CLABBY ANALYTICS

Research Report

The Impact of Placing IBM's Integrated Facility for Linux on IBM Power Systems

Executive Summary

IBM's Integrated Facility for Linux (IFL) is a processing environment that integrates IBM processors with memory and virtualization facilities in order to create very efficient, dedicated Linux engines. These engines are usually deployed on enterprise class systems that are used for database serving, business transactions, and batch processing. IFLs simplify the deployment and improve the performance of Linux – and run on IBM z and POWER microprocessors.

> ***The reason why IBM offers IFL facilities on its advanced processors/servers is to present its customers with more powerful and more scalable alternatives to Linux running on x86 architecture. In the past, Linux buyers have demonstrated a strong preference for Linux on x86 architecture. But with more power, more capacity, more scalability — and with very competitive pricing — IBM is hoping to change this preference.***

At *Clabby Analytics*, we believe that IBM's new Power Systems IFLs will be wildly successful for two reasons: 1) historical precedence; and, 2) the need for more powerful enterprise-class Linux servers (servers that can offer more capacity, better throughput, and richer Qualities-of-Service [QoS] than offered by alternative x86 processor/system designs):

- ***Historical precedence*** – since the introduction of IFLs on mainframes, IBM has seen a steady rise in IFL usage (at present, 36% of IBM System z users make use of IFLs — and this number is growing steadily). This strong adoption rate shows that the value proposition of IFLs clearly resonates with mainframe users – and we expect it to also do so with Power Systems users.
- ***The need for more powerful/richer RAS/richer QoS enterprise-class Linux servers*** –As applications running on Linux move from the edge-of–the-network (such as file serving, Web, and email) into the enterprise-class (mainstream business applications such as SAP, mobile, social, Oracle, business analytics, etc.), information technology (IT) executives know that they will need "enterprise class" servers that offer higher performance, greater scalability, higher reliability, availability, and serviceability (RAS), richer virtualization, and better security than has typically been found on x86-based servers. Power Systems are an excellent example of this type of RAS/QoS-rich enterprise-class system environment.

> ***If Power Systems IFLs mimic the mainframe IFL growth path, then it is reasonable to expect that Power Systems IFLs will be most successfully deployed in scenarios that can take advantage of Power Systems scalability (such as Linux consolidation, co-location, and cloud-in-a-box configurations). IFLs will also be adopted by IT executives who want to capitalize on Power Systems performance/throughput and Quality-of-Service (QoS) advantages over x86 architecture (advantages such as the deployment of mission critical workloads on Linux on high performance/highly secure Power Systems).***

In this *Research Report*, *Clabby Analytics* takes a closer look at IBM's new Power Systems IFL technology, its pricing strategy, related use cases and services. And based upon our research — as well as based on the mainframe IFL precedent — we predict that Power Systems IFLs will be hugely successful in the Linux consolidation server market, as co-location servers, as cloud/open source servers, and as advanced reliability, availability, serviceability RAS/QoS servers supporting mission critical Linux applications.

1st Things 1st: It Is Very Important to Understand How POWER- and x86 Architectures Differ

There are some very substantial differences in microprocessor capabilities and associated systems designs when comparing POWER-based servers with x86-based servers. These differences impact system utilization, processing speed/throughput, service level expectations, and more.

POWER Is a More Powerful/More Efficient Processor w/Access to More Memory

From a microprocessor perspective, the POWER microprocessor has several distinct performance and efficiency advantages over x86 processor architecture in processor speed, and processor efficiency. From a processor speed perspective, POWER can achieve a 4.4 GHz speed as compared with Intel's Sandy Bridge-EP Xeon processor at 3.6 GHz – giving POWER processors a twenty percent performance advantage. From an efficiency perspective, Power Systems can deliver twice as many threads per cycle as Intel's Xeon processors (4 vs. 2). In total number of threads per server, Power Systems have a huge advantage over Intel servers (a high-end Power System can process 1024 threads simultaneously; while today's high-end Intel servers tend to max-out at 64 threads). See Figure 1 for a graphical comparison of these processor environments.

