




Take a Closer Look at the New IC922 Server

Power Systems expert Jaqui Lynch gives an in-depth look at the technical specifications for the latest POWER9 box.



By Jaqui Lynch

03/02/2020

On January 28, 2020 IBM announced the latest in their POWER9 series of servers, the IC922 (9183-22X), where the IC stands for “inference and cloud” and the I can also stand for I/O. It’s a 2U server with up to 20 cores and up to 2TB of memory that is architected to provide the best performance for AI inference models. It integrates with the AC922 training system to provide a high scalable modular insight platform.

Technical Specifications

The IC922 is a two-socket server with both sockets fully populated with identical SCM (single chip module) processors. Processors are either 12, 16 or 20 cores, which means the server will have a total of 24, 32 or 40 active cores. Each core provides for up to four threads (SMT4). The Processor SMP interconnect supports two inter-node SMP X-bus links with a peak inter-node performance of 120 GBps. The inter-node SMP X-bus connection has twice the performance of the LC922 server, as it has two X-busses whereas the LC922 only has one.

The IC922 also supports NVIDIA T4 GPUs to boost inference performance. There are two options for the GPUs:

#EK4L 0-4 NVIDIA T4 GPU with PCIe Gen 3 x16 LP (16 GB)

#EK4M 0-2 NVIDIA T4 GPU with PCIe Gen 3 x16 (16 GB)

The NVIDIA T4 PCIe GPU is a single slot 6.6-inch PCIe gen 3 adapter that can be used in an x8 or x16 slot in the system. The FC EK4L is a short, low-profile adapter. The FC EK4M is a short full height adapter.

There are a total of 32 DIMM slots, providing for up to 2TB of memory using 16GB (#EM62), 32GB (#EM63) or 64GB (#EM64) memory features. To attain maximum memory bandwidth at least 16 of the 32 DIMMs should be populated. Memory DIMMs are populated as follows: 4, 8, 12, 16, 24 or 32 memory DIMMs (up to 2TB).

The minimum configuration includes two processor modules, 64 GB of memory, two power supplies, two line cords plus some other features.

A minimum configuration would include items such as:

2 x #EK00	12c 2.8-3.8ghz OR
2 x #EK01	16c 3.35-4ghz OR
2 x #EK02	20c 2.9-3.8ghz
4 x #EM62	16GB dimms 2666mhz 4Gb ddr4
1 x #EC16	Open Power Abstraction Layer (OPAL
1 x #4650	Rack Indicator -- Not Factory Integrated
2 x #EKMP	AC Power Supply, 2000 Watt (220 V)
2 x power cords	Select two power cords from supported list. Feature 4558 is defaulted. (FC 4558 Power Cord To PDU/UPS, (100-240V/16A))
1 x #2147	Primary OS Linux
1 x #93xx	Language Group Specify

The IC922 has up to 3 backplanes (#EK36) to provide for up to 24 x 2.5" SAS/SATA disks or SSDs. At the time of announce NVME disks were not an option.

There are 10 PCIe slots available for adapters. Four low profile slots are mounted directly on the system board and there are an additional six slots mounted on two riser cards (#EK30). Slots 1 and 6 are directly connected to the processors and support PCIe Gen4 x16 (at 32GBps). Slots 4 and 9 are also directly connected but are PCIe gen 3 x16 and can support double wide adapters. The remaining six slots are connected over PCIe switches (PEX). Slots 2 and 7 are Gen3 x16 slots but are limited to Gen3 x8 (8GBps) speed. The connection from the PEX to the processor has a bandwidth of Gen 3, x8 and the remaining slots 3, 5, 8, and 10 have a maximum throughput of a Gen 3, x8 connection up to 8 GBps. This means that adapters requiring the best performance should be in slots 1 and 6 if possible. Additionally, slots 1,4,6, and 9 are CAPI 2.0 enabled. The PCIe adapter guide for the IC922 clearly lists the limits and placement recommendations for all of the potential adapters for the server.

The IC922 can be ordered with anywhere from zero to 24 disks or SSDs. Each backplane can support up to eight disks or SSDs. However, the IC922 does not have an integrated disk controller. The number of controller HBAs that you need to install will depend on the number of backplanes and disks/SSDs that are installed. Possible configurations for disk could be:

- Maximum of 8 x 2.5" SAS/SATA disks: one #EK41 or #EK47 HBA in PCIe slot 5 and one drive backplane (#EK36). The Broadcom MegaRAID 9361-8i HBA (#EK47) supports RAID levels 0, 1, 5, 6, 10, 50, and 60, and Just a Bunch of Disks (JBOD). The other HBAs only support JBOD.
- Up to 24 x 2.5" SAS/SATA disks - One #EK41 or #EK47 HBA in PCIe slot 5, one #EK42 in PCIe slot 10, and three drive backplanes (#EK36). This means that you should plan to lose slots 5 and 10 if you plan to have internal disks in the backplane. There's no support for I/O drawers so all internal disks and adapters need to fit into the maximum of 10 slots and 24 disks. Additional disks can be accessed from the SAN.

The rear of the IC922 has a VGA port, a serial port (RS232), a BMC (Baseboard Management Controller) port, two USB 3.0 ports, and two 1 Gbps RJ45 ports. The front of the server has a USB 3.0 port. They are all integrated into the system board.

The server has two hot-swap and redundant power supplies. They're 2000 Watts 220 Volt AC power supplies and have C20 receptacles, not C14. The PDU they are connecting to must support C19 connections (the default is usually C13). When planning for power you should plan for a maximum draw of 1.855kVa (about 9.3amps) on either of two PDUs with a maximum thermal output of 6143 BTU/hr.

