

Smarter technology for all

Lenovo agile infrastructure for next-gen workloads

Grigory Pryalukhin
Senior Business Development Manager | Lenovo ISG

Lenovo



Made in Europe for Europe

- Lenovo delivers customized solutions in the heart of Europe from our manufacturing plant in Ullo, Hungary.
- Our Ullo plant shipped one million workstations & servers in its first year, enabling deeper collaboration, transparency, and understanding with Lenovo customers across the EMEA region.





The widest portfolio

One of the world's most comprehensive portfolios of AI-ready devices, AI-enabled solutions, and AI-optimized experiences

Lenovo ThinkReality	Lenovo ThinkSmart	Lenovo ThinkBook
Lenovo ThinkStation	Lenovo ThinkVision	Lenovo ThinkPad
Lenovo ThinkPad X1	Lenovo ThinkShield	ThinkPhone by motorola
Lenovo YOGA	 motorola	Lenovo LEGION

Global footprint and supply chain security

#1 Asia/Pacific Supply Chain Top 10
Gartner

#3 High Tech Supply Chain
Gartner

#8 Global Supply Chain
Gartner

Lenovo Zero Trust Supply Chain

Lenovo Trusted Supplier Program

ISO 9001:2015
Certified

30+ Manufacturing sites
180 Markets

Secure
Facilities & shipping

- Research center
- Manufacturing
- Company location

Our Commitment: investing in EECA markets

**10
Markets**



**>100% YoY
Growth**

**200+
Covered
Customers**

**30+
ASP
(Service
Providers)**



**>\$1M
Investment in
Premier Support
with Local Stocks
of Spare Parts**



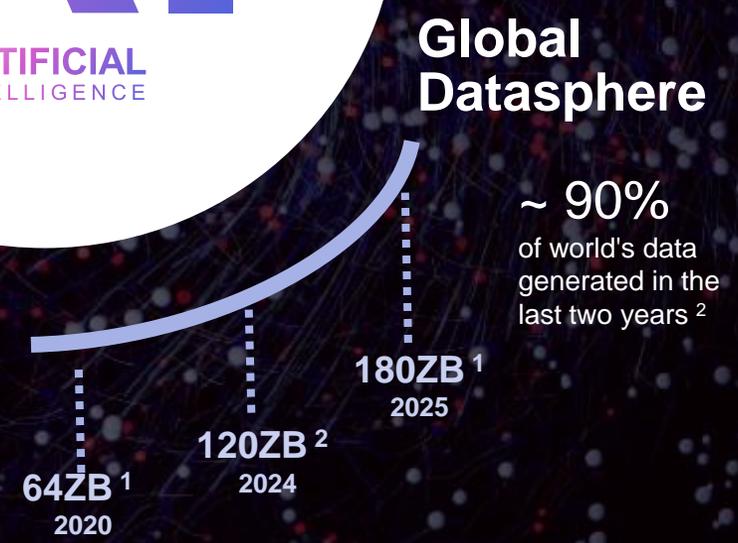
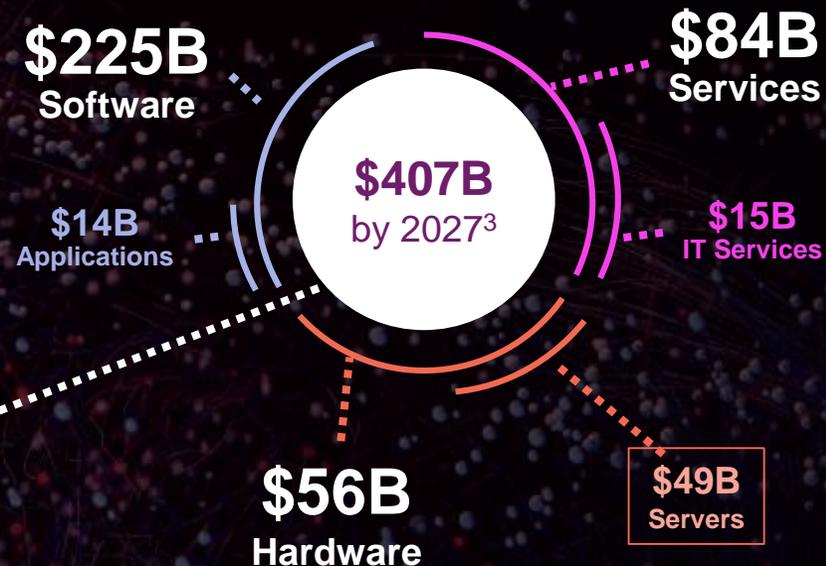
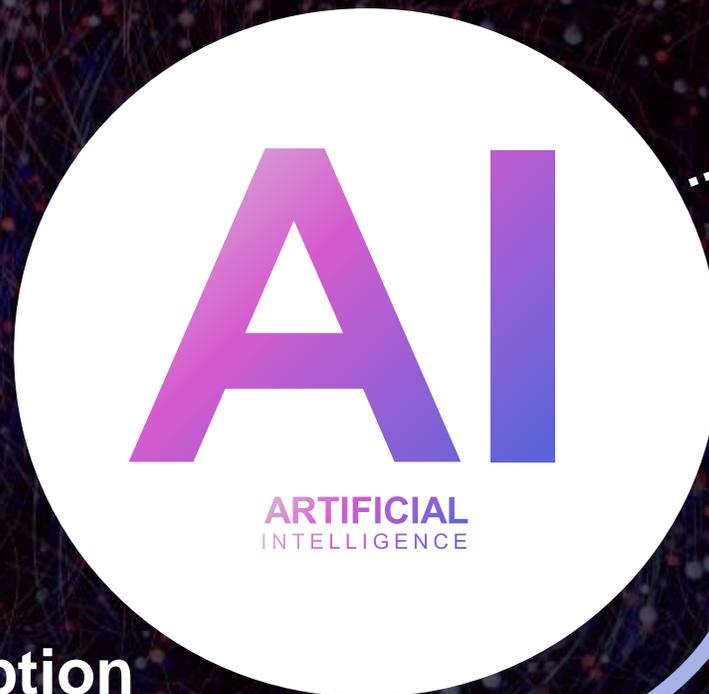
**30+
Gold & Platinum
Partners**