> *To Clabby Analytics, one of the most important differences between POWER- and x86-processors can be found in the efficiency of the processing cores. More efficient processor cores means that more work can be done on each core — and this means that fewer software licenses will need to be purchased. It is possible to save millions of dollars in software costs by using Power Systems running Linux as compared with less powerful/less efficient x86-based systems.*

Figure 1: Major Differences in Intel Xeon vs. POWER Architectures

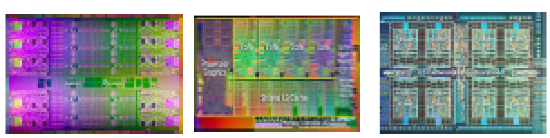

	Westmere EX	Sandy Bridge EP	POWER Systems
Clock rates per processor	1.8–2.67 GHz	1.8–3.6 GHz	3.1–4.4 GHz
Symmetric multi-threading per core	1,2	1,2	1, 2, 4
On-chip L3 Cache	30 MB	20 MB	80 MB
Maximum Memory per server	4 TB	1.5 TB	16 TB
Max Threads per server	160	64	1024

More processing power per core (Clock rates per processor; Symmetric multi-threading per core)

← Faster cache intensive workloads (On-chip L3 Cache)

← Run more memory intensive workloads (Maximum Memory per server)

← Larger servers for consolidation (Max Threads per server)

Source: IBM Corporation, December, 2013

Finally, as shown in Figure 1, POWER processors also offer more on-chip cache (which is important because the closer data can be located to the processor, the faster it can be processed). In Figure 1, Power Systems have access to four times as much on-chip cache than the nearest Sandy Bridge-EX processor.

> ***Simply stated, the reason that all of these differences are important is that POWER-based systems can process more work more quickly than high-end x86-based servers. More on chip cache enables Power Systems to execute cache-intensive workloads faster – and access to larger amounts of main memory can speed the processing of memory-intensive workloads. Further, the ability to process significantly more threads better positions Power Systems for Linux consolidation serving.***

Other Power Systems Advantages – Security, Virtualization, Memory Design, LPAR Bandwidth
In addition to the processing speed and on-chip memory/main memory advantages described in the previous section, Power Systems have several other design advantages when compared to x86 architecture. These include greater core efficiency; a stronger security track record, deeper virtualization infrastructure and management, denser memory, and faster I/O performance (especially logical partition [LPAR]-to-LPAR bandwidth).

Security
It is important to look at security from two perspectives: 1) risk management; and, 2) compliance. When we consider Power Systems from a risk perspective, the first thing we look at is the number of vulnerabilities that are reported each year. The number of vulnerabilities reported on Power Systems averages around 10-15 per year (and none are ever reported on the PowerVM hypervisor). Contrast this with the hundreds of vulnerabilities reported each year on Windows- and Linux-based x86 servers. Clearly, from a risk management perspective, Power Systems are a better choice.

The second area that we look at is security compliance. Security compliance is important from a regulatory perspective to meet governmental and special industry requirements – and it is also important for protecting personal data. To ensure security compliance on its Power Systems, IBM offers a product known as Power Security and Compliance (PowerSC) that provides monitoring and reporting tools optimized for virtualized environments on Power Systems. Functionally equivalent software can be found in the x86 ecosystem – but the level of integration of that software may vary. IBM's PowerSC has been specifically designed for and integrated with the POWER microprocessor and associated virtualization environment.

Virtualization
Virtualization is the process of logically representing physical resources such that unused capacity on those physical resources can be found, pooled and used by other programs. IBM's Power Systems virtualization environment is based upon its PowerVM architecture – a combination of hardware, firmware, and software that provides CPU, network, and disk virtualization.

The most important evaluation criteria when it comes to virtualization is how well does the virtualization software perform (how much sustained utilization can be had running the virtualization software). As this Solitare Interglobal report shows, the threshold for maximum sustained usage for Power Systems runs at almost 70%, whereas the utilization rates for x86-based servers max out in the 28-41% range. In short, what this means is that Power Systems users can run more work on their servers than x86 users, helping Power Systems users obtain a higher return on their investment. It also means that Power Systems users need not buy as many software licenses (because each POWER core is doing more work than the x86 cores – and software is usually priced per core). Note that better utilization also lowers data center costs (not as many machines running and taking up space). And, finally, fewer machines mean fewer access points (important from a security perspective).

Cache Memory

With respect to cache memory, Intel-based x86 servers use static random access memory (SRAM), for cache and dynamic random access memory (DRAM) for main memory. Power Systems use for eDRAM (embedded DRAM) for cache and DRAM for main memory. The important difference between these two types of RAM is that DRAM and eDRAM are denser ─ requiring only two transistors to implement a bit, as compared to four or six transistors in SRAM. This denseness enables IBM to cram a lot more L3 eDRAM cache on POWER chips as compared with Intel x86 designs. (Note that SRAM is faster than DRAM, but having more L3 memory to work with on-chip can result in big increases in overall performance). As a result of denser memory, IBM can put 80MB of L3 cache on a POWER7+ chip as compared with typical Intel Xeon (the E7, for instance) at 30MB (IBM has an almost 3-to-1 advantage).

