

From here to AI with IBM Power enable professionals the world needs today and tomorrow

AI-DL Skill Build - Opportunity

Around 133 million new jobs will be created by AI by 2022 according to a World Economic Forum (WEF) report.

Is your university sufficiently prepared to train the workforce of tomorrow? Does it have the differentiated infrastructure and the associated skills required to bridge the gap between industry and academia?

IBM with its decades of experience in reinvention and innovation is here to help. It'll help you set up the infrastructure and build the skills among students and your faculty for the AI era.

This will impart outstanding industry-grade skills to our young students and will make them prepare for any challenges that the Smarter planet will bring to them in their exciting careers ahead.

Why IBM Power Systems?

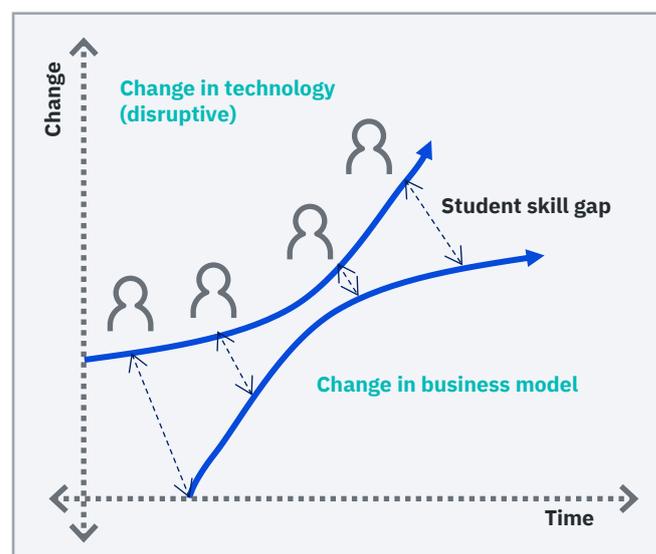
The IBM Power® System AC922 is the next generation of the POWER9 processor-based systems, which are designed for deep learning and artificial intelligence (AI), high-performance analytics, and high-performance computing (HPC) that will help in bridging the gap between industry and academia.

The system is co-designed with OpenPOWER Foundation members and is deployed at the most powerful supercomputer on the planet with a partnership between IBM, NVIDIA, Mellanox, and others.

The POWER9 Architecture is premier Acceleration Platform designed for big data and AI workloads with embedded NVlink in CPU and has superior performance. It provides the current technologies

that are available for HPC, improving even more the movement of data from memory to GPU accelerator cards and back, enabling faster and lower latency data processing.

Unprecedented performance and application gains with advanced IO interfaces integrated into the P9 processor delivering capabilities not available on x86



POWER9 Features

Enhanced core and chip architecture for next-gen workloads

- Optimized for emerging algorithms to interpret and reason
- I/O processor chip provides superior bandwidth, scale and capacity to ingest and analyse data

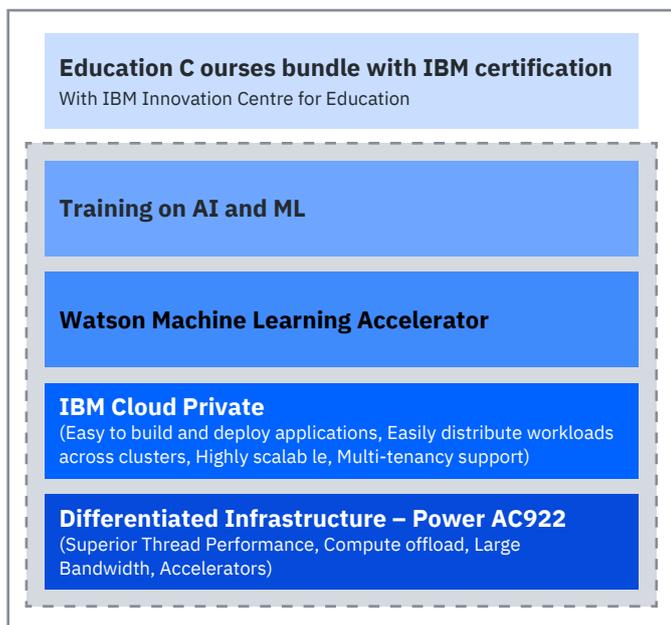
Newest processor chip family with optimized silicon for a range of platforms

- Scale out for HPC and next-gen apps
- Scale up to 16 sockets

Premier platform for accelerated computing with maximized bandwidth

- Ultra-high bandwidth CPU-GPU interconnect delivers ~ 4x faster training
- OpenCAPI™ interface offers high bandwidth communication to NICs, FPGA accelerators, and storage controllers

