single swing of the bat.

Returning to the Next-Gen IT story, they noted that "it's clear, that the next-generation of data center will be a bastion of virtualization, consolidation, and automation technologies". And so it is that I see the strategy of HP unfolding – no longer dependent on a single outrageously successful product, but rather, from the compilation of many successful packages built from a bladed architecture. We do know we are heading towards a future where the basics will be simpler, and where we can't load up on the data center staff to look after it all.

And this leads me to a couple of closing thoughts.

While Winston forecasts Integrity NonStop servers will include potentially multiple Linux components – can we realistically rule out other HP Integrity server offerings not including NonStop components? Can we rule out NonStop becoming part of every BCS offering?

I can already see a day when the option to run a data base may come complete with a NonStop foundation – but would there be any limits to how Nonstop evolves? I just don't think so and see the NonStop server line living well beyond any previous expectation I we had.

Tuesday, February 12, 2008

"My Wish" for NS Blades

This weekend I drove out to the Southern Californian coastline, and it was a real change from earlier weekends spent in Colorado. Sure it's winter across the US but out along the coast, it was balmy days in the mid 70s. I was driving the roadster, with the top pulled back, and the picture here is of me beside the Pacific Coast Highway –or PCH as it is known locally - just north of Malibu. The drive reminded me of those articles you find in the airline's magazines – three perfect days in "wherever" - and I had days about as perfect as you could experience anytime in LA.

During the previous week, I was emailing folks in the UK as I was preparing to catch up with them next month. As these emails had been flying around, I made a minor mistake in the subject line – rather than proposing to meet for "drinks", I was suggesting we meet for "rinks"! Of course this led to Sean Bicknell of XYPRO, and a BITUG committee member, to respond "I'll bring my ice skates then shall I!?!?" Driving the PCH, enjoying the sun and thinking how freezing it was back in Colorado, led me to recall the lines of the song "All Star", from the first Shrek movie, by Smashmouth: "The ice we skate is getting pretty thin, the waters getting warm so you might as well swim."

Talking about skating on thin ice, I have been taking a really close look at HPs press releases, and have pulled down a number of Roadmap presentations from

the ITUG web site. I have also been deep into other HP web site pages. I really wanted to get a better understanding of HP's plans for NonStop. I am putting together a new presentation to use at ITUG events, and it all started as I recalled a comment made at last years' HPTF&E by Martin Fink.

Martin was talking about how pleased he was that NonStop was being openly accepted by BCS development. He related that, as the blades developers working on the new cClass bladeserver chassis began to leverage the IP of the NonStop group, they asked Martin "Can we call it the NonStop Backplane"? Martin told the audience his response was "why not – go right ahead"! And today, if you go to the web and pull down the pdf of the HP BladeSystem c-Class architecture "technology brief", you will find on page 8 a paragraph titled "NonStop signal midplane provides flexibility".

As we continue to hear more about the first bladed architecture products, something a lot of us simply call NS Blades, it too will be using this chassis. When I asked a well-placed NonStop engineering manager whether this was the case he responded "the NS Blades system does indeed use the new c-Class platform (BladeSystem c7000) enclosure", adding enthusiastically "everything to blades"!

Predicting the future at HP, particularly when it comes to the future of NonStop, is a risky business and I have spent my time out skating where the ice is thin. I haven't always been right but I have been close most of the time. But rather than talking about what will be coming from product management, it may be more interesting to lay out my own wishes for the NonStop family.

I began to lay out some of my thoughts in earlier blogs and late in January (The check-in desk two-step!) I wrote about Mission Critical applications and of their long association with the NonStop platform. I went on to add in the next posting in January (NeoView; a new view?) how transaction processing will always have mission critical elements and that the support of these elements will always benefit from the NonStop architecture. And then in an early February posting (Disruptive technologies and radical innovation) I reported that NS Blades have the potential of becoming a classic disruptive technology in that its support for any OS, including NonStop, means that your choice of OS can be arbitrary and

driven solely by the demands of the application.

I have three items on my wish list and all involve the NS Blades. I am working with the belief that HP will begin shipping NS Blades sometime mid-to-late summer. I am also confident that these will utilize the c-Class blade enclosure (BladeSystem c7000). And so my first wish is to see HP BCS deliver on the slideware Martin Fink first unveiled as the “Shared Infrastructure Blades” package. This is where any mix of NonStop, HP-UX, Linux, and Windows Server OS’s will be supported.

