

## **POWER9 – should I wait?**

**By Jaqui Lynch**

### ***Introduction***

At the recent Hot Chips 2016 conference (August 2016) IBM released some details on the 2017/2018 POWER9 chips. The new chips are going to be even more powerful than POWER8 and are architected to optimize both scale out and scale up technology. This leaves many who are on POWER7 or earlier with a conundrum – do I upgrade to POWER8 now or do I wait for POWER9 and skip a generation or 2? This decision is key to current budgeting but it is impacted by many other factors.

### ***What do we know about POWER9?***

Since the first release date is not till second half of 2017, many things could change from what was discussed at Hot Chips. For now, the design for POWER9 consists of 24 core chips that are initially optimized for 2 socket scale out servers. The chips are being reduced from 22nm (POWER8) to 14nm and, potentially even smaller in 2018. The servers will also use PCIe Gen 4 slots for I/O connections and the CAPI2.0 interface for coherent accelerator and storage attach along with a new 25Gbps advanced accelerator attach bus. This is very cool but what does it really mean to the business?

First of all, the reduction in chip size means smaller, faster servers that can handle a great deal more work. No information on power and cooling requirements is available yet so we will have to see what that looks like closer to the time. Performance and parallelism will be significantly increased – this means you can pack more virtual machines onto the same server and service more users. And the use of accelerators such as GPUs and FPGAs provide even more performance enhancements. All of this means you can have fewer, smaller, high performing servers in your data center that are capable of providing a reliable and scalable environment for your workloads. So do you wait or upgrade now?

### ***Factors in determining upgrade decisions***

One key factor in determining whether to upgrade now or whether to wait is clearly business requirements. If the business needs to upgrade now, due to performance or technology or support needs, then upgrading to or implementing POWER8 today is the only option. POWER8 is a great technology offering reliability, scalability and performance. If you need to upgrade from your current technology or you need additional resources within the next 12 to 18 months, then I would upgrade to POWER8.

You may also want to consider upgrading to POWER8 as a staging position until POWER9 is available. Given that AIX 7.2 does not run on POWER6 and the new HMC v8 does not support POWER5, this makes POWER8 a very attractive platform to get to the new technologies such as Flash cache that come in AIX 7.2, while also preparing your systems for a POWER9 update later.

If you do upgrade to POWER8 now, then I would also ask about upgrade protection to see if there are any offerings available for POWER9 upgrades when POWER9 is formally announced.

A second factor is how risk averse your company is. Many companies use technology for competitive advantage and want to be the first with the new technology. Other companies like to be about 6 months behind when a new technology comes out so that it has time to be field tested. And the last group of companies only upgrades when they have to and they may stay one technology behind. Obviously, it is important to know where you stand on this issue. If you are one of those who likes to be first, then it is worth asking about beta programs to see if you can get early access to the technology. Or you can wait till IBM announces the product with the final specifications and be one of the first to implement it. If you are very risk averse, then this is not the path for you.

Other factors affecting upgrades will also include vendor support. Right now your vendor will most likely support POWER8 as a platform. If this is a mission critical application, then you would want to confirm with them when they will support POWER9, and if moving to POWER9 will change any of the licensing pricing or strategies that they use. IBM prides itself on their binary compatibility between the POWER servers but you still need to confirm when the vendors will support POWER9. This may mean going to POWER8 rather than POWER9 until POWER9 has been out for a few months.

### ***POWER9 USES***

Two of the growing markets for data centers are the high-performance computing and the hyperscale markets. IBM is targeting huge growth in these two markets and POWER9 is designed to assist in this. For high performance computing, cloud, analytics and other similar applications parallelism and performance are critical. The new POWER9 processors will be highly parallel, supporting up to eight threads per core and each POWER9 processor is expected to come with up to 24 cores, making the chip well suited for applications that benefit from being run in parallel. Additionally, the POWER9 chip has been designed to have a range of connectors to attach FPGAs (field programmable gate arrays), GPUs and ASICs (application specific integrated circuits). This allows the use of co-processors that are designed for very specific tasks related to databases, artificial intelligence, business intelligence, data mining and other similar applications. Memory will be DDR4 DIMMs. POWER9 will also be the first processor to provide PCI-E 4.0 which can provide for 16Gbps (2xPCI-E3). It will also use the newer NVLink 2.0 which provides a bandwidth of up to 25Gbps for the NVidia GPUs. NVLink allows CPUs to communicate with other components at speeds 5-12x faster than PCI-E gen3 x16. The chip also includes on-chip acceleration for compression and encryption.

IBM has also committed, through the OpenPower foundation, to license the blueprints to the CPU's architecture, server hardware and software out to other

vendors. This opens the door for what IBM is calling “Partner chips” to be produced that would allow other vendors to build POWER8 and POWER9 servers. Already, Google and Rackspace are working on POWER9 computer blueprints for a system called Zaius, and Google appears to be endorsing POWER9 as well.

### ***Summary***

POWER9 is going to be a phenomenal technology and the servers are going to provide great performance for mission critical workloads like cloud, analytics and other highly scalable applications. Given that POWER9 is at least 12-18 months away, depending on the server you need, I would not wait for POWER9 at this time. Instead I would recommend putting in place a strategy that brings your current systems up to the current POWER8 technology along with the latest AIX, so that you are well positioned to take advantage of POWER9 when the server specifications become available. If you are running anything prior to POWER8 today, then now is the time to design a strategy to get current. There are significant savings in power, cooling, and maintenance by getting to the current architecture. And, for those of you concerned about AIX v5.3, IBM recently provided a support option to run AIX v5.3 on POWER8. Personally, I recommend you get to AIX 7.1 at a minimum, but at least you now have an option other than WPARs if you are stuck at v5.3 for a while.

### ***References***

AIX v5.3 on POWER8 Support

<http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=897/ENUS216-155&infotype=AN&subtype=CA>