IBM Centre of Excellence for Educational Institutes



Power AC922 Key Features

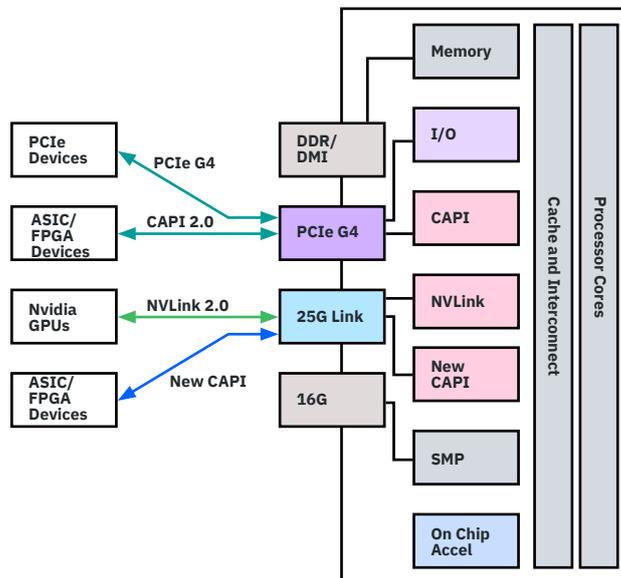
- 1. An acceleration superhighway**
Unleash state of the art IO and accelerated computing potential
- 2. Designed for the AI era**
Architected for the modern analytics and AI workloads
- 3. Delivering enterprise - Class AI**
Flatten the time to AI value curve by accelerating the journey to build, train, and infer deep neural networks
- 4. Faster iterations**
Designed for building more accurate models with high data throughput and faster AI model training time.

5. Flexible deployment

Build and train AI models with flexible deployment options

6. Open innovation

Use popular open frameworks, software and tools



Power AC922

- High-performance tensor processing is supported mainly by GPUs. V100 GPUs are the latest product that support deep learning.
- Only IBM can provide the CPU-GPU NVlink which is required for the project.
- Built for the AI-Era, the POWER9 supports up to 5.6x more I/O and 4x more threads than its x86 contemporaries.
- In Power AC922 NVlink 2.0 communication happens between CPU and GPU which gives major performance boost compared to x86 where NVlink communication is only between 4X GPU's.
- There are coherent interfaces that enable the GPUs to directly communicate with the CPU memory structures with full coherency (Coherent access to system memory)
- Power AC922 has provision for up to 256 Gbps of input bandwidth on single IO slot with PCI x16 Gen4 (allowing for future proofing).

State of the art I/O and acceleration attachment signaling

- PCIe Gen 4 x 48 lanes – 192 GB/s duplex bandwidth
- 25G Link x 48 lanes – 300 GB/s duplex bandwidth

Robust accelerated compute options with OPEN standards

- On-Chip Acceleration – GZip x1, 842 Compression x2, AES/SHAx2
- OpenCAPI – High bandwidth, low latency and open interface using 25G Link
- NVLink 2.0 – Next generation GPU ← → CPU bandwidth and integration

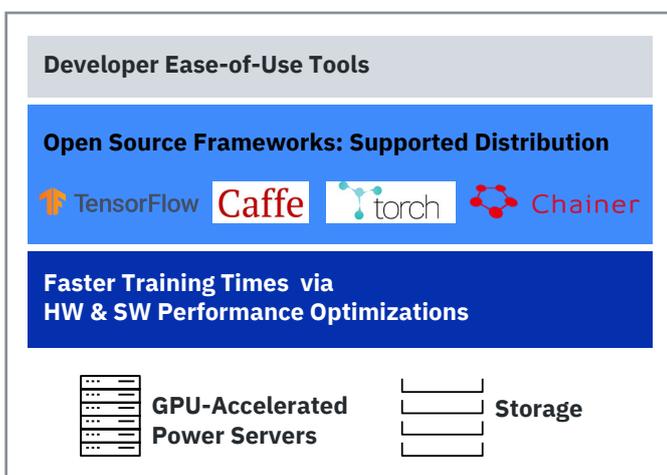
Watson Machine Learning Accelerator

It is a complete environment for data science as a service, enabling educational institutions to bring AI applications into production. It bundles IBM PowerAI, IBM Spectrum Conductor, IBM Spectrum Conductor Deep Learning Impact

It combines popular open source deep learning frameworks, efficient AI development tools, and accelerated IBM® Power Systems™ servers.

Higher productivity for data scientists- Speeds training time as there is no need to navigate setup operations in addition complex tuning and optimization challenges.

Customer results have proven this process 3-5x “faster to deployment” than a normal “do it yourself” deep learning implementation.



WML-Accelerator Features

- **Enables rapid deep learning deployment**
Accelerate time to deploy deep learning workloads and get results quicker from initial prototype to enterprise-wide deployment.
- **Delivers faster insights at scale**
Exploit unique data architecture and turn deep learning performance into value with IBM Power Systems and other available hardware.
- **Supports larger, more accurate models**
Leverage Large Model Support and Distributed Deep Learning to support models at higher resolution, up to almost any scale.
- **Business ready and dependable**
Expand your deep learning deployment to multiple users and lines of business with multi-tenancy and role-based access controls. Protect them with end-to-end security.

IBM Cloud Private

Tool provided for full lifecycle management of containers. It leverages open- source Kubernetes.

Component suggested is Community Edition (CE).