LPAR-to-LPAR Bandwidth

With respect to LPAR to LPAR bandwidth, communications amongst partitions within a single Power System operates at 1.21 gigahertz (GHz) – much faster than trying to communicate amongst x86 servers on a network (delays from slower network speeds, firewalls, and additional layers in the software stack of each server can be avoided when communicating partition to partition).

A Closer Look at IFL Technology and Packaging

IFLs are sold as bundled "virtual solution stacks" ─ a single IFL solution stack consists of 4 processor core activations, 32GB of memory activations, 4 processor core licenses of PowerVM for Linux (the Enterprise Edition is required for IFLs) and reduced software and hardware maintenance. The "virtual" element of these solution stacks refers to the way an IFL is deployed (as a software solution used to deploy Linux virtual machines). IFLs are implemented using IBM's Capacity Upgrade on Demand (CUoD) technology ─ a technology that exploits additional processors that are packaged with certain models of Power Systems that can be activated as needed to deliver additional computing power to Power System customers (who pay only for the additional computing power that they need).

Finally, once IFLs are deployed, IBM customers also need to license their processor activations via the Linux distribution of their choice (Red Hat and SUSE offer special pricing options for IFL environments). Also note that Power Systems IFLs are available only on Power 770, 780 and 795s (these are the most scalable models of IBM's Power Systems line).

Pricing

As is the case with IBM mainframe IFLs, Power System IFLs are priced substantially less than the standard general purpose POWER cores. IBM's motive for doing this is twofold: 1) IBM wants to make it extremely enticing for prospects and customers to choose Power System IFLs instead of x86-based servers; and, 2) IBM wants to capture new Linux workloads on its Power Systems.

To demonstrate how aggressive IBM is about Power Systems IFL pricing, consider what the cost for PowerVM, POWER processor core activations, memory activation at full price and Linux software would cost. The US list price per IFL as part of the above mentioned virtual solutions stack is $8,591 (with out the server hardware and PowerVM/Linux OS/hardware support) ─ as compared with a regular list price for all of these elements at $49,568. *This means that IBM Power*

Systems IFL customers can save over $40,000 by taking advantage of IBM's special bundled solutions!

> ***This significant drop in price due to bundling is indicative of how much IBM wants to capture new Linux workloads on Power Systems. The mainframe takes a similar approach with low cost IFL hardware and software — but also creates "solution editions" that consist of discounted database, infrastructure, and application solutions all bundled, tested and made available for a fraction of their list price. We expect that, over time, the Power Systems marketing organization will create similar solution editions around Linux as they have already done around AIX and IBM i.***

As for how IBM Power Systems IFLs compare in price to x86 architecture, IBM offered the following pricing scenario (Figure 2) that compares the cost to build-out a Linux environment on an existing Power System as compared to deploying an equivalent Linux solution on an HP DL560. This scenario assumes that an IBM customer has a Power System 770 already running — and that there are 32 additional POWER cores and 256GB of memory that can be activated. It then compares the cost of activating the same 32 core/256GB memory environment on an HP DL560 server (running slower cores, by the way), to the cost to run an equivalent environment on the IBM Power System 770. After looking at street and list prices, this pricing model shows that it costs about 4% less to implement an equivalent Linux environment on a Power System with IFL facilities as compared to a native x86-based server.

Figure 2 — An HP x86-based Linux Server Compared to a Power Systems IFL Environment

Client currently owns a Power 770 with 64 cores & 1 TB of memory installed, 32 cores active + 512GB active
Add 4 x Power IFL features activating another 16 cores + 128GB memory

HP DL560 2.7GHz
32-cores, 256GB

Hardware	**$30,584**
OS	$19,494
VMware Lic & SWMA	$21,731
HWMA 24x7 & SWMA	$1,811
TCA List	$73,620
TCA Street @ 25% disc	**$55,215**
SPECInt Rate Performance .	**1,240**
Price/Perf List	$59.37
Price/Perf Net	**$44**

8 Power IFL on Power 770
Activate 32 cores, 256GB on 64c 3.8GHz Power 770 system*

Hardware	**$68,727**
OS (Sub/Support provided by TSS)	$19,629
PowerVM SWMA	$8,720
HWMA 24x7	$25,455
TCA List	$120,420
TCA Street @ 35% disc	**$78,273**
SPECInt Rate estimate	**1,838**
Price/Perf List	$65.52
Price/Perf Net	**$43**

Prices are for concept illustration only and are subject to change. Example prices shows 3 yr TCA of purchasing 4 x Power IFL @ $8,591 each, HWMA, PowerVM & Linux Software subscription & support (24x7) on an existing 64-core, 512GB Power 770, based up on scaling estimates of SPECint_rate workload, including Power 780 64-core result of 3,730 and other results published Oct 2012.