The importance of this feature – and getting the manageability all sorted out – is that customers will be able to deploy any mix of web, application, and data base servers in a single enclosure. NS Blades customers will be able to add additional blades as growth in any one area dictates. Indeed, I am expecting that they should be able to reconfigure and “redeploy” under-utilized blades to other OSs as workload mixes change over time. The manageability component is very important – you only want a single command interface to all the blades. A separate console, per OS instance, kind of misses the point. And I am wishing for a programmatic interface into all of this as well.

My second wish is to see a hypervisor introduced where NonStop can be configured as a “guest OS” in much the same way z/VM is used on the IBM mainframe. The trick here is to see this introduced without marginalizing the traditional association between NonStop and the hardware with respect to being fault tolerant. In other words a processor, within a guest NonStop configuration, needs to know when to initiate a take-over. Perhaps exposure to the hardware may not be right answer these days – as so much self-checking and correction goes on at the chip level – but NonStop users will not want to see a hypervisor becoming a potential “single point of failure” either. While it may not be as simple as configuring a number of hypervisor processors under the management of a check-pointed “hyper-mon” process, my own logic wants to take me down this path.

If you check out hypervisors on wikipedia, you will see references being made to Type 1 (native, or bare-metal) and Type 2 (hosted) hypervisors. I am a lot more interested in the Type 1 implementations, and while some OSs may want to run

the hypervisor as a guest of the OS (e.g. Microsoft Virtual Server – a virtual machine on the Windows Server 2003 operating system), traditional NonStop users may view the properties of Pathway providing something similar!

But it is my last wish where I really want to go out on a limb. If you assume Martin is successful and a shared infrastructure blades package becomes available with a native, or bare-metal, hypervisor (NonStop as a guest with no loss of its NonStop attributes), then wouldn't it be advantageous to users if interrogation of the incoming transactions would direct mission critical transactions to NonStop, important informational but not quite mission-critical to a Unix or Linux, and voluminous inquiries to Windows? A variation on today's workload balancing products, but supporting a transaction profiling capability that once set up, learns about the overall mix of transactions, and automatically adjusts the OS configurations on the fly? No longer would you agonize over the ratio of processors assigned to any of NonStop, Linux, or whatever – the system would learn enough over time and adjust the composition accordingly!

Returning from the drive along the PCH, the song "My Wish" by Rascal Flats began to play and as I listened to the verse: "And if you're faced with a choice, and you have to choose, I hope you choose the one that means the most to you. And if one door opens to another door closed, I hope you keep on walkin' till you find the window". Followed by the chorus "Your dreams stay big, and your worries stay small". I just knew they were singing about NS Blades!

With the flexibility inherent with the design of NS Blades, and the doors this will open for us all, I can't see placing any limits on how big our dreams will be! Or, as Smashmouth sang, "My world's on fire how about yours, that's the way I like it and I never get bored"!

Saturday, March 1, 2008

Thirty years on - a new generation!

I had only been back in Sydney for just over a year when IBM launched the “Glendale Series” of mainframes – the IBM 4331 and IBM 4341. The year, 1979, was a very exciting time for IT in Australia as the federal government was allowing a capital investment allowance of 40 percent, as an incentive to industry, which included new computer purchases. However, this great deal was due to expire on June 30th, 1979.

So, with the New Year, corporations were waiting for the announcement of this new mainframe and were anxiously standing by their fax machines eager with anticipation. As January 30th dawned on the US East Coast (January 31st in Sydney), the news finally broke and the price points were stunning. It was a circus! IBM worldwide quickly became overwhelmed with orders and introduced the first lottery – allocating new machines according to a random draw.

These new mainframes featured high-density logic chips, with up to 704 circuits per chip. And with the older IBM 370/158 rated as a 1.0 “old style” MIPS machine (based on IBM’s figures of a cycle time of 115 nanoseconds, which is about 8.7

MHz.), this new one, the 4331 mainframe, rated at about .3 MIP, and its big brother, the 4341, at a shade under 1 MIP – an incredible price/performance achievement at the time.

I was reminded of this as I was in Orlando, FL at an IBM user event to see the launch of IBM's tenth generation mainframe – the IBM System z10 Enterprise Class (z10EC), and I was anticipating its launch in much the same way as I had some thirty years before. As the events of Monday wound down, and the trade-show floor emptied, IBM engineers descended on the z9 mainframe that was on the stand supporting attendee labs, unplugged it from the storage, and after 4 hours of frantic work replaced it with the new z10. By midnight, it was running, and as the attendees returned for early morning labs they found themselves running on the new mainframe. Whereas the other four sites around the world, where IBM held its press conferences, had only z10EC shells - it was at the user event where a fully operational z10EC made its public debut!

