



Market Trends

The HP/UX/Itanium to Linux/x86 Multi-core Migration Trend

Executive Summary

Over the past few years, many customers have implemented x86 servers for applications that might have previously been implemented on Unix servers. However, based on revenue and market share, it's clear that those Unix losses have been at the expense of Oracle/Sun's and HP's Unix systems sales — not IBM's. Hardest hit by this market shift are Oracle (its SPARC server sales dropped 6% last quarter) and Hewlett-Packard (HP) — a company whose flagship Itanium-based server line has come under serious fire as of late (Microsoft, Red Hat, and now Oracle have all withdrawn future development on the Itanium platform).

Another analyst company has noticed this trend as well. IDC, a leading information technology (IT) research firm, recently published a report entitled: “Unix Migration Accelerates in the Face of Rising Adoption of Windows and Linux Solutions on x86 Servers” in which it also describes this Unix to Linux/x86 migration trend. (Their report is based on telephone interviews with 407 mid-sized and large organizations. We confirmed this trend by speaking to several vendors, including Red Hat, and IBM).

To us, the migration situation at Hewlett-Packard is the most interesting migration scenario. HP appears to be in denial that x86/Linux servers pose a huge threat to its HP/UX/Itanium-server base — as exemplified in this HP document: <http://h20195.www2.hp.com/v2/getdocument.aspx?docname=4AA3-1768ENW.pdf>. In this report, Hewlett-Packard describes how its new DL980 G7 Intel Xeon multi-core-based servers make “an excellent platform for scale-up consolidation and for virtualization of legacy UNIX® and Linux workloads”. It goes on to point out how its DL980 G7 is a good alternative to IBM's AIX (Unix) Power Systems environment and Oracle's Solaris (Unix/SPARC environment) — but makes no mention of migration from HP/UX to Linux on the DL980. Are HP's HP/UX/Itanium-based servers somehow invulnerable to Linux/x86 erosion?

Given HP's earnings and ecosystem struggles, it should be clear to even the most casual observer that HP's Itanium-line of servers (and the related HP/UX operating environment and accompanying system software) are in serious trouble. Itanium architecture (the basis for HP's Integrity line of servers including Superdome) is under attack from Intel's own x86 multi-core Xeon architecture — and if Itanium sinks, HP/UX will follow.

In this *Market Trends* report, *Clabby Analytics* takes a closer look at HP/UX to x86/Linux migration. And what we conclude is that:

- HP's business critical systems group (the organization that sells HP/UX/Itanium-based servers) is tremendously exposed to erosion by x86/Linux servers;
- Due to certain hardware and software shortcomings, HP's x86 servers are not a “shoe-in” for capturing all of its customers' migration; and,

The HP/UX/Itanium to Linux/x86 Multi-core Migration Trend

- HP is at a competitive disadvantage as compared with IBM from a systems, systems software, and developmental perspective (because HP does not control its own software destiny).

The Situation

Based on discussions with leading IT vendors as well as numerous IT buyers — and also confirmed by other research firms — *Clabby Analytics* sees a clear migration trend away from Unix on SPARC and Itanium architectures. And, as for IBM's Power Systems (IBM's Unix server line), Power Systems seem to be benefiting from the overall Unix migration trend (in Q1 2011, IBM reported 164 competitive wins against HP — with most Unix migrations going to AIX on Power Systems).

Still, when we probed members of IBM's Migration Factory (a professional services organization that has performed over 6,500 competitive displacements over the past 5 years), we found that approximately 20% of their Itanium migrations are going to Linux on x86 — and that this number appears to be growing quickly quarter-by-quarter. Further, it is worth noting that not all of IBM's migrations use their Migration Factory services, especially migrations to x86 architecture — so, it is our belief that the actual number may be closer to 30%.

Why Is This Happening?

When Intel introduced its x86 multi-core architecture (code-named “Nehalem”) back in 2008, it very significantly changed the competitive dynamics in the server market. Multi-core technology (which places two or more processor cores onto a single computing complex) gave x86 processors (specifically Intel Xeon-class processors) the ability to scale in terms of performance to levels that could compete with midrange and high-end processors (such as Itanium, SPARC, and POWER processors). With this change, IT buyers could now purchase highly-scalable information systems based on low-cost, industry standard x86 architecture.

