

The background features several decorative images of server racks, each enclosed in a rounded rectangular frame. These frames are arranged in a staggered pattern across the slide. The server racks themselves are dark with various components like fans, lights, and cables visible.

Lenovo AI/ML Innovations

15 November, 2018 | Dave Weber, Director & CTO, Lenovo

Lenovo – AI End-to-end Strategy

DATA CENTER & CLIENT SOLUTIONS

\$1.2 Billion
AI Investment

“AI has changed everything and big data analysis is what large companies depend on.”



Yang Yuanqing
CEO - Lenovo



AI Research

Providing tools and expertise to accelerate AI innovation.



Enterprise AI Solutions

Empowering organizations to launch their AI initiatives.



AI Exploration

Democratizing AI for all.

100+ Data Scientist
& Developers

AI Research &
Innovation Centers

AI HW & SW
Platforms

End-to-end Solutions



IoT



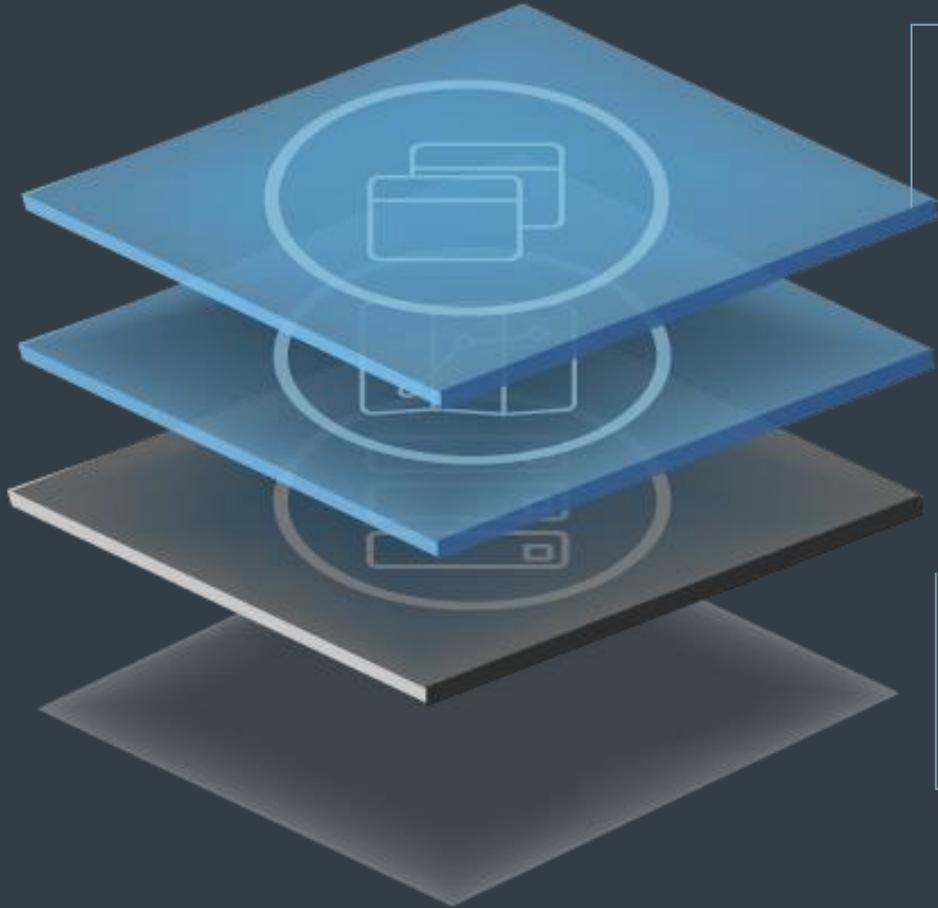
Big Data



AI

LiCO - A Platform For AI Development

LiCO - Lenovo Intelligent Computing Orchestration



Web Portal



Pre-trained Models



Workflow Templates



Distributed AI Training



AI Frameworks



Open-Source Cluster Management



Optimized HW Libraries



LiCO AI Software Roadmap



Lenovo in Artificial Intelligence

Driving AI in focus verticals



Healthcare & Life Science

Health Screening & Early Diagnosis



Academia & Research

AI Powered Research



Manufacturing

Quality control & Predictive Maintenance



Energy / O&G

Improving uptime & Efficiency



Financial Services

Risk Monitoring, Anomaly Detection,
Augmented Trading Strategies

Other Use
Cases

Develop on Lenovo AI Products

The right resources to get started

TRAINING PLATFORMS



ThinkSystem SR650
2U, 2 Intel Xeon,
Big Data Platform



ThinkSystem SD530
2U, 2 x 2 Intel Xeon,
2 x 2 NVIDIA V100



ThinkSystem SR670
2U, 2 Intel Xeon,
4 x NVIDIA V100



LiCO v5.1

Increasing performance 



LiCO for AI
A single software stack
to efficiently manage
HPC and now AI
workloads