**#1 OEM for
Nutanix**

**Local
Professional
Service
Engineers**



The AI revolution is radically changing the IT landscape



Bring AI to the Data



\$2.2B

AI Investment
Commitment



#3

Global AI
Infrastructure
Provider



80

AI Optimized
Infrastructure
Products



4

Global
AI Innovation
Centers



165

AI Solutions in
AI Innovators
Program

AIFI Aifi Refresh
Aifi Build
Aifi To-Go

ADD-FOR Sybille Automotive

alwaysAI alwaysAI Smart Retail
alwaysAI Smart Construction
alwaysAI Smart Manufacturing
alwaysAI Smart Venues

ARTISIGHT Patient Room of the future
Operating room of the future
Clinic of the future
Command Center of the future
Assisted Living Facility of the future

Aupera Aupera VMSS for Retail
Aupera VMSS for Smart Cities
Aupera VMSS for Manufacturing

ASCENDION AI Services

BIKAL Bikal AI Platform
Bikal AI Security
Bikal AI Services

byte LAKE CFD Suit
Cognitive Services Suit (Manufacturing)
Cognitive Services Suit (Automotive)
Cognitive Services Suit (QSR)

C2RO C2RO Entera for Manufacturing
C2RO Entera for Retail / Mall
C2RO Entera for Smart City

chooch Chooch AI Vision Platform
Chooch Image Chat

consight Consight Fortify
Consight RESTful API

cnvrg.io cnvrg.io platform

deci Deci AutoNAC

Dataloop Dataloop Data Management Platform
Dataloop LIDAR
Dataloop Automation Platform

EAGLE EAIGLE Automated Vehicle Access Control (AVAC)
EAIGLE Access Control Solution for Visitors
EAIGLE Density Analytics
EAIGLE Thermography Monitoring Solution

everseen Evercheck Visual AI Platform
Everdoor Visual AI Platform
Evershelf Visual AI Platform

Fingermark Supersonic
Eyecue
Digital Menu Boards

FedML FedML AI App Ecosystem
FedML Open source
FedML Octopus
FedML Parrot
FedML Beehive
FedML Spider
FedML Cheetah

nybl.

GRAYMATICS LabVista
UrbanVision
siteinsights
Vista Mart
Vista Karma
G-Safe

HU/EX Automated Intelligent Drive-Through Assistant (AIDA)
Analytics Dashboard
Conversational AI Core Engine

DEEPBRAIN AI

Infinia ML Infinia ML Document Processing

InstaDeep™ DeepChain
DeepPack
DeepPCB

kinetic vision Industrial Metaverse
Retail Metaverse

IBM maximo IBM Maximo Asset Management
IBM Maximo Predictive Maintenance
IBM Maximo Remote Monitoring
IBM Maximo Visual Inspection

NEARBY COMPUTING NearbyOne AI Platform

NR NEUREALITY NR1 NAPU
NR1 M
NR1 S

>50
AI Software Providers

>30k
Channel partners

Lenovo
AI Innovators
Partner Ecosystem

>165
Enterprise AI Solutions

180
Countries served

vmware Vsphere
Aria UI

v7 V7 Labs Labellig Platform

VSaaS Video Surveillance as a Service AI Platform

trigo StoreOS
EasyOut
EasyStock

SUNLIGHT Sunlight HyperConverged Edge
Sunlight NexCenter
Sunlight NexVisor

wandisco Data Activation Platform
Edge to Cloud
Data Migrator

WaitTime WaitTime AI Platform

VSBLTY DataCaptor
VisionCaptor
VSBLTY Vector
Retail Hardware Solutions

VOCALYTICS AI Machine Listening Software

VISTRY Vistry AI Drivethru
Vistry AI Kitchen
Vistry AI Platform

WATAD Robotics Process Automation
Computer Vision
Supervisory Control and Data Acquisition System (SCADA)
Power Substation Automation (PSA)

THE EDGE COMPANY BCMS Ventur
BCMS Ventur X

intel.

NVIDIA

AMD

Qualcomm

vmware

Red Hat

Azure Stack

IBM Watson Studio

Portfolio Overview - From Far Edge to Core

Edge Client



SE10/SE10-I



-20C to +60C
Intel ATOM
Up to 32 GB RAM



M90n-1



0°C to +50°C
Intel Celeron, i3
Up to 8GB RAM



SE30



-20°C to +60°C
Core™ i3, i5
Up to 16GB RAM



SE50



0°C to +50°C
Core™ i3, i5, i7
Up to 32GB RAM

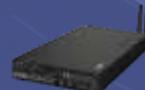


SE70



-20°C to 60°C
NVIDIA Jetson
Up to 16GB RAM

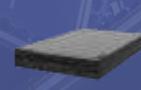
Edge Server



SE350*



0°C to 55°C
Xeon® D-2100
Up to 256GB RAM



SE350V2



0°C to 55°C
Xeon® D-2700
Up to 256GB RAM



SE360V2



-20°C to 65°C
Xeon® D-2700
Up to 256GB RAM



SE450



-5°C to 55°C
Xeon® SP 3rd Gen
Up to 1TB RAM

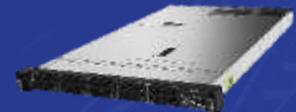


SE455 V3



-5°C to 55°C
AMD EPYC 8004
Up to 576GB RAM

Datacenter



SR630, SR645



0°C to 35°C

Gateways

Diverse Portfolio

From ultra compact gateways to data center grade products.