Source: IBM Corporation — January, 2013

We like this comparison because it shows that Power Systems IFL pricing is very close to x86 pricing – so it helps to remove the price discussion from the conversation so IT executives can instead focus on platform efficiency, bandwidth, RAS – and all of the other advantages that Power Systems deliver as compared with x86-based server environments. On the other hand, we don't like this model for two reasons: 1) it doesn't show how much more efficient the IFLs on Power Systems operate (and it is reasonable to assume that they can process information twice as fast as their Intel competitors because the POWER processor can execute twice as many threads per clock cycle); and, 2) we would rather see an apples-to-apples Power Systems/x86 server comparison that shows what each server costs running the same workload.

The New Power Systems Linux Workloads

IBM believes that Power Systems IFL's are ideal for addressing four use cases: 1) private cloud/consolidation environments; 2) collocation environments; 3) environments that need high QoS for mission critical business applications; and, 4) environments that can benefit from simplified operations by aligning business process flows with shared services using underlying system "policy groups" (see Figure 3).

Figure 3 — Target Markets for Power Systems with IFLs

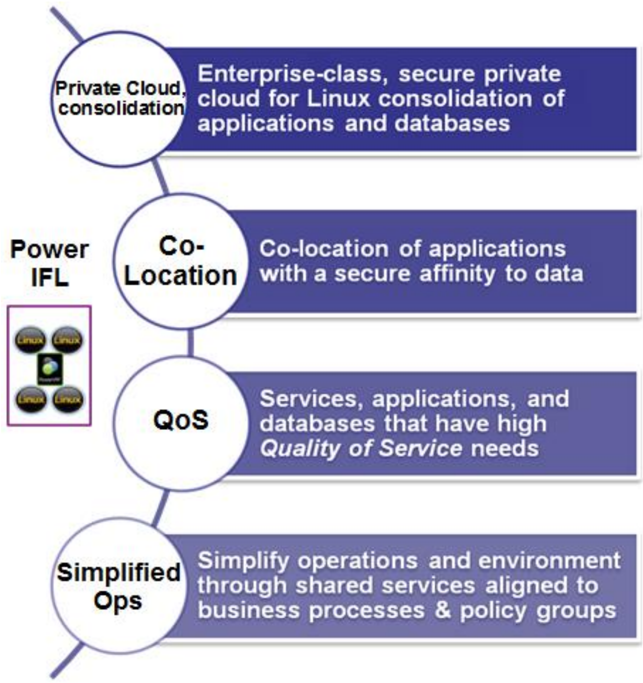

Source: IBM Corporation — January, 2014

At *Clabby Analytics*, we see the target markets for Power Systems IFLs slightly differently. In short, we are expecting the adoption of Power System IFLs to mirror mainframe IFL adoption scenarios. And this means that we expect to see Power Systems IFLs deployed:

- ***On unused processors in existing Power Systems that have spare capacity*** (as shown in Figure 2). The primary benefits of using IFLs on these systems is that IFLs will help increase overall system utilization rates — enabling the enterprise to attain a better return on its system investment while also eliminating the need to purchase additional Linux servers.
- ***For Linux consolidation purposes***. In this scenario, an enterprise can take advantage of high-end Power Systems scale-up architecture to run hundreds to thousands of Linux instances within the same server (thus reducing distributed x86 system sprawl). Because Power Systems offer greater utilization rates than x86 alternatives — and because all of these Linux instances can be managed using common management tools — enterprises will

not need as many x86-based servers to execute Linux workloads, nor will they need as many distributed Linux server managers to administrate those servers. Further, Unix (AIX), IBM i, and Linux workloads can be easily consolidated and managed on the same server — leading to simplified operations.

