Press the ML Easy Button

Unleash data science innovation.

Unleash data science innovation with Validated Designs for AI with automatic machine learning and streamlined machine learning operations.

The Dell Validated Design for AI — Automatic Machine Learning (AutoML) automates machine learning (ML), making it easier for data scientists to train the best model in the least amount of time in a VMware®-virtualized environment.

The integrated solution includes H20.ai AutoML to automate ML, cnvrg.io for ML operations (MLOps), NVIDIA® AI Enterprise Suite™ for cloud-native artificial intelligence (AI) development and deployment, delivered on VMware vSphere® with Tanzu™ deployed on an engineering-validated and optimized Dell infrastructure stack.

This integrated solution simplifies AI, speeding time to AI insights delivered with proven engineering expertise to deliver automation at scale.

Easy for IT

Run AI workloads alongside existing applications.

Many organizations already use VMware vSphere for server virtualization and NVIDIA GPUs for workload acceleration. However, as data grows, data centers are starting to hit performance and scale limitations — especially when it comes to AI.

Building separate systems for AI workloads — with different hardware, software and toolsets — leads to management complexity, poor utilization rates and lengthy provisioning times. Data scientists might have to wait weeks or even months for IT to provision the resources required for AI development.

Dell Validated Designs for AI are jointly engineered and validated to power AI workloads alongside existing applications in VMware-virtualized data centers with near-bare-metal performance.

Disaggregating resources at a more granular level optimizes hardware investments — including server GPUs — increasing utilization and creating an incredibly flexible compute foundation for AI, all on infrastructure that is easy for IT to deploy and manage.

DELLTechnologies

These solutions offer:

AI simplified
Jointly engineered and validated solutions make it quick and easy to deploy optimized hardware and software stacks.

Faster AI insights
Automated ML and streamlined operations delivered with the AI tools and frameworks in an optimized infrastructure enables faster time to production for development and for IT teams.

Proven AI expertise
Confidently deploy an engineering-tested AI solution backed by world-class Dell Technologies and NVIDIA services and support. Select ProSupport Plus for a single point of contact for software and hardware support.

Easy for data scientists

As AI becomes more strategic, data science teams need to streamline and simplify workflows to continue adding value at scale. AutoML with H20.ai Driverless AI and MLOps with cnvrg.io speed and streamline bringing the best AI models into production.

Harness MLOps to improve AI results with cnvrg.io.

Moving ML models into production is complex, with data science teams often spending more time dealing with infrastructure, feature engineering, deployment and monitoring challenges than on creating groundbreaking algorithms.

cnvrg.io MLOps was built by data scientists, for data scientists to streamline the ML process so they can focus on innovation, not infrastructure. MLOps addresses the challenges of moving ML models into production by reducing friction and bottlenecks that impede AI success.

This Dell Validated Design for AI, built in collaboration with cnvrg.io, standardizes ML pipelines to minimize friction for data science and engineering teams from research to production.

Learn more about the Dell Validated Design for AI — MLOps.

Train ML models directly from GitHub.

