



IBM Power Systems

By Jean S. Bozman

When IBM rolled out its line of Power7 systems in 2008, it might have been difficult to imagine what Power systems would become in 2022 – or to know whether there would be a Power10 processor at all.

On Sept. 8, 2021, those questions were answered. The 7nm IBM Power10 processor was described at the Hot Chips conference, highlighting its on-chip AI inference engines. And, with the Sept. 8 news, it's clear that the Power systems' capabilities now include on-chip AI engines for data analytics, reduced energy requirements, and enhanced support for SAP HANA, clustering, workload isolation, containers, cloud computing and end-to-end security.

If the capabilities of Power systems were not fully appreciated in the late '00s, it was likely because x86 servers were so widely installed in the customer base. Any yet, Power systems found ways to stand out, running scalable datacenter, web, ecommerce and the first wave of emerging cloud workloads. In the 2008/2009 financial downturn, Power-based systems supported the first big wave of enterprise cloud migration – and they are positioned to take on the 2020-2022 wave of cloud migrations for customers' hybrid clouds.

The New Power E1080 Systems

Now, IBM is unveiling and rolling out the first of a [new generation of Power systems](#) based on the Power10 processor. The first model, the IBM Power E1080, is being introduced as the fall season begins – a time of year when IBM often makes its strongest IT systems announcements.

Four key categories topped the company's list of design points: flexibility, security, business value and reliability. These characteristics were identified by IBM's survey of its own customer base worldwide. But they are not the only features, which have a much longer list.

AI/ML processing was a key factor in the IBM Power E1080 design. A new embedded feature for AI and machine learning can support five times the amount of AI inference processing, compared to the earlier IBM Power E980 server. It accomplishes this due to packaging four AI engines per core, embedded in the Power10 chip itself. That is significant, given that AI and analytics drives a high percentage of annual IBM Power systems shipments.

Consolidation capability is another aspect of Power10 systems. Customers will be able to consolidate more workloads per footprint in the new Power systems, reducing their energy



costs and datacenter “real estate” requirements by installing the new systems. Those include enterprise applications and data managed by a variety of operating systems (IBM AIX, Linux and IBM i), SAP applications, and a range of relational databases residing at customer sites.

The new Power E1080 systems include:

- Embedding four Matrix Math Accelerator (MMA) engines per Power 10 core
- Reducing energy consumption (33%), compared to Power9 systems
- Supporting SAPS benchmarks (8-socket system with 120 cores) up to 955,000 SAPs
- Supporting up to 16 sockets for scaling-in-place deployments
- Adding 50% more capacity than Power9 systems
- Providing 2.5 times faster encryption than Power9 systems
- Supporting end-to-end, key-enabled security for distributed and clustered systems

As announced, the IBM E1080 will be available for on-premises installation, and accessible from the IBM Cloud public cloud service. Customers can subscribe to a pay-per-use consumption model in on-premises data centers with by-the-minute metering of capacity usage through Power Private Cloud with Dynamic Capacity. Companies can add Power E1080 and Power E980 into a private cloud pool, allowing them to share pay-per-use capacity in the pool.

Virtual Server Deployment Options

Consumption models include more Power Virtual Server options for the IBM Cloud. IBM refers to this form of deployment as Power Virtual Server co-located with IBM Cloud. Customers who prefer to access Power systems as cloud resources can choose among 14 data centers worldwide, spanning North America, EMEA and Asia/Pacific/Japan. This will support pay-per-use of Power Virtual Server LPARs (logical partitions) in geographic areas that are closest to customers’ physical locations, for faster performance and reduced latency.

Customers running their workloads on Power E1080 installed on-premises can extend their workloads in Power Virtual Server without requiring any additional application or middleware refactoring.

Key Takeaways

So, given the four top design principles for Power10 systems, how did IBM do? It responded to those four key characteristics it identified: flexibility/agility; increased security, gaining value from AI/ML analysis of business data; and high availability/resiliency.



Now that the Power10 generation of systems are here, IBM could build on this base by marketing with an emphasis on providing more “paths” to Power systems use-cases. This could take many forms: Building on existing partnerships with co-location services, installing more Power E1080s for use in the IBM Cloud (Virtual Servers); selling more Power E1080s for behind-the-scenes, cloud-enabled AI/ML capabilities; support for end-to-end corporate-wide security; and providing as-a-service models for IBM Power E1080s deployed in customers’ private clouds.

These options are all viable, and they should all be presented as use-cases. In a cloud-centric world, customers do not have to follow any given path. That is clear for any customer’s hybrid cloud strategy – and it is a pragmatic approach to marketing in the New Normal economy.

Customers are shopping for systems capabilities, not for specifications, in the hybrid cloud world. It’s all about applications, analytics performance, security, capabilities, and consumption modes. Rather, customers should consider multiple paths for leveraging the Power E1080 systems’ capabilities. That would give Power10 systems a higher profile in IBM’s largest accounts. Importantly, it would make Power10 systems more “visible” in a wider range of customer accounts -- including longtime IBM customers and prospective customers worldwide.