Base Station / DC Closet

Highest Performance

CPU & GPU rich systems for ultimate performance.

Data Center

Flexible Deployment

In harsh environments with ruggedized devices & unique cooling capabilities.

Delivering tomorrow's oil & gas experience today with AI

Lenovo AI Innovators

It provides certified AI solutions, and helps customers build up their own through our AI Discover Labs



- Behavior analysis
- Customer journey
- Front & back of Store analytics



Adv. Multi-phase simulations

GRAYMATICS

- Employee Safety and compliance
- Automated Security perimeters



- Well-Lifecycle
- Predictive Maintenance
- Reservoir Management

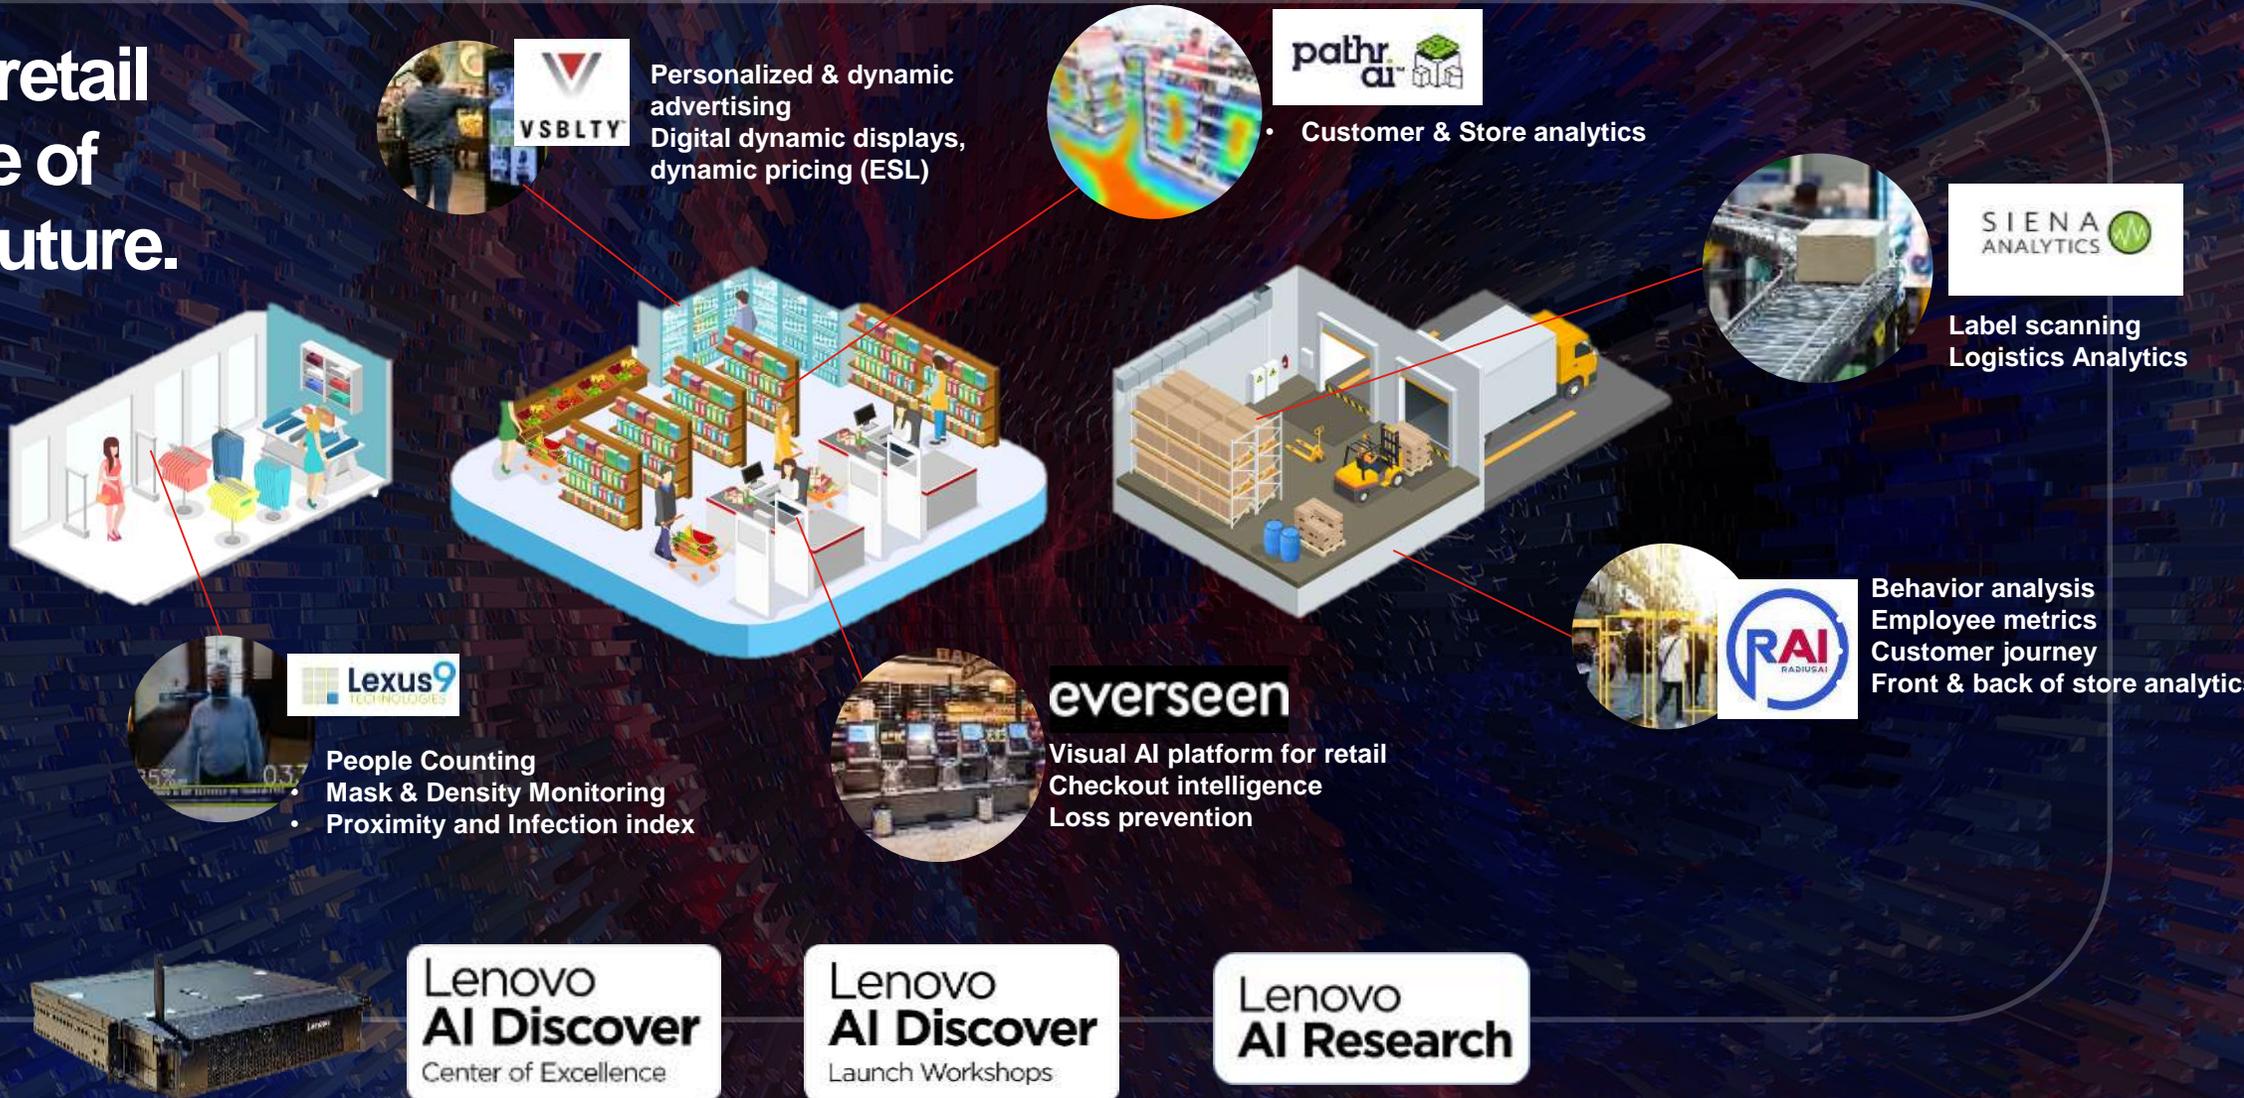
Or build your own AI solutions with:



- Automated Visual Inspection
- Comp. Fluid Dynamics
- Bespoke AI Solutions

Delivering tomorrow's retail experience today with AI

The retail store of the future.



Lenovo delivering tomorrow's Banking experience today with AI

The Bank of the Future

- Data Security
- Video & Audio Recognition
- Fraud Prevention & Security



- CX Focused Transaction Zones
- Smart Devices for Customer Service



- AI Ready Servers
- Virtual Assistants
- Remote Conferencing
- Edge Computing



- Smart ATMs
- Assisted Online Banking
- Automation
- Digital Twins

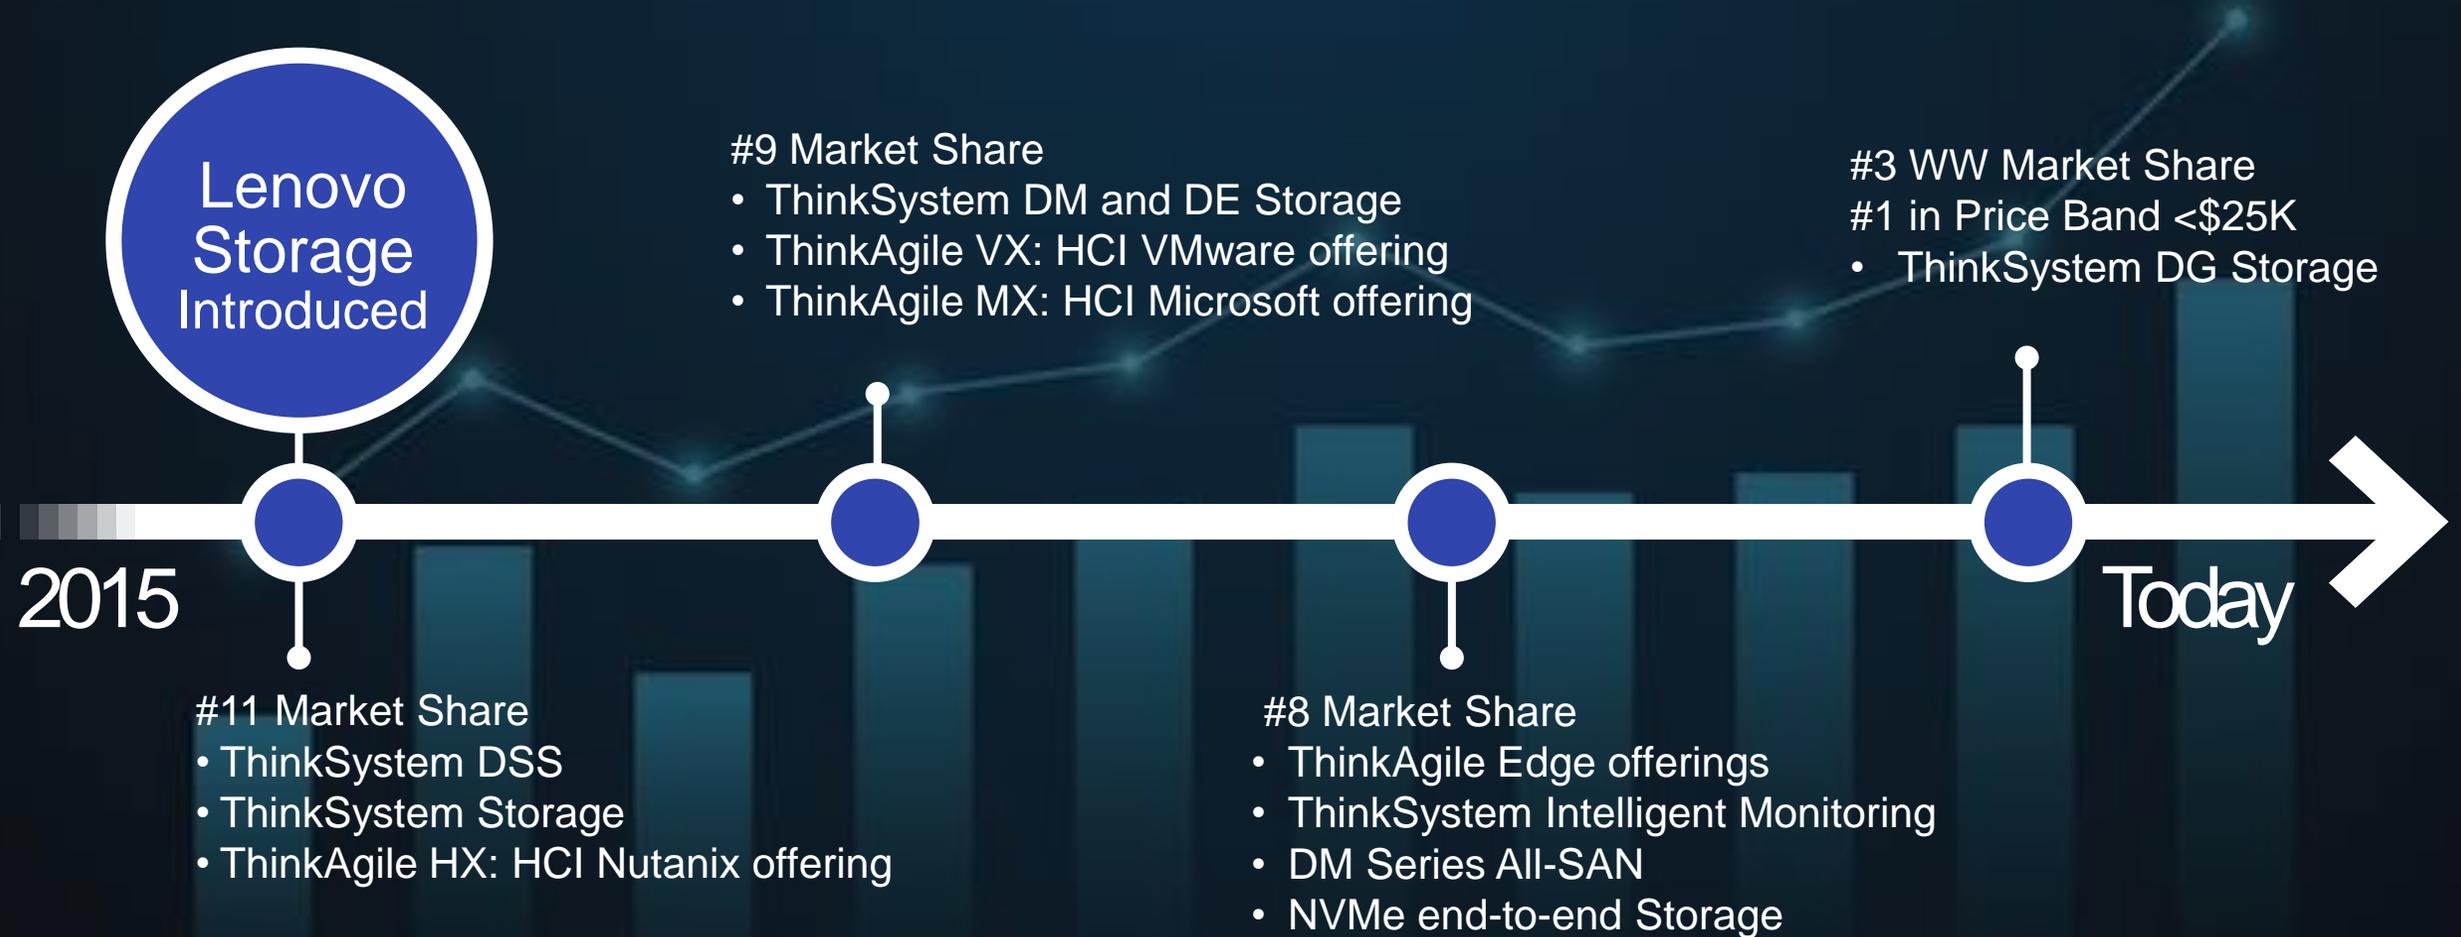


Lenovo
AI Discover
Center of Excellence

Lenovo
AI Discover
Launch Workshops

Lenovo
AI Research

Lenovo Data Management and Storage History



Lenovo Storage Introduced

2015

#11 Market Share

- ThinkSystem DSS
- ThinkSystem Storage
- ThinkAgile HX: HCI Nutanix offering

#9 Market Share

- ThinkSystem DM and DE Storage
- ThinkAgile VX: HCI VMware offering
- ThinkAgile MX: HCI Microsoft offering

#8 Market Share

- ThinkAgile Edge offerings
- ThinkSystem Intelligent Monitoring
- DM Series All-SAN
- NVMe end-to-end Storage

#3 WW Market Share

- #1 in Price Band <\$25K
- ThinkSystem DG Storage

Today

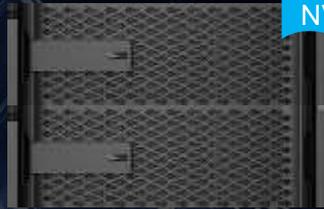
ThinkSystem

Performance SAN



NVMe

DE Series



NVMe

DM Series

ThinkSystem

Enterprise Unified



QLC

DG Series



NVMe

DM Series

ThinkSystem

Data Infrastructure



Storage
Networking



Tape Solutions



D Series

ThinkSystem

Software Defined Storage (SDS) Solutions



DSS-G



WEKA

ThinkSystem DE Flash & Hybrid Arrays

Workload focus: Databases, AI, analytics, backup and archive, media, video surveillance

Performance

Best value for performance
All-Flash NVMe Array up to 2M IOPS,
44GB/s throughput

Scalability

All-Flash scalability to 1.84PB. Hybrid scalability to
8.64PB



Reliability

Fully redundant, active-active
architecture. Up to six nines availability.

Data Protection

Simplified snapshots and asynchronous
and synchronous replication

Efficiency

Easy to configure, manage, and scale
Improve the overall efficiency of your
IT operations.

Deployment Flexibility

