

A Closer Look at the PureFlex System

July 2012 | by [Jaqui Lynch](#)

On April 11, IBM announced the PureSystems platform. The first two members of this family are PureFlex System and PureApplication System. The former is an infrastructure as a service (IaaS) platform, combining server, storage and networking hardware with a management infrastructure that provides single pane of glass management for all of the system's compute nodes.

The PureFlex hardware itself consists of a 19-inch T42 rack with network hardware, a Storwize v7000 storage system and up to four chassis, each of which can hold up to 14 half-width compute nodes. These look similar to blades but are horizontal rather than vertical. So why not call it a blade system?

Some fundamental differences exist: For starters, a single management infrastructure allows for provisioning and control of all of the nodes within the rack—a huge step forward in flexibility and control. Additionally, everything is pre-cabled, integrated and tested within the PureFlex System before it's shipped. The customer only needs to provide power and final network connections, and to set up the final LPAR provisioning. This level of integration makes a huge difference in time to deployment. It's no longer necessary to coordinate with multiple groups to get storage-area network (SAN) ports, switches, cables, etc. It's all in the rack.

Additionally, unlike traditional blades, the networking north-south problem has been resolved. This refers to communication between two blades, which must go to the rack switch and to the receiving blade—leading to network latency that can affect performance in latency-sensitive applications. PureFlex System compute node-to-compute node communications can occur within the chassis (east-west rather than north-south) if attention is paid to the placement of nodes that must communicate with each other, thus significantly reducing latency.

The Technology

PureFlex technology also allows compute nodes to be mixed within a chassis. So Web, application and database servers can all be within the same chassis, taking advantage of the latency reductions. Additionally, those servers could be whole nodes, LPARs within compute nodes and could reside on different OSs. Within a PureFlex chassis, you can run partitions with AIX, IBM i, Linux and Windows. Multiple hypervisors are supported with initial support for PowerVM, kernel-based virtual machine (KVM), VMware and HyperV. Three Power and two Intel nodes are available to choose from today, and the system has been designed to be upgradeable to extend its lifetime and thus reduce total cost of ownership (TCO).

PureFlex Systems come in three sizes—express, standard and enterprise. The minimum express system consists of the rack, chassis, 1 x 10Gb network switch, 1 x 8Gb SAN switch, the Flex Manager node, two power supplies, four fans, redundant chassis management modules (CMMs) and a redundant v7000 system with two SSDs and eight HDDs. Additionally, a customer must select at least one of the compute nodes. The standard edition doubles the fiber switches, adds eight more HDDs, two power supplies, and two fans. Finally, the enterprise edition, compared to standard edition, adds another 10Gb Ethernet switch, two more power supplies and two fans.

The chassis itself takes 10U in the rack and holds 14 half-wide nodes (p260, p24L, x220, x240) or seven full wide nodes (p460) or any combination thereof, including the half-width Flex Management node. Three

options for the Power compute nodes are available—the p260, p24L (Linux only) half-width and the p460 full-width nodes. The p260 / p24L are is a dual-socket nodes that supports 2 x 4 core 3.3ghz sockets or 2 x 8 core 3.2 or 3.55ghz sockets, and up to 256GB of memory. There's room for two HDDs or two SSDs, however, the use of HDDs limits the memory to 128GB because they are larger than SSDs and encroach on the space for memory. There are two mezzanine cards, one for networking and one for fiber—this node only supports one virtual I/O (VIO) server. The p460 is a four-socket node that supports 4 x 4 core 3.3ghz sockets or 4 x 8 core 3.2 or 3.55ghz sockets and up to 512GB of memory. There's room for two HDDs or two SSDs, but again the use of HDDs limits the memory to 256GB. With four mezzanine cards, it's possible to support two VIO servers, although one will need to boot from SAN as the disks cannot be split.

The Intel half-width, two-socket nodes are the x220 and the x240. The x220 offers two, four, six or eight cores per socket, supporting up to 192GB memory. The x240 offers four, six or eight cores per socket and supports up to 768GB memory. The Intel nodes also support 10 Gb Ethernet, 16 Gb FC, and Fourteen Data Rate InfiniBand expansion cards.

Key Drivers

So what's it all mean? Does it really work? Are blades and servers now defunct? The short answer is the PureFlex Systems are clearly aimed at environments in the two to 32-core range, especially where both Intel and Power nodes need to communicate with minimal latency. Additionally, simplified management and pre-delivery integration are key drivers as they reduce the complexity of setup. It's much easier to set up a 10-node p260 (each 16 cores) solution than to set up a similar p740 complex (including rack, stack, cabling, etc.). Individual servers still have a role, especially the p770 and higher, depending on whether horizontal or vertical scaling is the target.

With respect to whether it really works, I had access to a system just before announcement, and, even though it was not the production version of the Flex Manager, I was still able to test many of the capabilities and navigate around the system without needing a training session. It was far more intuitive than Systems Director and performed very well. Is it perfect? No, but IBM has done a great job publicizing the elements that will come later. I'm waiting on my lab system now so I can work on more testing, especially as it relates to latency. In the meantime, it's worth exploring the PureFlex System as an option in your data center.

IBM Systems Magazine is a trademark of International Business Machines Corporation. The editorial content of IBM Systems Magazine is placed on this website by MSP TechMedia under license from International Business Machines Corporation.

©2019 MSP Communications, Inc. All rights reserved.