So much has changed since the Glendale Series rolled out, but the performance improvements provide the biggest contrast and with the z10EC, it's all about the new z10 Processor Chip and the z10EC Multi-Chip Module (MCM). IBM is leveraging the POWER6 Dual-Core chip – no, it's not the same chip (CISC for z10EC Processor Chip), as its not s a RISC chip as you will find in other IBM offerings – with IBM suggesting that they are "siblings, (but) not identical twins" as there is small number of elements common to both.

Instead of the high density logic chips, with 704 circuits, that we saw on the 4331 and 4341 mainframes, we now have ceramic MCM blocks (103 ceramic layers) each with 5 Processor Units (PU) chips, where each PU had four cores, and where there were now 994 million transistors per PU chip – somewhere near 5 billion transistors associated just with the PUs on the MCM block. Throw in the two Storage Control (SC) chips that add an additional 1.6 billion transistors, and you end up with something like 7 billion transistors per MCM block. A fully populated z10 with four "Books", IBM-speak for very fat blades, each with one of these MCM blocks, gives you a total of 64 PUs that are capable of processing some 30,000+ old-style MIPS.

Putting this into context, my rough, back-of-the envelope calculation, suggests

that just one of these mainframes now has more MIPS than was shipped across the entire life of the Glendale Series. All in a machine that's not much bigger than the earlier 4241 stood upright! Certainly not the behemoth many of us had grown accustomed to seeing holding court in the middle of the data center. And did I mention that the contrasting color flash down the side of this new mainframe has been changed from red to green – a stark reminder of everyone's concerns over energy requirements.

But in the press reports that came out on the day of the announcement, there were still some concerns about IBM's continued investment in mainframes. According to a report by Associated Press "analysts said IBM's advances in chip technology and software are helping the mainframe stay competitive against lower-cost competitors. But they caution that because of price IBM still faces challenges in luring in new customers." And in the same story, they quoted Brad Day, a Forrester Research Vice President, as having said "this is definitely not a slam dunk - the math still has to be there. The life-cycle-cost-of-ownership argument still has to be there."

The picture I have included here is of me beside the working "internals" of the z10EC. What you may not pick up all that easily are the very fine optic cables coming down to the multiplexers that support all the external storage. These are running InfiniBand – a first for IBM. When I was talking to an HP NonStop engineering manager recently, and we were talking about InfiniBand, he talked openly about how "the ServerNet footprint is all over InfiniBand." While IBM is only using it to connect to I/O subsystems, as well as supporting its own Parallel Sysplex interconnect option, I found it highly encouraging for IBM to now deploy a technology with which all of us working on NonStop are familiar.

IBM continues to make extensive use of virtualization – the concept of logical partitioning, or LPARs, is the first layer of abstraction above the metal and is based on earlier releases of VM. "LPARs, on (the z10EC are) not really virtualization with overhead, because it is truly built into the hardware," suggested a colleague attending the show and helping me sort through all the terminology. In an LPAR you can run IBM's zVM hypervisor with its support for any number of guest OS's including zOS and zLinux. While you may want to configure a "relatively small number of LPARs, compared to the number of VM guests," he

added, “there is a significant advantages when running as a VM guest.” IBM’s support of virtualization remains a very powerful feature of the z10EC and is a technology where I do anticipate hearing more from HP, including HP NonStop, in the future.

The timing of my trip to Orlando, coming as it did back-to-back with my trip to South Africa for SATUG, where HP and Intel had provided updates to their product roadmaps, I couldn’t help but notice many similarities, and how the technologies of HP and IBM appear headed in very similar directions. I know that this statement might surprise some of my IBM colleagues – but the more I looked at the z10EC and the more questions I asked, the more I saw how the “top-of-the-line” offerings from both vendors were beginning to look alike.

For me, the biggest difference comes with the decision taken by HP to standardize on single chip architecture, and to ride that vendor’s roadmap. IBM still uses different chip architectures and while there’s some shared portions, it still requires two different design groups and two separate fab facilities. I don’t get that, and I can’t see that being a sustainable model for the long haul!