iSCSI, Fibre Channel , NVMe/IB,
NVMe/RoCE, and NVMe/FC

ThinkSystem DE Series Storage Array Portfolio

All-Flash



DE4000F (2U24)



DE6000F (2U24)



DE6400F
(2U24 NVMe)



DE6600F
(2U24 NVMe)

Hybrid



DE2000H (2U24)



DE2000H (2U12)



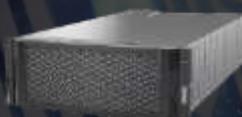
DE4000H (4U60)



DE4000H (2U24)



DE4000H (2U12)



DE6000H (4U60)



DE6000H (2U24)



DE6400H
(2U24 NVMe+SAS Expansion)



DE6600H
(2U24 NVMe+SAS Expansion)

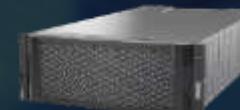
SAS Expansion



DE120S (2U12)



DE240S (2U24)



DE600 (4U60)

ThinkSystem DM Flash & Hybrid Arrays

Workload focus: AI, analytics, virtualization, VDI, unstructured datasets, edge consolidation

Performance

Superior all-flash performance with NVMe over TCP & Fibre Channel

Efficiency

Unified file, block, and object storage for advanced data efficiency

Scalability

All-flash scalability to 88.5PB. Hybrid scalability to 57PB



Data Protection

Integrated ransomware protection and immutable snapshots and copies

Reliability

Fully redundant, active-active architecture, synchronous replication (MetroCluster), and snapshots

Hybrid Multi-Cloud

Integrated hybrid multi-cloud: Amazon, Azure, Google, and IBM

ThinkSystem DM Series Storage Array Portfolio

All-Flash



DM5000F



DM5100F



DM7100F

Hybrid



DM3000H



DM3010H



DM5000H



DM7100H

DM Series Expansions



DM120S (2U12)



DM240S (2U24)



DM240N (2U24)



DM600S (4U60)

ThinkSystem DG Arrays

Workload Focus: AI/ML and analytics, data lakes, media/rendering, backup consolidation

Price/Performance

Up to 6x faster at 50% less cost* with budget-friendly high-capacity all-flash QLC storage

Scalability

Up to 17.6PB raw capacity in 24-node NAS cluster NAS and up to 8.8PB raw capacity in 12-node SAN cluster



Sustainability

Delivering substantial power savings versus hybrid arrays, enabling workload consolidation to reduce rack space and datacenter footprint

Security

Best-in-class data security and ransomware protection

**Internal calculations comparing Lenovo ThinkSystem DG Enterprise Storage Arrays to previous generation hybrid models (June 2023).*