But, in order to exploit these x86 multi-cores, IT executives immediately recognized that they would most likely have to shift operating environments away from Unix to Linux or Windows. (Linux and Windows dominate the x86 market — and HP/UX and AIX don't run on x86). Some executives immediately seized the opportunity to migrate away from Unix to Linux/x86. Others took a wait-and-see approach...

If Itanium Fades Away — Then What?

If Itanium sales continue to struggle at Hewlett-Packard, HP will have little choice but to cease investment in Itanium system design/sales/marketing. As for HP's Itanium customers, should Itanium fail, we see three migration options:

1. *Migrate to a new version of Unix on a more viable platform (and to us, this means moving to AIX (IBM's Unix) on the Power Systems platform;*
2. *Migrate to IBM's System z (mainframe) environment (which runs several operating environments including z/OS, zVM, and Linux); or,*
3. *Migrate from HP/UX on Itanium to Linux on x86 architecture (because HP/UX does not run on x86 architecture).*

The HP/UX/Itanium to Linux/x86 Multi-core Migration Trend

Unix-to-Unix Migration

Hundreds of HP/UX customers have already opted for option number one (migrating from HP/UX on Itanium to IBM's AIX on Power). We describe this market trend in depth in this report on IBM's Migration Factory (the professional services organization that has performed over five thousand migrations over the past four years): For an in-depth look at Unix-to-Unix migration, please see our recently published report at:

<http://www.clabbyanalytics.com/uploads/MigrationFactoryFinal.pdf>

Unix-to-Mainframe Migration

A much smaller number have opted for option number two — migration from HP/UX/Itanium to System z (mainframe) architecture — but we think this situation is about to change due to the introduction of IBM's new z114 low-cost mainframe server environment (hardware pricing starts in the \$75K-\$100K range). We expect that HP/UX customers who require extremely high levels of Quality-of-Service (QoS) will be enticed by the System z's best-in-the-industry virtualization, its MTBF (mean time between failure), and its security characteristics (the System z is the only commercially available server to have ever achieved EAL Level 5 security certification). For more on this topic, please see: <http://www.clabbyanalytics.com/uploads/z114Final.pdf>.

Unix to Linux/x86 Migration

The third (and most rapidly growing) group of HP/UX users are those who have chosen to migrate from HP/UX/Itanium to Linux or Windows on x86 multi-cores. This group believes that Intel's Xeon-class x86 multi-core servers combined with industrial-strength Linux, are now capable of running workloads formerly assigned to HP/UX/Itanium servers.

Choosing the Right x86 Multi-core: A Look at IBM System x vs. HP x86 Environments

When we compare one server vendor's platform to another, we tend to look closely at system cost, designs (how they use memory, the systems bus, redundant components, etc.), and software. After comparing HP's x86 system designs and systems software to IBM's x86 designs/systems software, we have concluded that:

1. When comparing "systems costs" (not just server/chassis costs — but also related costs such as interconnect, license, and power costs), IBM has several distinct design advantages as compared to HP servers that lead to significant cost savings;
2. IBM's x86 server designs are more innovative than HP's designs;
3. HP is at a distinct competitive disadvantage in the area of system software; (the software stack that typically includes middleware, management software, and other extensions such as memory management); and,
4. IBM offers better systems designs (better use of systems resources, access to more resources, more flexible expansion options [especially related to memory], and more configuration choices).

Costs

We measure a server's true cost by considering server hardware costs, chassis hardware cost, interconnect costs (switches and adapters), license costs, power costs, and maintenance costs.

From our perspective, server prices are generally the same (x86 vendors tend to match their competitor's server prices). Likewise, power costs are usually within \$1000 of each other.

The HP/UX/Itanium to Linux/x86 Multi-core Migration Trend

And, there are some differences in maintenance (IBM can be around \$5000 less than HP in blade maintenance, for instance). Just remember that, if ordering a large number of blades, these costs differences can really start to add up.