The NonStop architecture may be better positioned than the IBM mainframe as HP looks to position NonStop as a “configuration option” on future HP bladed architecture products. In my recent presentations I have simply been stating: “rather than asking the question ‘who will be using NonStop’ perhaps the more appropriate question may be ‘when do I take advantage of the NonStop?’” There will always be a subset of transactions that will be viewed as mission critical, and routing them to mission critical applications deployed on bullet-proof NonStop servers, at a much lower price point, will continue to give HP an edge.

I realize that this blog posting is a little more technical than I have written in the past, and I trust you will indulge me on this occasion. But has the arrival of z10EC changed my view on any of this or become a disruptive technology that will force HP to change its roadmap – well, not exactly! Without taking anything away from IBM – it’s a great system with some remarkable technology that will be welcomed by many of IBM’s customers – it’s not a must-have “killer system”! And I just have to wonder if the green stripe down the side of the box is a slight tinge of envy as much as a testament to its reduced environmental impact!

Sunday, March 9, 2008

From higher altitudes!

I took a weekend break in Boulder. The first for many weeks, and I needed it. I had just returned from Orlando and on Sunday, it was back on the plane for a trip to Prague – one of my favorite cities in all of Europe.

The weekend along the Rockies proved to be unseasonably warm so it was a good time to grab the motorcycle for a ride up into Wyoming and to revisit some of my favorite front-range roads. The picture I have included here is of me parked at a rest stop just outside Laramie, as Interstate 80 (I80) crests the 8,600 foot summit – the highest pass on I80 as it connects New York with San Francisco.

While the temperature for most of my trip had been in the 70sF (20C), passing over this summit brought it down to the low 40s (5C) with winds pretty constant at 45 mph (60+ kpm). It was also at the midpoint of the 300 mile (480k) loop I had chosen and the views from this altitude were spectacular. I took the time to have a good look at the surrounding mountains and to peer down into the alpine pastures and valleys.

At one point on this ride across Wyoming I had an anxious moment. I just didn't know for sure where the next town was and, with no gas gage on the motorcycle, I was not quite sure if I had enough gas to make it into Laramie. As luck would have it, I came across the town of Buford and it had a gas station. It has a population of 1, making it the smallest township in the U.S. Finding it was

completely by chance and I am sure I would have been less stressed if I had have consulted a map before leaving Cheyenne.

I suppose it's not the smartest thing to do, but I never take maps. Big-twin cruiser riders (as opposed to our touring brethren) never need maps – it just doesn't look good to be seen stopped along the side of a highway, checking a map, and looking lost. We just keep riding till something looks familiar, and we don't turn around just to be safe. And we never, ever, stop at gas stations to ask directions!

As I was sitting in a Starbucks in Ft Collins I heard a familiar song, by Jim Morris, called Southward where he sings "It was chilly through Virginia, but I hit the Outer Banks, and my attitude was toasty as a fire … Took the Cedar Island ferry on a morning bright and clean, what's behind me never crossed my mind at all." Listening to the music also reminded me of something similar from the '70s movie - Gumball Rally. In perhaps the most memorable scene Raul Julia, playing the Italian race car driver, Franco, delivered his famous one-liner "the first rule of Italian driving?", as he tore off the rear-view mirror, "what's behind me is not important!"

Yes, I had ridden across mountain summits and I had been cold. I had taken time to enjoy the scenery from that altitude. But now, with my latte and a warm corner to curl up in, what was behind me was no longer important and riding without a map hadn't proved disastrous. As I mellowed, and began to anticipate the next leg of my ride, I began to think about all the roadmap presentation I have seen of late.

I had sat through roadmap presentations on NonStop at SATUG, as I had sat through Intel's chip roadmap. I had sat through IBM's plans for the System z as the z10EC mainframe rolled out last week. I had sat through many GoldenGate presentations as well. Yes, product roadmaps are very important, and they provide a wealth of information if you just know what to look for! There's not a product manager on earth that doesn't get passionate about something they are directly involved in, and they are always eager to press home the point of how such-and-such feature is going to dramatically improve the fortunes of the user, the technology, and their own development team.

And it reminded me that sometimes, watching presentations on product roadmaps, there was much to be learned from what was not said as from what was actually said. Likewise, you could also learn a lot from recalling what had been included in previous presentations, but had never materialized as product. In today's litigious world, what actually makes it onto a roadmap has been so well-sanitized that it often lacks drama or excitement. Resources are committed - there are real dates, test plans are in place - and you can anticipate an early adoption program following shortly.