Customers want to transform from hybrid to all flash

But price is often a limiting factor

Hybrid Flash



DM7100H



DM5000H



DM3000H

Performance Flash



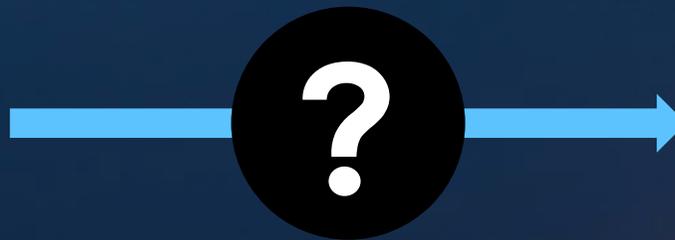
DM7100F



DM5100F



DM5000F



5–10ms latency, \$'s

Sub-ms latency, \$\$\$'s

ThinkSystem DG Series: Accelerate All-Flash Transformation

Workload Focus: AI/ML and analytics, data lakes, media/rendering, storage consolidation

Hybrid Flash



DM7100H



DM5000H



DM3000H

Capacity Flash



DG7000



DG5000



Performance Flash



DM7100F



DM5100F



DM5000F

5–10ms latency, \$'s

2–4ms latency, \$\$'s

Sub-ms latency, \$\$\$'s

ThinkSystem DG Series: TCO & Sustainability

Reduce up to 80% of rack space with high-density capacity flash



Hybrid Flash Array



*Based on internal testing results. Lenovo ThinkSystem DG5000 vs. previous generation hybrid array, August 2023

ThinkSystem DG Series Storage

Business Benefits



Up to 6x faster at 50% less cost compared to HDD arrays



Reduce rack space and achieve better energy efficiency compared to traditional HDD arrays



Keep data available, protected, and secured with Lenovo Unified Complete



Manage data on-premises or in the cloud



DG is performance optimized and rapidly scalable to meet the demands of the tuning and inference stages in Enterprise AI LLMs



DM/DG Series is born Feature Rich

Protocols	Efficiency	Anti-ransomware	Data protection	Tenancy and scale
<ul style="list-style-type: none"> NAS / File <ul style="list-style-type: none"> NFS SMB SAN / Block <ul style="list-style-type: none"> FC iSCSI NVMe/FC NVMe/TCP Object (S3) 	<ul style="list-style-type: none"> Leading raw-to-usable Deduplication, compression, compaction Thin provisioning Data tiering to hybrid, object, or cloud 	<ul style="list-style-type: none"> Near-real-time detection Tamper-proof Snapshot recovery Multi-admin verification Blocking of known malicious files 	<ul style="list-style-type: none"> App-integrated backup Business continuity Sync/async disaster recovery Snapshot Clones Immutable 	<ul style="list-style-type: none"> Multitenancy QoS Clusterwide scaling

ONTAP Data management software

Delivering automation, efficiency, data protection, and security capabilities for file, block, and object



At-a-Glance: DG Series

All-flash, high-capacity flash storage systems

Sustainable

- Outstanding price/performance
- Leading raw-to-effective data efficiency
- Tier to hybrid flash or cloud
- Better energy efficiency

Secure

- Ransomware protection
- Business continuity, disaster recovery
- Application-integrated backup and recovery

Scalable

- Entry capacity aligned with needs
- Clustered scalability
- Connect to performance flash or hybrid flash

Cloud-connected

- Disaster recovery to cloud
- Back up to cloud
- Tier to cloud



Form factor: 2U/4U (+2U Expansions)

Max drive count:

DM7000: 96x 30TB SSD = 2.94PB Raw

DG5000: 48x 30TB SSD = 1.47PB Raw

Storage Protocol support:

NAS (File): NFS, CIFS/SMB

SAN (Block): iSCSI, FC, NVMeoF (FC & TCP)

OBJECT: S3

Lenovo Software Defined Storage (SDS) Portfolio

High Performance File Solution (HPFS) powered by WEKA

- High-performance, large-scale workloads
- AI/Analytics, Finance High-frequency Training, Genomics/Life Science, Media & Entertainment

Increase
productivity
up to

90%

Improvement in deep
learning time

Improve
resource ROI

75%

Reduction in workload
wall clock time

Lower
storage costs

65%

Savings over traditional high-
performance computing



Object Storage Solution powered by Clouidian

- Capacity-intensive enterprise workloads
- Backup/Archive, Video/Images, Healthcare/Life Sciences, Research, IoT, Video Surveillance, Data Lakes

Lower
Storage costs

70%

Compared to traditional
storage systems

Reduce primary
backup tier cost
up to

40%

Superior RTO/RPO for fast
restore times

Reduce TCO
by

33%

Compared to legacy storage
systems

ThinkSystem



Lenovo ThinkSystem Ready Nodes for HS350X V3

Workload focus: Object and scale-out file, backup and recovery, Big data analytics, multimedia streaming, surveillance, AI/data lakes, data observability, video capture/retrieval, and content delivery networks

Optimized Infrastructure

Single-socket storage server built for maximum capacity.

Powerful Efficiency

Supports the latest Intel Xeon Scalable processor family for maximum efficiency and peak computational resources.



Maximum Density

Supports 24 3.5" hard drives in a 2U form factor (8 drives per tray).

Hybrid Cloud-ready

Scalable, efficient, and cost-effective infrastructure for your on-premises or cloud workloads.

Reliable

Lenovo continues to lead the industry in best uptime among all x86 platforms.*

Innovative Design

The three-tray design allows users to access the hot-swap drives without powering off the system.

Lenovo SAN Storage Solutions



ThinkSystem Storage

- Simplified data management
- Entry to High performance block
- Flash and Hybrid models
- Easy to configure, manage, and scale.

ThinkSystem Fibre Channel HBA and SAN support 32 Gb/s, NVMe, and now Autonomous.