Advantages of ICP-CE

- **Unified installer**
Rapidly set up a Kubernetes based by using an Ansible based installer.
- **ICP management console**
Manage, monitor, and troubleshoot your applications from a secure management console.
- **Isolate tenant networks**
Calico allows for improved performance and network isolation inside your cluster.
- **Robust monitoring and logging**
ICP uses Elasticsearch, Logstash, Filebeat, and Heapster for the collection, storage, and querying of logs and metrics.
- **Flexible management**
Containerized applications in combination with the private registry allow for flexibility

- **Ideal platform for managing DL workloads**
Container managers are a good fit for handling long-running yet interactive workloads
- **Private Docker image registry**
This local registry allows you to restrict which users can view or pull images.

IBM Innovation Centre of Education

IBM together with the leading Universities across the world is offering a series of futuristic B. Tech. CSE/IT & MBA programs, under the IBM Innovation Center for Education program.

There is more focus on Project based learning and Industry interaction. This will ensure that the students' work is constantly under review, and the IT specialist and architects from the industry help them in their learning.

Adequate training to the faculty and a support structure for their constantly being in touch with industry subject matter experts.

There are touch points included from the first semester itself. Moving onto later semesters, the labs and project-based learning as well as the industry interactions increases proportionately.

A Futuristic Delivery Framework

We ensure delivered support, in a manner that makes the learning effective for the students.

- **Hands-on learning**
The theory courses are augmented by adequate labs and project work, which will help the students apply the knowledge they gain in classroom.
- **Campus sessions**
Campus sessions are organized through IBM ICE Day, Hackathons, Poster competition, Summer/ Winter School, Gurukool and Project Days.
- **Student monitoring and tracking**
The platform tracks the work the student puts into the ecosystem. This includes all the project completion, the support that he/she provides/

receive to other team members, and the comments they receive from their guides. The program has been designed much like the real-life scenario where a professional is judged against some set parameters at every stage.

- **Vibrant concurrency of the courseware**
The courses will be reviewed and adjusted at periodic intervals for their concurrency. This is one of the key advantages of IBM's participation. This support comes from the IBM Labs resources and other connected groups and partner organizations.
- **Online discussion forums**
The learning will be further enhanced by providing adequately structured Discussion forums against relevant topics /subjects. These may be moderated by IBM and industry experts as needed.
- **SME visits and interaction**
Visits to the Campus will be planned by IBM and partner Subject Matter Experts (SME's), so that students can learn from their experiences and skills.
- **Webinar sessions**
There will be webinars conducted by IBM and partner SMEs. During these interactive sessions, the SMEs will share the latest in technology with the students.

This is a key element in technology concurrency at the Campus. When the students and industry specialist interact, in addition to the learning they receive from the skilled faculty and high-quality courseware, enhancing the skills many fold.

A Futuristic Delivery Framework

We ensure delivered support, in a manner that makes the learning effective for the students.

- **Hands-on learning**
The theory courses are augmented by adequate labs and project work, which will help the students apply the knowledge they gain in classroom.

- **Campus sessions**

Campus sessions are organized through IBM ICE Day, Hackathons, Poster competition, Summer/ Winter School, Gurukool and Project Days.

- **Student monitoring and tracking**

The platform tracks the work the student puts into the ecosystem. This includes all the project completion, the support that he/she provides/ receive to other team members, and the comments they receive from their guides. The program has been designed much like the real-life scenario where a professional is judged against some set parameters at every stage.

- **Vibrant concurrency of the courseware**

The courses will be reviewed and adjusted at periodic intervals for their concurrency. This is one of the key advantages of IBM's participation. This support comes from the IBM Labs resources and other connected groups and partner organizations.

- **Online discussion forums**

The learning will be further enhanced by providing adequately structured Discussion forums against relevant topics /subjects. These may be moderated by IBM and industry experts as needed.

- **SME visits and interaction**

Visits to the Campus will be planned by IBM and partner Subject Matter Experts (SME's), so that students can learn from their experiences and skills.

- **Webinar sessions**

There will be webinars conducted by IBM and partner SMEs. During these interactive sessions, the SMEs will share the latest in technology with the students.

This is a key element in technology concurrency at the Campus. When the students and industry specialist interact, in addition to the learning they receive from the skilled faculty and high-quality courseware, enhancing the skills many fold.

What's in it for you?

Leadership

First amongst peers to start AI/Deep Learning Elective/ Mandate courses with the AI global leader

Branding Associated

- IBM co-Branded lab at the campus
- Joint Press release with IBM
- Social media endorsements
- "IBM AI Expert" certification* for student

Opportunity for industry collaboration

- Opportunity for selected faculties to visit IBM research labs.
- Opportunity for students to showcase their deep learning projects

Opportunity for short term training programs

Opportunity to offer high in demand short term training programs to industry professionals and outside students with the cost.

Student industry readiness & employability

Prepare students for challenges and skills required that an AI driven economy and computing model will demand

Deliverables

- Power AC922 Server
- 3 years warranty for HW & SW
- 24 X 7 OEM support
- One-time installation of HW & SW
- On-Premises training
- 80-400 hours courses on AI-ML

