

Introducing the POWER9 Scale-up Servers: E950 and E980

August 2018 | by [Jaqui Lynch](#)

On August 7, 2018, IBM announced the scale-up POWER9 servers: the 9040-MR9 E950 and the 9080-M95 E980. These servers are designed to fit into a 19-inch rack. The E950 will GA on August 17, the 1-2 E980 node will be GA September 21 and the 1-4 node E980 will be GA November 16.

Scale-up POWER9

The POWER9 processor used in the scale up servers is almost the same in the E950 and the E980. The processor supports core multi-threading up to SMT8 and is designed as a single chip module (SCM). Using an SCM provides many benefits including reduced path lengths due to a single hop for the processor fabric interconnect. Additionally, there are benefits where a vendor does chip-based licensing versus individual core or socket based licensing. Each processor is populated with either 8, 10, 11 or 12 cores per socket with varying speeds depending on the cores and the model. The E950 will be a small software tier and the E980 will be a medium software tier.

E950 Technical Specifications

The E950 server is a 4 socket 4U server and, initially, it comes with either two or four processors in 8, 10, 11 or 12 cores per socket configurations. A fully populated E950 can support up to 16TB of DDR4 memory with a memory bandwidth of 230GB per second per socket. These servers have PCIe Gen 4 I/O directly integrated into the processors which reduces latency and increases I/O bandwidth.

Processor options are as follows:

48 cores	Up to 4 sockets of 12 cores	3.15 to 3.8 GHz
44 cores	Up to 4 sockets of 11 cores	3.2 to 3.8 GHz
40 cores	Up to 4 sockets of 10 cores	3.4 to 3.8 GHz
32 cores	Up to 4 sockets of 8 cores	3.6 to 3.8 GHz

In the above, the 44 and 48 core models are designed for maximum throughput and the 32 core model is designed for maximum core performance. A minimum of one socket full of cores must be active and at least half the memory must be active (minimum is 128GB).

Each E950 has 11 PCIe slots. There are 10 PCIe Gen 4 slots for adapter cards – 2 x16 per socket, plus an additional 2 x8. The 11th slot is a pcie3 that is reserved for the default 2 port 10Gb network adapter. The numbering for the slots starts with C10-C12 on the first socket, then C6-C9 on the second socket, C4 and C5 on the third socket and C2-C3 on the fourth socket. The x8 slots are C9 and C12. Slot C6 is the PCIe3 network card slot.

There are also four 25Gb/s accelerator ports for future use. As with previous servers, there are three possible backplanes: a single backplane, a split backplane and a dual raid, write cache backplane. For

those running VIO servers, the split backplane is the typical choice. The backplane supports up to 8 x SFF SAS drives plus the 4 x NVME. With the split backplane option the 8 x SAS drives get split into 2 sets of 4 drives for allocation to two LPARs. So, it would be possible to support multiple VIO servers (up to six) with four booting from the NVME and two from the SAS drives. There's also an external powered DVD option (feature EUA5) that I recommend purchasing with the first server.

The E950 comes with a default warranty of 9-5 for one year, but initially IBM is providing a free uplift to 3 years of 24-7 support.

E980 Technical Specifications

The E980 will initially support up to two nodes, although additional nodes (up to four total) will be available mid November. Each node is 5U in the rack and there is also an SCU (control unit) that requires 2U. If you plan to expand to four nodes in the future, then you need to plan for 22U in one block, so it is important to reserve that space for future use. The E980 will need to be installed towards the bottom of the rack due to weight restrictions. You'll also need to reserve space for horizontal PDUs.

Processors are available in 8, 10, 11 or 12 cores per socket configurations and each node has four sockets. A fully populated two node E980 has eight processor sockets and can support up to 32TB of DDR4 memory with a memory bandwidth of 230GB per second per socket or 920GB per second per node. The four node E980 will support up to 64TB of memory. These servers have PCIe Gen 4 I/O directly integrated into the processors which reduces latency and increases I/O bandwidth.

Processor options per node are as follows:

48 cores	Up to 4 sockets of 12 cores	3.55 to 3.9 GHz
44 cores	Up to 4 sockets of 11 cores	3.58 to 3.9 GHz
40 cores	Up to 4 sockets of 10 cores	3.7 to 3.9 GHz
32 cores	Up to 4 sockets of 8 cores	3.9 to 4.0 GHz

A minimum of eight cores and 256GB of memory must be active.

Each E980 has eight PCIe4 slots and they are all x16. The numbering for the slots starts with C1-C2, then C3-C4 plus p2-c1 and c2 (NVME), then C5-C6 and P2-C3-C4 (NVME), and C7-C8 as well as C13 which is the USB slot. Each NVME drive is individually addressable, which means that each NVME drive can be allocated to a different LPAR if so desired. Initially the E980 will support 2 x EMX0 I/O drawers per node and in fourth quarter of 2018 this will go to four per node.

The E980 comes with a default warranty of 24-7 support for one year and comes with redundant clocks and redundant flexible service processors (FSPs).

