

A Closer Look at the Power Systems AC922

January 2018 | by [Jaqui Lynch](#)

Back in 2014 we experienced the rollout of the first POWER8 servers; the whole line has been available for some time now. Throughout 2017 IBM released information on the new POWER9 chip and then, on December 5, 2017, IBM announced the first POWER9 server—the AC922 HPC system.

The AC922 (8335-GTG) is a 2U Linux-only server that's designed for HPC (high performance computing) and AI (artificial intelligence) workloads. It is a 2 socket server with either 32 or 40 fully activated cores. Its architecture is intended for deep learning, AI, high performance analytics and HPC users. Although this box is targeted to a specific audience, it's a good indicator of the technology to come in 2018 with the rest of the POWER9 lineup.

Technical Specifications

The AC922 is a 2-socket server that goes into a standard 19-inch rack. It requires that both sockets be populated with POWER9 SCM (single chip module) processors. There are two options for the processors:

1. 16 core, 2.6ghz (3.09ghz turbo)
2. 20 core, 2.0ghz (2.87ghz turbo)

The AC922 provides for up to four NVIDIA air-cooled Tesla V100 attached with second generation NVLink to the processors. NVLink is a high-bandwidth, energy efficient interconnect between CPUs and GPUs, and between GPUs. It was initially introduced on the POWER8 S822LC and is enhanced in the POWER9 AC922. The IBM POWER is still the only server line that provides not only GPU to GPU acceleration, but also CPU to GPU acceleration.

There are 8 DDR4 memory DIMMs per socket and all DIMMs must be filled and activated. DIMM sizes are 16GB, 32GB or 64GB and cannot be mixed. This means the server will have either 256GB, 512GB or 1TB of memory. Having identical memory DIMM sizes and filling all the DIMMs ensures the best memory performance and bandwidth for the server.

1. Slot 1—x4 slot with an x8 connector
2. Slot 2—x8 slot with an x8 connector, CAPI enabled
3. Slots 3, 4—x16 slots, CAPI enabled

Slot 2 should be reserved for the new EDR IB (Infiniband) adapters if you plan to use them. These new adapters are designed for very high performance and to accept an x8 bus from each processor—the only slot that they can be installed into is slot 2.

There are two x 2.5-inch internal SATA disk bays that can hold 1TB or 2TB hard disks or 960GB, 1.92TB, or 3.84TB SSDs. There are also two optional NVMe high performance storage adapters. If NVMe is chosen, then the 1.6TB NVMe flash adapter is the default, but it is possible to use an RPQ and order the 3.2TB adapter. An external DVD can be attached as needed.

There are a number of I/O adapters available—these are all low profile adapters. Currently adapter options include:

- 4 port 1Gb network
- 2 port 40/100GB network
- 2 port 10Gb network
- 4 port 10/1gb (2 x 10Gb and 2 x 1Gb—both SR and SFP+ options for 10Gb)
- 1 and 2 port EDR 100GB IB
- 2 port 16GB/s fiber channel adapter
- 1.6TB NVMe flash adapter or optional via RPQ 3.2TB flash adapter

The AC922 uses two hot swap 2200-watt power supplies in a redundant configuration. According to the announcement there is no power redundancy if GPUs are installed.

The server comes in an Open Power non-virtualized configuration which allows Linux to be used on bare metal with no virtualization. The supported operating system is Red Hat Enterprise Linux 7.4 for Power LE (IBM POWER9), or later. IBM provides Supportline support for RHEL for 1 year or 3 years with either 9-5 or 7/24 coverage. Hardware maintenance defaults to 9-5 nbd (next business day) for 3 years but can be uplifted to same day or 7/24. Service is CRU (customer replaceable unit) but can also be uplifted to IBM on-site repair.

Other specifications

At the front of the server you will see the operator interface which consists of 1 x USB 3.0 slot, a power button and the service LEDs. There are also four cooling fans and the two disk bays that use SFF4 carriers. The hard drives and fans can be replaced without powering off the server.

At the rear of the server is the BMC (service processor card), which has the interfaces for IPMI, the VGA connection, 2 x 1Gb network connections and 1 more USB 3.0 slot. The IPMI connection is a serial connection that can be used with the IPMI tool to manage the server. The top network connection is the network for the IPMI/BMC connection and the bottom network connection is for direct OS usage.

According to the site and planning guide, the following are the physical and environmental specifications:

	2 GPUs	4 GPUs
BTU/hr	6485	8533
Watts	1900	2500
kVA	1.96	2.575

This is with single phase 220-240v power.

Minimum Configuration

The minimum Power AC922 initial order must include two processor modules, 256 GB of memory, two power supplies, two line cords, rack-mounting hardware, a system software indicator, a rack integrator specify, and a Language Group Specify. The operating system is Linux and there is also a line item for “Open Power non-virtualized configuration”—this is the OPAL (Open Power Abstraction Layer) firmware that allows Linux to run on bare metal on POWER.

Management

The service processor (BMC or FSP) supports IPMI 2.0 (Intelligent Management interface), DCMI 1.5 (Data center management interface) and SNMP v2 and v3 (simple network management protocol). The IPMItool can be used to interface with the service processor. The OpenBMC tool also provides a communication method to the BMC using a CLI (command line interface). It can be used from either a

remote Linux system or from the host operating system console window. The IPMItool has the ability to monitor or manage the server it is running on or any other server on the LAN that supports IPMI. There are multiple versions of the IPMItool including ones that run on Windows and Linux. There is also a systems management GUI available through a web browser depending on the state of the system.

Summary

The AC922 is the first in a series of POWER9 servers to be rolled out. While it's targeted at the high throughput and performance needed for Linux workloads in the HPC, deep learning and AI spaces, it's a clear indicator of the kind of technology and performance that should be expected from the new POWER9 line. If you have any plans in the AI, deep learning or HPC space then the AC922 is well worth investigating.

References

PDF files for the AC922

https://www.ibm.com/support/knowledgecenter/en/8335-GTG/p9hdx/8335_gtg_pdf_files.htm

AC922 Site and Planning Guide

http://public.dhe.ibm.com/systems/power/docs/hw/p9/p9ia4_83x.pdf

AC922 Announcement

<https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=AN&subtype=CA&htmlfid=897/ENUS117-111&appname=lenovous&language=en>

AIX Virtual User Group (lots of topics including POWER9)

<http://www.tinyurl.com/ibmaixvug>

OpenBMC

http://public.dhe.ibm.com/systems/power/docs/hw/p9/p9eih_openbmc.pdf

IPMItool

<https://www.ibm.com/support/knowledgecenter/en/linuxonibm/liaai.ipmi/liaaiipmitool.htm>

NVIDIA NVLink on POWER

<http://www.nvidia.com/object/nvlink.html>

<https://blogs.nvidia.com/blog/2016/09/08/ibm-servers-nvlink/>

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.