From the time Martin Fink stepped to the podium in Berlin to address the European ITUG event, it was clear that the paths of the Open Systems / Linux and NonStop would intersect. At first, I took his comments to imply that at some point, we may have elements of NonStop integrated with Linux. Something modeled on the work first done during Bill Heil's watch - the ServerWare (later NonStop Software) project. Back in the mid '90s a lot of energy went into overlaying Windows with enough of NonStop's message system to provide Microsoft with the foundation for a clustered fault tolerant offering. Perhaps, what went on before should be considered more seriously. Perhaps what was behind us remains relevant, after all!

While this lab project never saw the light of day, and ServerWare development has been abandoned following the Compaq acquisition, I suspect that much of what was learnt has now been revisited but this time with a focus on Linux rather than Windows. While Martin may not bring a fault-tolerant Linux to market I wouldn't be at all surprised to see some lower-level components and routines within NonStop draw from Linux, and the experiences gained from ServerWare.

Another similar project I really thought was fully funded was what Cupertino internally referred to as Hybrid Super Clusters (HSC). Fueled by what some of the biggest customers had been doing (e.g. Sabre, as well as some stock exchanges), and where a mix of NonStop, Linux and HP-UX were deployed essentially as clusters, there was some consideration given to productizing some of this. As Martin presented his vision of "shared infrastructure blades" that supported a mix of blades and operating systems, I thought that this was really a sign of what was to come. I could envisage a marketplace that really could take advantage of such a package as and when it materialized.

Turns out, this was never on the product roadmap, in any official sense. From the lofty heights HP NonStop executives move in, this was more a framework for discussion than anything else. Or, was it? Recently, I was taking to a development executive who pointed out to me that “work (internally on HSC) has now been completed … one of the values that resulted from this project was our recently released NonStop Cluster Essentials product; it provides an integrated management interface for monitoring, managing, and controlling heterogeneous clusters of NonStop and Linux.”

This development I find very important. In the roadmap presentations given at SATUG, a senior product manager broke down the key areas of development into just three categories – Open Access, Scale and Availability, and Manageability and Compliance – and where the issue with Manageability and Compliance was the need to support cross platform deployments. We know that specialty controllers will begin to appear shortly – Linux on blades hosting communications and other I/O controller software - and for them to be successfully deployed, they would have to be seamlessly integrated into the overall manageability scheme.

Perhaps this is where we begin to see the pieces all fitting. As we see the Intel roadmap being followed, with dual-core, quad-core, and eventually multi-core chip technology being supported, we could see the morphing of NonStop with Linux appearing in a more subtle ways – integrated as part of the lower-level components of NonStop, as the OS of specialty controllers, and all managed from a single, integrated management environment. Is this what we expected after hearing from Martin? Is this the kind of Hybrid and Hybrid-Clusters we were anticipating? Again, not really! But on closer examination – is this of more immediate value. Absolutely!

Projects, whether on the roadmap or not, always are leveraged at some point and the products we finally see being delivered frequently employ components and routines reminiscent of earlier lab activities. I can only now speculate that when we see future bladed architecture product offerings then, to paraphrase the Rolling Stones, “we may not get what we want … but we may just get what we need!”

About The Author

I am Richard Buckle and I have enjoyed a long association with User Groups dating back to the late 70s. I am particularly passionate about the grass-roots activities of communities and have enjoyed working with regional groups in many countries.

I spent 7 years as a Director on the board of ITUG and served as its Chairman for 2 years (2004 - 2005). I have worked in the Tandem / NonStop marketplace for more than two decades, initially as a Tandem Computers employee before it was acquired by Compaq, and then for Infrastructure ISVs supporting the NonStop platform. Before joining Tandem, I had spent the previous two decades in the IBM community as an end-user as well as a vendor. I am currently the Senior Director, Business Development, for GoldenGate Software out of San Francisco.

I am active in a number of social networks as well as a frequent visitor of virtual communities. Most recently, I have begun blogging on a new social network “The HP User Group Community Social Networking Site” that can be found at:

http://hpusercommunity.org/mypage.aspx

Finally, I am very grateful for the continued support of the senior management of GoldenGate – without their support I would not be able to undertake this task. On several occasions I have solicited commentary and each time they have been extremely cooperative and timely, and for this I want to thank everyone involved at GoldenGate.