Code
Plan
Deploy
Operate
Build
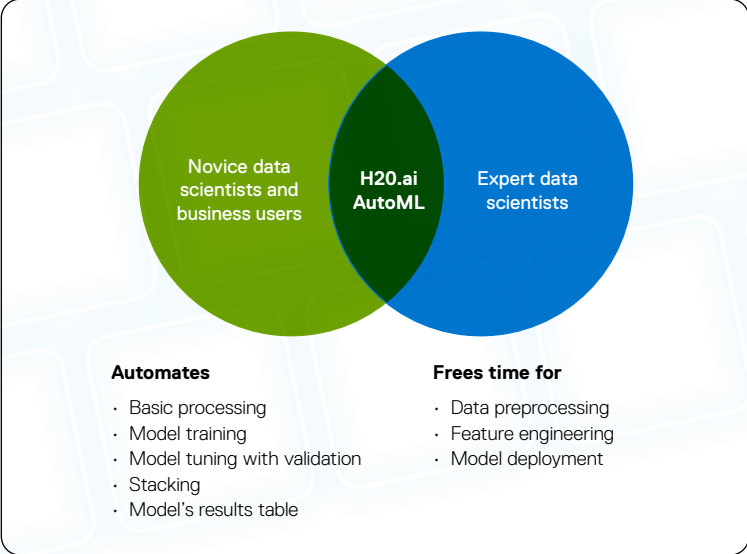Release
Test
Measure
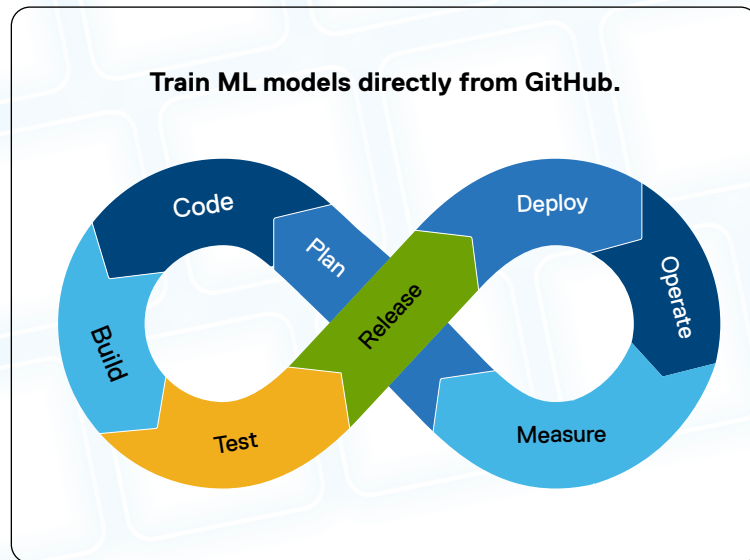
Source: cnvrg.io

Automate ML to scale AI success with H2O.ai AutoML.

As AI and ML are becoming more essential, the shortage of ML experts threatens to throttle innovation. Even for experienced data scientists, it can take a lot of time to train and tune AI models. H2O.ai strives to automate ML, making it easier and faster to create and launch groundbreaking AI applications.

AutoML from H2O.ai automates algorithm selection, feature generation, hyperparameter tuning, iterative modeling and model assessment. Automating these repetitive tasks makes it easier and faster to train and evaluate ML models and enables teams to focus on the business problems they are trying to solve.

The Dell Validated Design for AI — H2O.ai AutoML automates machine learning, making it easier for data scientists to train the best model in the least amount of time in a VMware-virtualized environment.

Learn more about the
Dell Validated Design for AI — H2O.ai AutoML.

Automates
- Basic processing
- Model training
- Model tuning with validation
- Stacking
- Model's results table

Frees time for
- Data preprocessing
- Feature engineering
- Model deployment

Source: H2O.ai

AI for virtualized environments

Configuration options

Configurations are based on Dell Technologies NVIDIA-Certified Systems™ with NVIDIA AI Enterprise suite, vSphere virtualization and VMware Tanzu containers. Fast, ample storage is provided by Dell PowerScale all-flash or hybrid storage arrays.

Compute

4x Dell VxRail HCI V670

or

4x Dell PowerEdge R750xa

with

NVIDIA A100 or A30 GPUs

Networking

Dell PowerSwitch 25GbE S5248F-ON

or

NVIDIA Spectrum® SN3700 and out-of-band PowerSwitch S4148T-ON

Storage

PowerScale F600

or

H600

Software

• H2O.ai AutoML
• cnvrg.io MLOps
• NVIDIA AI Enterprise
• VMware vSphere with Tanzu

NVIDIA technologies are built in.

The VxRail hyperconverged infrastructure (HCI) or PowerEdge servers at the heart of your solution comes with integrated NVIDIA technologies that help speed ML workloads — and results.

vGPUs

NVIDIA Virtual GPUs (vGPUs)
NVIDIA vGPU software enables sharing GPU resources across multiple virtual machines to make them accessible to any device, anywhere.

MIGs

NVIDIA Multi-Instance GPUs (MIGs)
NVIDIA MIGs expand the performance and value of GPUs by partitioning them into as many as seven instances to support every workload and extend accelerated resources to more users.