Chassis prices are usually pretty close too (within \$1000-\$2000 of each other). But when comparing chassis, there is more here than meets the eye. If one vendor's chassis has fewer slots than the other's, then adding the next blade may require an entirely new chassis if no slots are available. Adding this chassis may require a new rack, or a new row — and certainly more power. Pay close attention to the number of slots and density of the chassis that you buy. Buying less dense chassis can have a significant, negative financial impact.

Another significant difference between IBM servers and HP servers is in the area of interconnects between blades and chassis — and chassis and other chassis or other devices. In November, 2007, we published a report entitled “Why You Need to Pay Attention to How Your Blade Vendor Handles Virtual I/O” that explains why HP interconnects cost more than IBM blade interconnects (this report is still available at: <http://www.clabbyanalytics.com/uploads/PayattentiontovirtualIOfinalrevisionupdatefinal.pdf>). In essence, this report found that HP's hardware-oriented approach to the management of virtual I/O led to more expensive hardware and management solutions than IBM's software-oriented approach.

*Read this report for more technical details — but if you're primarily interested in the bottom line, it is this: the last we looked, when comparing 10x HP BL620 G7 blades (with two chassis due to chassis extension limits versus 10x IBM HX5 using one chassis (both systems having comparable memory) — **the cost difference for the interconnect hardware was about \$129,000!** (HP adapters and switches were approximately \$198,000 versus IBM's \$69,000). Our advice: pay very close attention to the interconnect costs of each environment...*

Also bear in mind that with the innovations described in the next section, buyers of IBM servers can load more virtual machines — meaning that IBM buyers can get higher utilization out of their IBM blade servers than HP buyers can. With superior virtualization/memory management, IBM buyers are essentially getting groups of virtual servers for free.

Innovation

We have long been fans of IBM systems designs. IBM is a company that has mastered design architecture — starting at the high-end with mainframes and then cascading advanced designs to Power Systems and System x architectures.

As we look more closely, System x has numerous innovative differentiators including:

- **MAX5** – Memory expansion with the external MAX5 memory option, decouples server memory from system processors to allow customers to optimize server performance by adding memory rather than buying additional servers. Memory capacity can expand up to 64 DIMMs standard and 96 DIMMs using MAX5 memory expansion per 4-socket server.
- **eXFlash** – Flexible hot-swap storage with up to 8 HDDs or up to 16 SSDs with eXFlash technology provides solid state drive technology that delivers faster I/O, with greater density and improved reliability.

The HP/UX/Itanium to Linux/x86 Multi-core Migration Trend

- **FlexNode** – Delivers the ability to re-deploy your server on a project-by-project basis for superior asset utilization and workload management. For example, a 4 socket server can be re-deployed as a two 2 socket server in order to match the server characteristics to the workload as it varies throughout a day.
- **Scalability** – Expand from a two-processor system up to four processors. Add a second system to create an eight-processor system. Start with two memory DIMMs and expand up to 192 with a dual-node system. Plenty of scalability...
- **eX5 memory management environment** — Some System x servers include a specially designed memory management subsystem (microprocessor and software) that manages swapping data in and out of memory — offloading the CPU from having to perform all memory management tasks. This memory management subsystem is a system design characteristic that plays an important role in System x data handling and performance.
- **IBM BladeCenter Open Fabric** is an integrated server I/O portfolio that offers high performance interconnects and rich management tools. It supports open standards and industry interoperability across multiple I/O fabrics, including Ethernet, iSCSI, Fibre Channel over Ethernet (FCoE), Fibre Channel, InfiniBand and Serial attached SCSI (SAS).
- **Specialized virtualization management ASIC** (application specific integrated circuit — a specialized microprocessor) on some of its x86-based servers that offloads the CPU from having to process virtualization instructions. By offloading the CPU from having to do this work, the CPU can be focused on other tasks (number crunching, for instance). Fujitsu has similar functionality as part of its blade systems offerings — but none of IBM's major x86 competitors offer this.

The Systems Software Situation

In January, 2001, HP acquired Bluestone, a maker of middleware software, for \$467.6 million. In June, 2002, Hewlett-Packard announced its plans to exit the middleware business and rely on software partners such as Microsoft, Oracle, and BEA to provide the middleware stack that would run on their servers.