Disks and I/O Drawers

Both the E950 and the E980 will support the current EMX0 PCIe I/O drawers with the EMXF fanouts and the EXP12SX and EXP24SX disk drawers. For the E950 support for the 5887 will be added later, but it will be available at GA for the E980. When connecting the E950 to the I/O drawers the cards to the drawers must be in the x16 slots. All of the slots in the E980 are x16 so this is not an issue there. As a part of the announcement, IBM will publish a list of initially supported I/O adapters as well as a list of those that will be supported later.

If you're ordering I/O drawers it is important to understand that only 4K disks are supported in the EXP12SX and EXP24SX drawers. I highly recommend ordering 4k disks for the internal disks as you cannot combine 512 byte and 4k disks in the same volume group, which can lead to problems in the future if you want to move volume groups around. You should also check whether there are any issues using 4k disks with the software you plan to run.

Management

PowerVM Enterprise Edition is now included with every POWER9-equipped box. It's a zero-cost item, although you will still need to purchase SWMA to get software support. Additionally, if you don't have PowerVM Enterprise Edition on your source POWER7 or POWER8 server, you can request a free 60-day activation for that server to facilitate migration.

An HMC or vHMC (virtual HMC) is required to manage POWER9 servers in a PowerVM environment. The HMC or vHMC is also required for concurrent firmware maintenance. For POWER9, the HMC must be a 7042 CR7 or higher, a 7063-CR1 or later, or a vHMC (x86 or POWER based). The vHMC requires four cores, at least 8GB memory and a minimum of 500GB of disk space. This means it's highly likely you'll be buying a new HMC. The new HMC at v9 can be used to manage POWER7 and higher servers. The new HMC will need to run HMC v9.2 to support the E950 and E980—this level doesn't support classic mode for the HMC.

NVME

In addition to regular SFF SAS disks, each E950 can include up to four NVME SSD drives and each E980 node can also have up to 4 NVME SSD drives. These drives are individually addressable, so the 4 drives can be assigned to 4 different LPARs. SSDs have a drive write per day rating—this is because they are really designed for read intensive workloads as high write workloads wear them out. On the E950 this rating is 2.4 and on the E980 it is 2.2. They should be used for VIO server or O/S boot drives, not for active data.

Gotchas

None of these servers have an internal DVD drive so you need to attach an external USB drive. Additionally, the POWER HMC doesn't have a DVD drive either. You can purchase an external DVD from IBM or you can try and find a USB DVD that will work with the HMC and server, but it must use external power. Additionally, the HMC and server can use USB sticks. So far, I haven't found them supporting anything larger than 16GB as they have to be formatted as FAT32. This is a really fast way to make firmware and HMC upgrades available, as well as to back up the HMC.

All of the POWER9 servers ship with the security patches for Meltdown and Spectre installed and enabled.

Power requirements are a little higher, so I highly recommend moving to the #EPTN 30 amp high function PDUs. These include 12 x C13 receptacles with a 20amp circuit breaker on each receptacle. The 9188/7188 PDU also has 12 x C13 receptacles rated at 10 amps, but there is a 15amp circuit breaker per pair of receptacles. These PDUs can be either 60 amp or 30 amp PDUs depending on the line cord attached.

You should also ensure you review your current power environment to ensure you have sufficient power available to support the new servers. Each E950 and each E980 node has four power supplies: two primary and two redundant.

Power to Cloud Rewards

IBM is also offering Power to Cloud rewards with each server purchased. Clients earn 5,000 points for each E950 and 10,000 points for each E980. Those points can be redeemed towards IBM Lab Services engagements starting with 10,000 points.

Prerequisite Software

The E950 and E980 require specific minimum releases of the supported operating systems in order to

run. Both servers support AIX, Linux and VIOS (PowerVM) and the E980 also supports IBM i. In particular, VIOS needs to be upgraded to 2.2.6.23 or later if you plan to use VIOS on the E950 and 2.2.6.31 for the E980. PowerVM 2.2.6.31 will be available 3Q and 2.2.6.32 and 3.1.0 will be available 4Q. There will also be new technology levels and service packs for AIX and other products.

LPM also requires specific minimum levels at the source in order to work. You can only migrate workloads from POWER7 or higher to a POWER9 box using LPM. Any POWER6 workloads will need to be migrated to a POWER7 or POWER8 and rebooted prior to using LPM to migrate to the POWER9.

Upgrades

IBM will be offering same serial number MES (or PEX) upgrades to the E980 from the E870, E870C, E880 or E880C servers. There's no same serial upgrade being offered to the E950. These upgrades will be GA November 16. When trying to determine what to upgrade to a good rule of thumb is that the p770, e850 and even a small 780 could be moved to the E950. Larger p780s, the e870/c and e880/c should move to the E980. Any p750s could move to either the S924 or the E950. This, of course, depends on what you want to consolidate and how many servers you want to consolidate from and to.

Summary

The POWER9 scale up server announcement is the third in the POWER9 series. The servers show incredible gains in performance and there are several new software releases that need to be considered. Prior to ordering any of these high-end servers, I would recommend requesting a TDA (technical delivery assessment) to ensure no power, HMC or software requirements are missed.

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.