- ***To reduce extract/transform/load (ETL) costs by collocating applications and data*** on the same server. One of most insidious hidden expenses in operating a data center is the cost of moving data. In order to move data to "less expensive" systems for processing, IT managers often purchase additional servers, storage and networking equipment in order to transfer files and process data. Then these managers spend hundreds of labor hours managing file transfers and associated data. Further, many machine cycles are wasted by having to extract data, transform it, and load it onto other servers. By collocating applications and data onto the same scale-up Power System, ETL costs can be eliminated (potentially saving enterprises millions of dollars annually — see this report for further details).
- ***To build a self-contained cloud-in-a-box architecture***. *Clabby Analytics* has written several reports on enterprises that recognize the merits of building and managing a secure cloud environment contained in a scale-up server (see here and here for examples). We expect Power Systems users to also recognize these same advantages when deploying Linux based cloud solutions (including Infrastructure-as-a-Service and Platform-as-a-Service environments) on highly secure and more easily managed Power Systems architecture.
- ***To address requirements for advanced RAS/QoS system designs for mission critical Linux environments***. While a case can be made that RAS (reliability, availability, and serviceability) has improved significantly over the past few years in x86 server design (thanks to the advent of "expert integrated systems" designs); there remain several distinct differences in the Quality-of-Service (QoS) characteristics of Power Systems versus x86 architectures — particularly in virtualization and security. From a virtualization perspective, Power Systems enable more robust virtual machines to be created as compared with Intel-based servers (thanks to more powerful computing facilities combined with access to more memory). Further, Power Systems offer advanced virtualization management facilities that allow system resources (CPUs, memory and storage) to be organized easily into shared pools for dynamic reallocation between multiple workloads. And Power Systems offer advanced mobility features (the ability to move workloads in progress to other servers). From a security perspective, as described earlier, consider that fewer than 10 vulnerabilities are reported yearly on Power Systems — and contrast this with the hundreds of vulnerabilities reported annually on x86 Windows- and Linux-based servers.
- ***As open server environments for new generation open/mobile/social applications***. From our perspective Power Systems can also be configured as centrally-managed host environments for new mobile, social and open source applications that have come and are coming to market.

Power IFL Services

Enterprises that want to maximize their investment in Power Systems IFLs might want to consider using IBM's PowerCare services to help deploy Power Systems IFLs. Tuning services for IFL environments include high availability, hosted infrastructure (cloud) services, performance tuning services and virtualization services.

-

- With respect to high availability, an IBM consultant can advise IFL managers on high availability design options — as well as help configure high availability services on Power IFLs;
- Linux on Power Hosted Infrastructure Services provides best practice advice for migrating and optimizing Linux hosted infrastructure services on Power IFLs;
- Linux on Power Performance assesses server and software configuration on Power IFLs, checking for overall system health and configuring for optimal performance; and,
- The Linux on Power – PowerVM service optimizes PowerVM virtualization features for Linux on Power IFLs, ensuring optimal dynamic resource sharing and utilization. Further, IBM's PowerVP (virtualization performance) provides performance monitoring that helps proactively address performance issues by mapping virtual workloads to physical hardware.

Summary Observations

IBM recently announced that it plans to invest $1B over the next several years in Linux solution development on Power Systems servers, as well as on expanding the OpenPOWER Consortium (an organization dedicated broadening the use of POWER processors around the world as well as expanding the POWER ecosystem). The company strongly believes in the future of Linux on Power Systems — and is putting its investment money behind that belief. IFL technology is but the tip of the iceberg in a rapidly developing Linux ecosystem on Power Systems.

The big challenge that IBM will face as it continues to drive its Linux on POWER initiative will be changing the ingrained buying habit of the Linux community that purchases x86-based servers as a matter of course without doing due diligence on whether Linux workloads can run more efficiently on other hardware platforms. With mainframe and Power Systems IFLs, IBM offers faster processors and significantly more memory for applications and databases — as well as richer QoS. For the right workloads (workloads that need scalability, that need faster processing and access to more memory and cache, and that need richer QoS), choosing to run Linux on a Power System should be a no brainer. But educating the Linux community on how to use the right processor and system design to best execute given workloads is going to be tough (some IT buyers are just not willing to drop their old buying habits).

Still, there is a segment of the IT buying community that understands just how important deploying applications on servers best suited to execute those applications actually is. *Clabby Analytics* has written about several of these exemplary organizations (including case studies on Swiss Re, on IBM, and on others). Enterprises interested in saving loads of money in systems acquisition cost, software costs, and management costs should be very interested in what IFLs can do to reduce their IT costs.

> ***The bottom line in this report is this – enterprises looking to lower costs through new levels of Linux consolidation; looking for higher QoS levels than provided by x86-based Linux and Windows environments; looking to reduce ETL costs through collocation and integration of applications/data/process flow; and looking for a centralized scalable platform to host next generation open/mobile/social applications – need to take a very close look at what IFL technology can do to improve services and lower costs.***

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

January, 2014