We saw this as a huge strategic mistake — and told the press so at that time. Here's why: computer makers that build their own middleware/software stacks can optimize their stacks in order to create major performance advantages for applications that run on their platforms. As an example, IBM's "smarter systems" (packaged, optimized hardware/software solutions) employ streamlined software paths through the company's middleware to the company's databases in order to deliver improved performance that is orders of magnitude greater than solutions that have not been software pathed. Software pathing, accordingly, creates a huge competitive advantage for IBM.

To be fair, HP may not provide middleware, but it does offer some systems software (systems/software management, virtualization, power management, several operating environments, and more). The problem, however, with HP's position in the systems software market is this: if Itanium is made redundant by Intel Xeon processor encroachment, then HP's system software strategy starts to fall apart because a lot of HP's systems software resides on HP/UX — an OS that doesn't run on x86 hardware.

The HP/UX/Itanium to Linux/x86 Multi-core Migration Trend

Comparing x86 Hardware Platforms

When we evaluate hardware platforms the first thing we consider is the workload that is to be run on a given platform. (Mainframes run heavy input/output applications/workloads better than any other system — while x86 architecture excels at the processing of many, light, fast thread workloads [and has also become an excellent architecture for compute-intensive tasks]).

When we compare x86 systems designs, *we don't start by looking at microprocessor characteristics* because the same microprocessors run across all the major vendors' systems. Instead, what we look for first and foremost is how much physical memory is available for application/data processing — and how that memory is managed. Our next point of comparison is the systems design itself. We look at the bus architecture; redundancies (for instance, redundant power supplies and fans); chassis/rack designs; associated interconnect (networking) hardware and software; flexibility, and configuration options.

As we compared HP's x86 systems design to IBM's, what we found, in general, is that HP does not offer access to as much memory as a comparable IBM System x (primarily due to IBM MAX5 memory management innovations); that HP does not match IBM in configuration options; and that there are big differences in chassis/rack designs (size/configuration options).

2 Socket Servers

In Figure 1, we compared IBM's 2 socket x86 servers to HP's equivalent 2 socket blade environment. And what we noticed immediately, major differences in memory configurability, flexibility, system design, and availability became readily apparent (see Figure 1).

Figure 1: HP vs. IBM in 2 Socket Systems

Point-of-Comparison	Hewlett-Packard	IBM	Impact
Memory	32 DIMM only Capped memory	16 or 40 DIMM <u>configs</u> MAX5 adds 6 memory buffers MAX5 adds 24 <u>add'l</u> DIMMS	IBM: more virtual machines leading to <i>higher</i> utilization; up-to 40X8GB DIMMS!
Flexibility	Monolithic design	<u>FlexNode</u> Partitioning	IBM <u>FlexNode</u> provides workload isolation, scheduled partitioning, s/w license optimization
System Design	8 Blades in 10U Chassis	14 Blades in 9U Chassis	IBM: more computing power in less space
Redundancy (for Availability)	Unrecoverable failure brings down system	<u>FlexNode</u> failover guards against this	IBM: greater availability

Source: Clabby Analytics — July, 2011

The HP/UX/Itanium to Linux/x86 Multi-core Migration Trend

4-8 Socket Servers

As we examined HP versus IBM 8 socket options, what we found was:

- IBM enables its customers to purchase a 4-socket x3850 that can scale-up to 8-sockets when needed. HP sells a 4-socket DL580 — but scaling up requires a different machine (the DL980).
- IBM's 8-socket 3850 offers more I/O bandwidth with 4 I/O hubs vs. HP's 2-3 I/O hubs — and IBM offers more x8 slots.
- The IBM 3850 design was more elegant (and simpler) as compared with the HP DL980. IBM's 3850 X5 uses 2 chassis and 4 cables; HP's DL980 uses 2 XNC boards, 2 interface boards, 2 link boards and 4 wrap cables.
- IBM's 3850 offers more HDD capacity with 16 2.5" hard drive bays vs 8 2.5" bays on the HP DL980.
- IBM's 3850 can offer better performance and consolidation than the DL980 by taking advantage of disk speed using IBM's eXFlash high IOPS 1.8" SSD (solid state drive) technology.
- IBM's 3850 uses fewer PDUs (power distribution units)/line cords (IBM has 4 power supplies — HP's DL980 uses 8).
- IBM's 3850 is easier to service (this may help account for the maintenance cost differences described earlier). IBM's 3850 allows top access to CPU and memory versus front access on the HP DL980 (and back access to 3rd Boxboro IO hub option). And,
- IBM supports x8 DIMMs (memory modules) that use less power than x4 DIMMs.