BMC

If you have worked with any of the Linux-only POWER servers or with the new POWER HMC (Hardware management console) then you will have configured a BMC (baseboard management controller) connection. The BMC is used for system service management, monitoring, maintenance, and control, and is a specialized service processor that monitors the physical state of the system by using sensors. The BMC also provides access to the system event log files (SEL). Access to the BMC is through the network. Once it is setup you can access it via the web (OpenBMC https), or you can use the ipmitool. The BMC will need its own IP address to allow this access. The default userid will be root and the default password will be OpenBmc (0 is a zero). You will be forced to change this the first time you login. Once it is configured you can also use SSH to access the BMC.

The IC922 server has three 1 GB network ports with RJ-45 connectors on the back of the system. The right-hand port is the dedicated BMC port which would need its own network connection. The left-hand port is a shared port that can be

used to create a connection to the BMC as well as a network port for the operating system. On the HMC the latter is called failover mode -there are two MAC addresses and two IP addresses (one for the BMC and one for the operating system) but only one physical connection. The port in the middle is also for usage by the operating system. Section 2.1 of [this Redbooks publication \(http://www.redbooks.ibm.com/redpieces/pdfs/redp5584.pdf\)](http://www.redbooks.ibm.com/redpieces/pdfs/redp5584.pdf) gives a great overview of the BMC and also discusses the ipmitool, providing examples of how to use ipmitool as well as the OpenBMC tool to communicate with the BMC.

Operating System

At announce the only supported operating system was RHEL (Red Hat Enterprise Linux) ppc64le version 7.6-alt for POWER9. It's important that the correct version be downloaded; it must be the alt version not the version for POWER8. Other distributions should be added at a later date.

Summary

If you are looking for a scalable, powerful compute platform in a small package then it is well worth looking at the IC922. It is designed for high-performance data analytics and high-performance computing workloads running on Linux. In particular it excels with workloads such as inferencing data, cloud and heavy storage-oriented workloads. It integrates into an AI environment where training is run on the AC922 and the models are then passed to the IC922 for inferencing.

References

IC922 Announcement

http://www-01.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_sm/1/897/ENUS9183-_h01/index.html&lang=en&request_locale=en ([http://www-01.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_sm/1/897/ENUS9183-_h01/index.html&lang=en](http://www-01.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_sm/1/897/ENUS9183-_h01/index.html&lang=en&request_locale=en)).

<https://www.ibm.com/downloads/cas/US-ENUS120-017-CA/name/US-ENUS120-017-CA.PDF> (<https://www.ibm.com/downloads/cas/US-ENUS120-017-CA/name/US-ENUS120-017-CA.PDF>).

IC922 Adapter Placement Guide

https://www.ibm.com/support/knowledgecenter/POWER9/p9iaf/p9iaf_pcie_slot_details.htm (https://www.ibm.com/support/knowledgecenter/POWER9/p9iaf/p9iaf_pcie_slot_details.htm).

IC922 Technical Overview Redbook

<http://www.redbooks.ibm.com/redpieces/pdfs/redp5584.pdf> (<http://www.redbooks.ibm.com/redpieces/pdfs/redp5584.pdf>)

IC922 Data Sheet

<https://www.ibm.com/downloads/cas/LVYP6RM2> (<https://www.ibm.com/downloads/cas/LVYP6RM2>)

Tesla T4 GPUs

<https://www.nvidia.com/en-us/data-center/tesla-t4/> (<https://www.nvidia.com/en-us/data-center/tesla-t4/>)

Nigel Griffith's Youtube videos on the IC922

<https://www.youtube.com/watch?v=Jc6cFoj2104> (<https://www.youtube.com/watch?v=Jc6cFoj2104>)

<https://www.youtube.com/watch?v=KEuOAu2x97o> (<https://www.youtube.com/watch?v=KEuOAu2x97o>)

<https://www.ibm.com/support/pages/node/1488969> (<https://www.ibm.com/support/pages/node/1488969>)

Nigel Griffith's Handout and replay on the IC922 for the PowerVM User Group

http://public.dhe.ibm.com/systems/power/community/aix/PowerVM_webinars/93_IC922_Hands_On_Experience_Part1.pdf
(http://public.dhe.ibm.com/systems/power/community/aix/PowerVM_webinars/93_IC922_Hands_On_Experience_Part1.pdf)

http://public.dhe.ibm.com/systems/power/community/aix/PowerVM_webinars/93_IC922_Hands_On_Experience_Part2.pdf
(http://public.dhe.ibm.com/systems/power/community/aix/PowerVM_webinars/93_IC922_Hands_On_Experience_Part2.pdf)

http://public.dhe.ibm.com/systems/power/community/aix/PowerVM_webinars/93_IC922_Hands_On_Experience.mp4
(http://public.dhe.ibm.com/systems/power/community/aix/PowerVM_webinars/93_IC922_Hands_On_Experience.mp4)

About the author

Jaqui Lynch has over 38 years of experience working with a projects and Oses across vendor platforms, including IBM Z, UNIX systems and more.

Related Content

[Community \(/community\) A POWER9 Roadmap →](#) (<https://ibmsystemsmag.com/Trends/03/2017/power9-roadmap>)

[Community \(/community\) About Processor Modes →](#) (<https://ibmsystemsmag.com/Trends/05/2017/about-processor-modes>)

[Community \(/community\) POWER9: What's Already Out There Says Plenty →](#) (<https://ibmsystemsmag.com/Trends/09/2017/power9-whats-out-there>)

IBM Systems ^{MEDIA}

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.

© 2020 Key Enterprises LLC. All rights reserved