

# Transform your business through automated artificial intelligence

Innovation requires an artificial intelligence (AI) strategy, and **H2O Driverless AI** on **IBM Power Systems™** is designed to accelerate your ability to create intelligent products and services that quickly deliver trusted results.

H2O Driverless AI on Power Systems can:

Scale with  
**2.6x**  
more RAM<sup>1</sup>

Transfer data  
**9.5x**  
faster versus PCIe 3.0<sup>2</sup>

Build  
**4x**  
larger AI models<sup>3</sup>

Process data  
**5.6x**  
faster<sup>4</sup>

The above claims are based on a comparison of the IBM POWER9 processor and x86 Intel Xeon Scalable Platform family.  
PCIe = Peripheral Component Interconnect Express

## The benefits of H2O Driverless AI and Power Systems can include:



### Automatic machine learning development

Use prepared exploratory data analysis to generate visualizations

Perform streamlined and customizable model documentation

Harness open and extensible automated machine learning (ML) optimization



### Faster and more accurate results

Provide predictions designed to be easily explained

Acquire reason codes and model interpretability in plain English

Apply interpretability for debugging and regulation



### Industry-leading interpretability

Increase accuracy with automatic feature engineering

Use GPU acceleration for 3.7x reduction in AI model training<sup>5</sup>

Build automatic reports including K-LIME, Shapley, Variable Importance and more



### Quick deployment into Power Systems

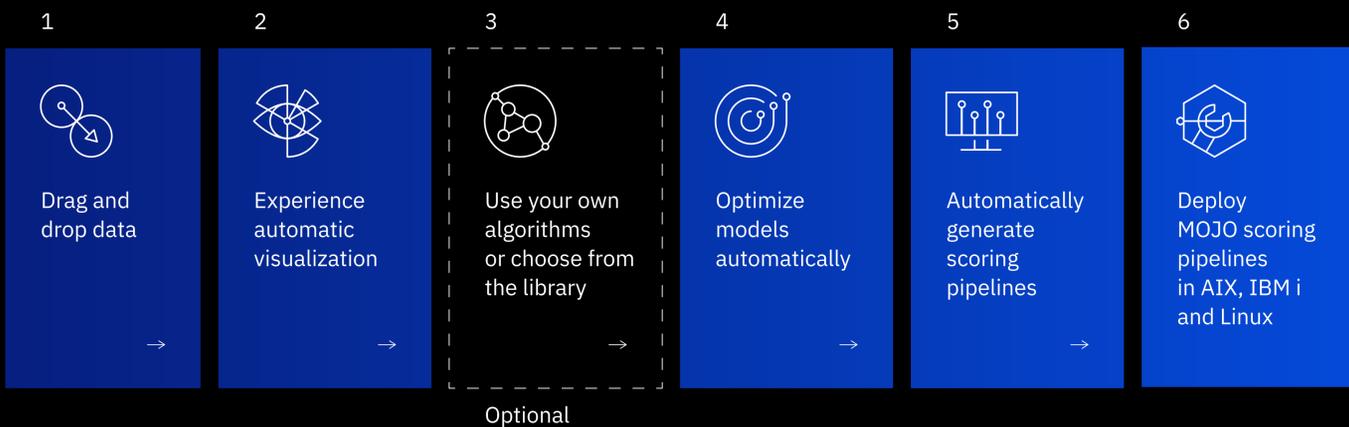
Launch Java and Python MOJOs<sup>6</sup> more easily

Add H2O AI MOJOs runtime to customer applications

Works with existing customer applications in IBM AIX®, IBM i and Linux® operating systems

## How Driverless AI works

By employing AI tools across your IT infrastructure, you can gain insights faster.



## The value of enterprise-ready Power Systems

Move AI applications into production faster and offer new capabilities and insights

Help reduce AI model training time with highly parallelized CPUs and GPUs

Experience systems designed to ease deployment for low latency models in Python, Java and R



CPU = computer processing unit  
GPU = graphics processing unit

## Why IBM?

Together, H2O Driverless AI with Power Systems helps differentiate you from your competitors through automated AI.

You can create AI-driven solutions through an automated ML platform and start transforming your business today. Contact your IBM representative or IBM Business Partner to help you build a plan to accelerate your journey to AI.

[Learn more](#) →



© Copyright IBM Corporation 2019. IBM, the IBM logo, ibm.com, AIX, and Power are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Red Hat®, JBoss®, OpenShift®, Fedora®, Hibernate®, Ansible®, CloudForms®, RHCA®, RHCE®, RHCSA®, Ceph®, and Gluster® are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

69028269USEN-01

1. 2.6X memory capacity is based on 4TB per socket for the IBM POWER9 processor family compared to 1.5TB per socket for x86 Intel Xeon Scalable Platform family - product brief: <https://www.intel.com/content/dam/www/public/us/en/documents/product-briefs/xeon-scalable-e-platform-brief.pdf?asset=14606>

2. 9.5X is based on POWER9 and next-generation NVIDIA NVLink peak transfer rate is 150 GB/sec = 48 lanes x 3.2265625 GB/sec x 64 bit/66 bit encoding compared to x86 PCI Express 3.0 (x16) peak transfer rate is 15.75 GB/sec = 16 lanes X 1GB/sec/lane x 128 bit/130 bit encoding

3. 4x compares 2TB system memory in AC922 with 4GPUs vs Xeon-based server with 16 NVIDIA Tesla V100 GPUs, each with 32GB RAM (16 x 32GB = 512GB GPU RAM)

4. 5.6x I/O bandwidth claim based on CUDA H2D Bandwidth Test conducted on a Xeon E5-2640 V4 +P100 vs Power9 with NVLink 2.0, V100 (12 GB/s vs 68 GB/s rated) - Power AC922 with POWER9 processor and NVLink 2.0 is the ONLY system in the industry with NVLink 2.0 from the CPU to the GPU

5. Results are based IBM Internal Measurements using Power AC922; 40 cores with NVLink 2.0 vs. 2-Xeon E5-2640 v4 with 20 cores 4XTesla V100 GPU. Using Chainer v3 with Large Model Support running on POWER9 shows the value of NVLink 2.0 and the performance capability that it can deliver when the problem set becomes larger than the memory supported on the GPU cards. These tests move large amounts of data between the CPU and the GPU and can reduce model training times by 3.7X

6. H2O can allow you to convert the models you have built to either a Plain Old Java Object (POJO) or a Model Object, Optimized (MOJO). A MOJO (Model Object, Optimized) is an alternative to H2O's POJO. As with POJOs, H2O can allow you to convert models that you build to MOJOs, which can then be deployed for predictive scoring in real time. H2O-generated MOJO and POJO models are intended to be easily embeddable in any Java environment