To us, the biggest difference is in the amount of memory that can be accessed by the operating environment (to enable more virtual machines to be run), and by applications that want to place more data in memory in order to process that data more rapidly. IBM's system design and its innovations lead to the ability to install and exploit more memory than comparable HP's servers. IBM's 4 socket blades and its x86 rack designs show the same type of characteristics — access to more memory, better memory management, better (in our opinion) redundancy and availability characteristics, and denser designs.

Professional Services

Both Hewlett-Packard and IBM offer a wealth of professional services for x86 servers. IBM's services include hardware maintenance services, managed technical support, server optimization and integration services (including server efficiency studies), network integration services, service management strategy and design services, software, storage optimization and integration services, and storage software maintenance and support.

IBM's services revenue last year was \$56 billion; HP's services revenue was \$40 billion.

Customer Scenarios

One of the most difficult tasks that research analysts undertake is trying to get real world customers to describe what technologies they are using, how they are using them, and what business results they are seeing. Interestingly, finding mini-case studies on Unix to Linux/x86 did not prove to be very difficult. For example:

The HP/UX/Itanium to Linux/x86 Multi-core Migration Trend

- Red Hat has also posted numerous customer case studies that describe HP/UX migrations to Red Hat Linux. These case studies can be found here: http://www.redhat.com/migrate/unix_to_linux/hpux-to-rhel/success-stories.html
- When searching IBM's website, we also found this example (MEVA — a large construction firm) that had moved from HP/UX to Linux on x86 architecture: <http://www-01.ibm.com/software/success/cssdb.nsf/CS/STRD-84SFQA?OpenDocument&Site=linuxatibm>

We expect to update this report later in the year with even more customer Unix to x86 migration scenarios.

Summary Observations

At *Clabby Analytics*, we do not believe that any single microprocessor or systems design handles all jobs optimally. We believe that IT buyers should closely examine their application characteristics and then decide which microprocessor/systems design can best service individual applications. QoS requirements also need to be weighed — as they too play a very important role in choosing a particular system environment.

To reiterate, IT buyers who are considering migration to a new platform should look closely at:

- *The requirements of the given workload (some workloads run best on servers that rapidly process thousands of lightweight application threads [such as on x86 architecture] — while others may have heavy I/O requirements and may run best on mainframe architecture);*
- *System design characteristics — such as component redundancy (to ensure high-availability), memory management (to provide applications with the ideal amount of memory required to execute tasks); processor characteristics (some processors are better at handling parallel computing tasks, while others are better at processing data intensive workloads and/or executing serial workloads); and so on...*
- *QoS requirements — Quality-of-Service requirements also play an extremely important role in helping determine which platform architecture to adopt. For instance, most banking and financial workloads require a high degree of security. Architectures that cannot meet security requirements (or performance targets, or availability requirements, etc).should not be considered to replace existing workloads that run on Unix.*

If workload analysis determines that x86 multi-cores can do the job, we suggest that IBM's System x line be evaluated. We find it superior in terms of breadth and depth to Hewlett-Packard and Oracle x86 offerings for a variety of reasons — particularly in system design, configuration offerings, and related systems software.

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

© 2011 *Clabby Analytics*
All rights reserved
July, 2011

Clabby Analytics is an independent technology research and analysis organization. Unlike many other research firms, we advocate certain positions — and encourage our readers to find counter opinions — then balance both points-of-view in order to decide on a course of action. Other research and analysis conducted by Clabby Analytics can be found at: www.ClabbyAnalytics.com.