A100

NVIDIA A100 GPUs
Accelerate AI workloads with up to 20X higher performance over the prior generation.[1] The A100 supports NVLink bridge, the world's first high-speed GPU interconnect offering a significantly faster alternative for multi-GPU systems than traditional PCIe-based solutions.[2]

[1] NVIDIA website, Accelerating the Most Important Work of Our Time, accessed June 2022. [2] NVIDIA website, NVIDIA NVLink™, accessed June 2022.

NVIDIA technologies are built in.

The VxRail hyperconverged infrastructure (HCI) or PowerEdge servers at the heart of your solution comes with integrated NVIDIA technologies that help speed ML workloads — and results.

AI

NVIDIA AI Enterprise
NVIDIA AI Enterprise is an end-to-end, cloud-native suite of AI and data science tools and frameworks optimized and exclusively certified by NVIDIA to run on VMware vSphere with NVIDIA-Certified Systems. It includes key technologies for the rapid deployment, management and scaling of AI workloads.

Certified
Systems

NVIDIA-Certified Systems
As NVIDIA-Certified Systems, VxRail HCI and PowerEdge bring together NVIDIA GPUs, NVIDIA ConnectX® smart network interface cards (SmartNICs) and NVIDIA BlueField DPUs in optimized configurations. These are validated for performance, manageability, security and scalability and are backed by enterprise-grade support from NVIDIA and Dell Technologies.

Solution highlights

VxRail HCI
Joint engineering between Dell Technologies and VMware combined with the simplicity of VxRail HCI System Software and the performance of the underlying next generation PowerEdge servers provides an ideal platform for virtualized AI. Factory installed VMware vSphere delivers powerful support for the latest NVIDIA GPUs, delivering enterprise virtualization for GPU powered AI workloads.

VMware vSphere with Tanzu
VMware vSphere with VMware Tanzu runs more than 70 million workloads across hundreds of thousands of organizations worldwide. VMware vSphere transforms bare-metal servers into centrally managed AI and ML infrastructure pools that can quickly provision virtual machines and containers on demand.

PowerEdge servers
PowerEdge servers are engineered to deliver unmatched performance and versatile configurations to meet the demands of AI workloads. Flash storage, the latest processors and memory with flexible local storage make PowerEdge servers a foundational choice for AI.

PowerSwitch networking
Based on open standards, PowerSwitch networking frees the data center from outdated, proprietary approaches. Our future-ready networking technology helps you improve network performance, lower networking costs and remain flexible to adopt new innovations.

PowerScale storage
Designed to reduce complexity and optimize results for data-intensive AI workloads, PowerScale gives you scalable, available storage that eases adoption and management. PowerScale is available in hybrid and all-flash models to quickly deliver large amounts of data to AI models for faster, more accurate results.

Why Dell Technologies

Collaborate at worldwide Customer Solution Centers.
Collaborate with Dell Technologies engineering teams at one of our worldwide Customer Solution Centers, tap into the resources of one of our HPC & AI Centers of Excellence or test and tune real-world systems at the HPC & AI Innovation Lab.

Consume AI as-a-Service with Dell Technologies APEX.
With simple and consistent cloud experiences delivered as-a-Service (aaS), APEX can help you get the AI-optimized solutions you need to fast-track intelligent outcomes everywhere. APEX can deliver a cloud operating model for AI on-premises, off-premises and at the edge, so you can create measurable value from data at any scale.

Speed success with services.
Dell Technologies Services include consulting, deployment, support and education to help drive the rapid adoption and optimization of AI environments from initial set up and upskilling of resources through to ongoing support. Managed Services and Residency Services can help reduce the cost, complexity and risk of managing IT so you can focus resources on digital innovation and transformation.

Learn more

DELLTechnologies

Get a feel for enterprise AI.
NVIDIA LaunchPad is a free program
that provides short-term access to a
large catalog of hands-on labs so you can
experience end-to-end solution workflows
in the areas of AI, data science, 3D design
collaboration and simulation, and more.

NVIDIA.