INFERENCE PLATFORMS



ThinkSystem SD530
2U, 4 x 2 Intel Xeon



ThinkSystem SR650
2U, 2 Intel Xeon,
2 NVIDIA V100



ThinkSystem SD530
2U, 2 x 2 Intel Xeon,
2 x 2 NVIDIA V100

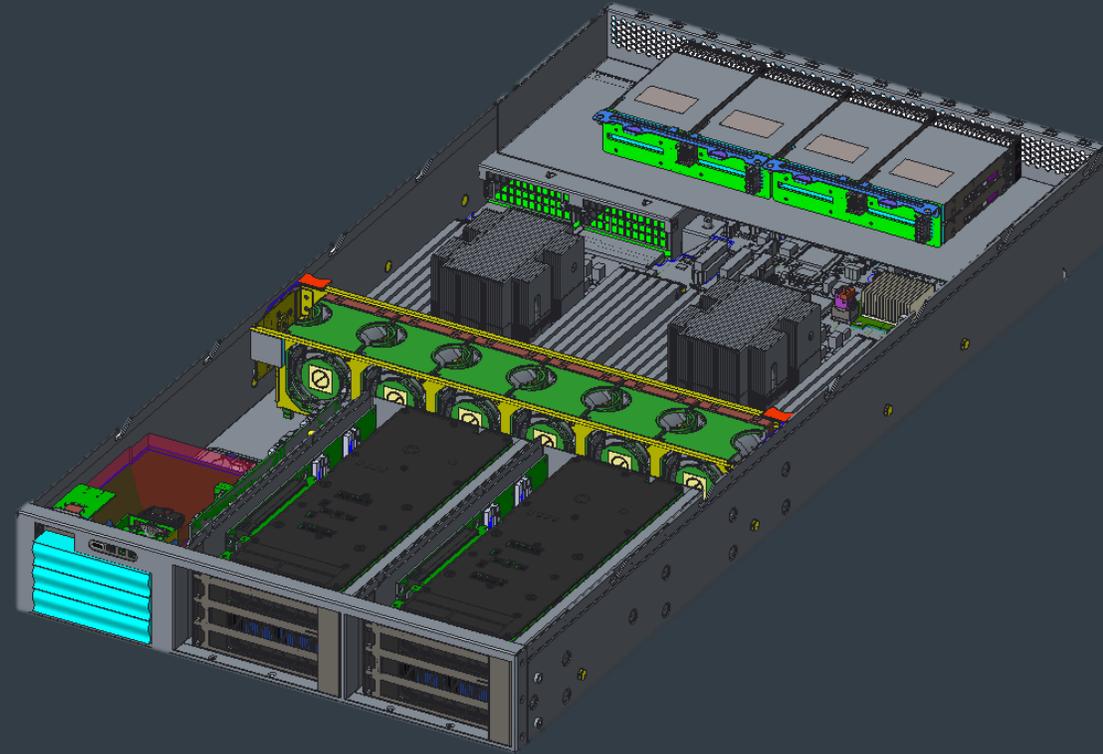


ThinkSystem

2018 Lenovo. All rights reserved

Lenovo AI/ML SR670: 2U, 2S, 4 GPU System

- Mechanical
 - 2U, 19" Rack Server
- CPU
 - 2 Sockets (including 205W sku's)
- Memory
 - Up to 24 DIMMs, 16/32 GB
- GPU
 - Up to 4 double wide passively cooled GPUs
 - Up to 8 passively cooled single wide FHHL GPUs
- Storage
 - M.2 Adapters (Single or Dual Drive)
 - 8 SAS/SATA HDD/SSDs
- Network
 - Up to 3 PCIe Adapters
 - Dedicated GbE Management Port
- OS Support
 - Ubuntu, RHEL7.4, SLES13, Windows, VMware



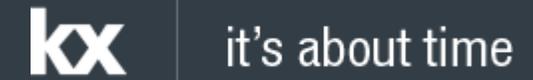
Lenovo AI/ML Drop In Cluster

- 1x 1U [SR630](#) management node
- 1x 2U [SR650](#) compute/ML inference / training node(s)
 - GPU-ready
 - NVMe & NVMeoF-ready
 - FPGA-ready
- 1x 2U [SR670](#) AI/ML training node(s)
 - Up to 4x GPU's / FPGA's per two sockets
- 1x 1U 32 port [Lenovo 100 GbE switch](#)
- [LiCO](#) 90 day eval SW license
- Lenovo on-prem, cloud-based dev-ops – Kubernetes/Docker via Red Hat OpenShift
- RHEL 7.4
- Lenovo Professional Services (LPS) for setup and initial customization
- 90 day Try & Buy for the servers and switch
- Optional: Lenovo Data Scientist consultations

Lenovo Servers – STAC[®] Leadership Performance



- SR950 config:
 - 4x 8180's (28C per), 3 TB DDR4
 - Kdb+ 3.5, RHEL 7.5
 - 6 TB Intel DC P4800X (Optane) SSDs
 - 56 TB Intel DC P4510 (3D NAND) SSDs.
 - Spectre & Meltdown mitigation fixes applied



- First-ever single 4S server to run STAC-M3[™] Kanaga suite (48 TB)
 - SUT ID: KDB181009: <https://stacresearch.com/KDB181009>
 - Set 5 new records, including:
 - 58% faster for 5 year snapshot
 - 40% faster for 50-user 12 day VWAB in Year 1





Different is better