Lenovo Servers & Emulex FC HBAs

Brocade
GEN6
FIBRE CHANNEL



DB610S



DB620S



DB630S

Brocade
GEN7
FIBRE CHANNEL



DB720S



DB730S



Directors



Lenovo Data Management

Multi-Workload
Enterprise Storage

Performance Tuned Storage
Single-Workload Exa-scale

Tier 0

- RPO 0, RTO 0, Sync Replication
- Low Latency
- High IOPS
- NVMe
- Application Integration for Snapshots

Tier 1

- Integrated Workflow Automation
- Snapshots Thin provisioning
- Data Efficiency (Compression, Compaction, Deduplication)
- Adaptive Quality of Service
- Application Based Balanced placement (Performance)
- Local Storage Caching

Tier 2

- Scale out Storage Clusters
- High Throughput
- Dynamic capacity expansion
- Remote Storage Caching
- Asynchronous replication
- Backup Software Integration

Tier 3

- S3 Object Storage Integration w/ Cloud
- Encryption & Trusted Platform Module (TPM) support
- WORM (Write Once Read Many) Data protection
- AWS, Google, Azure Integration and other Cloud Providers
- Dual Parity RAID Protection, Enhanced Parity (Erasure Coding, etc.)
- Ransomware Protection

ThinkSystem

DM\DG Series

ThinkSystem

DE Series

ThinkAgile

VX HX MX

ThinkSystem

Tape

Lenovo

DSS-G

ThinkSystem

WEKA

ThinkSystem

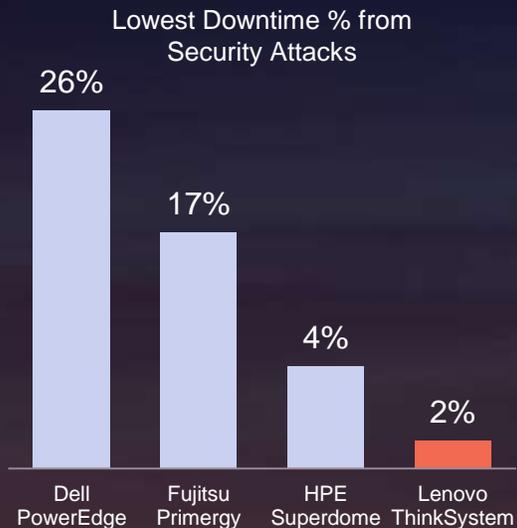
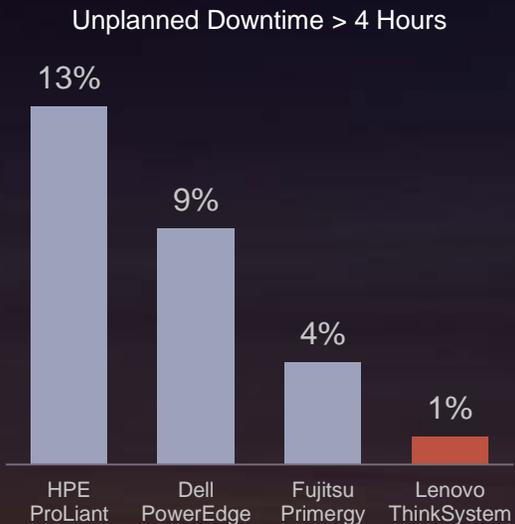
CLOUDIAN

- Millions of IOPS
- 100's Gbps of throughput
- nVidia GPU Direct
- Snapshots & Cloning

- Cloud Bursting
- Snap to S3
- WORM protection
- 10's Gbps of throughput

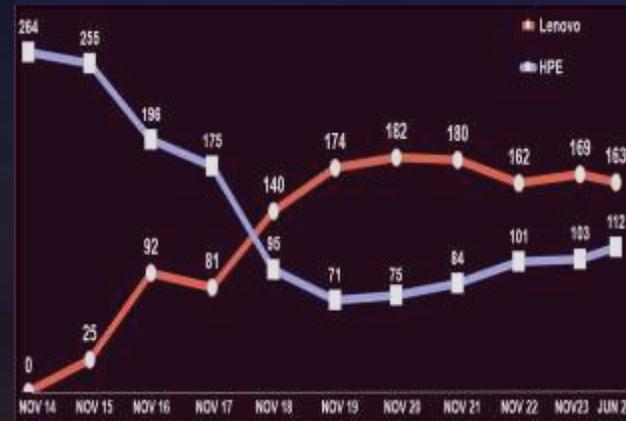
- Encryption
- Data Immutability

ISG Server Leadership



588
Current World Record Workloads

>6x as many #1 Results



#1

x86 Server Reliability for 10 Consecutive Years

#1

x86 Security for 5 Consecutive Years

#1

Current World Record Performance Workloads

#1

Top 500 Supercomputers for 6 Consecutive Years

Source: ITIC (Mar 2024), Benchmark Report (July 2024) & Top500.org (June 2024)

Lenovo™

#1 Server Reliability

ITIC, March 2024

(Information Technology Intelligence Consulting Corp.)



Note: 'To maintain objectivity, ITIC accepted no vendor sponsorship'*

'Lenovo's ThinkSystem servers provide the best reliability among all x86 servers for the 10th consecutive year and best x86 server security for the last five years'*

'averaging between six nines or 31.5 seconds of per server unplanned annual downtime and seven nines or 3.15 seconds of unanticipated per server yearly downtime'*

*Source: [ITIC 2023 Global Server Hardware, Server OS Reliability Report](#)

Information Technology Intelligence Consulting
ITIC

We've achieved **>6X** as many current #1 world records as any other server vendor

588 World records in workload performance

248

World records set by servers running 2nd 3rd and 4th Gen AMD EPYC™

340

World records set by servers with Intel® Xeon® Scalable Processors

71

World records set by servers with 4th and 5th Generation Intel® Xeon® Scalable Processors

184

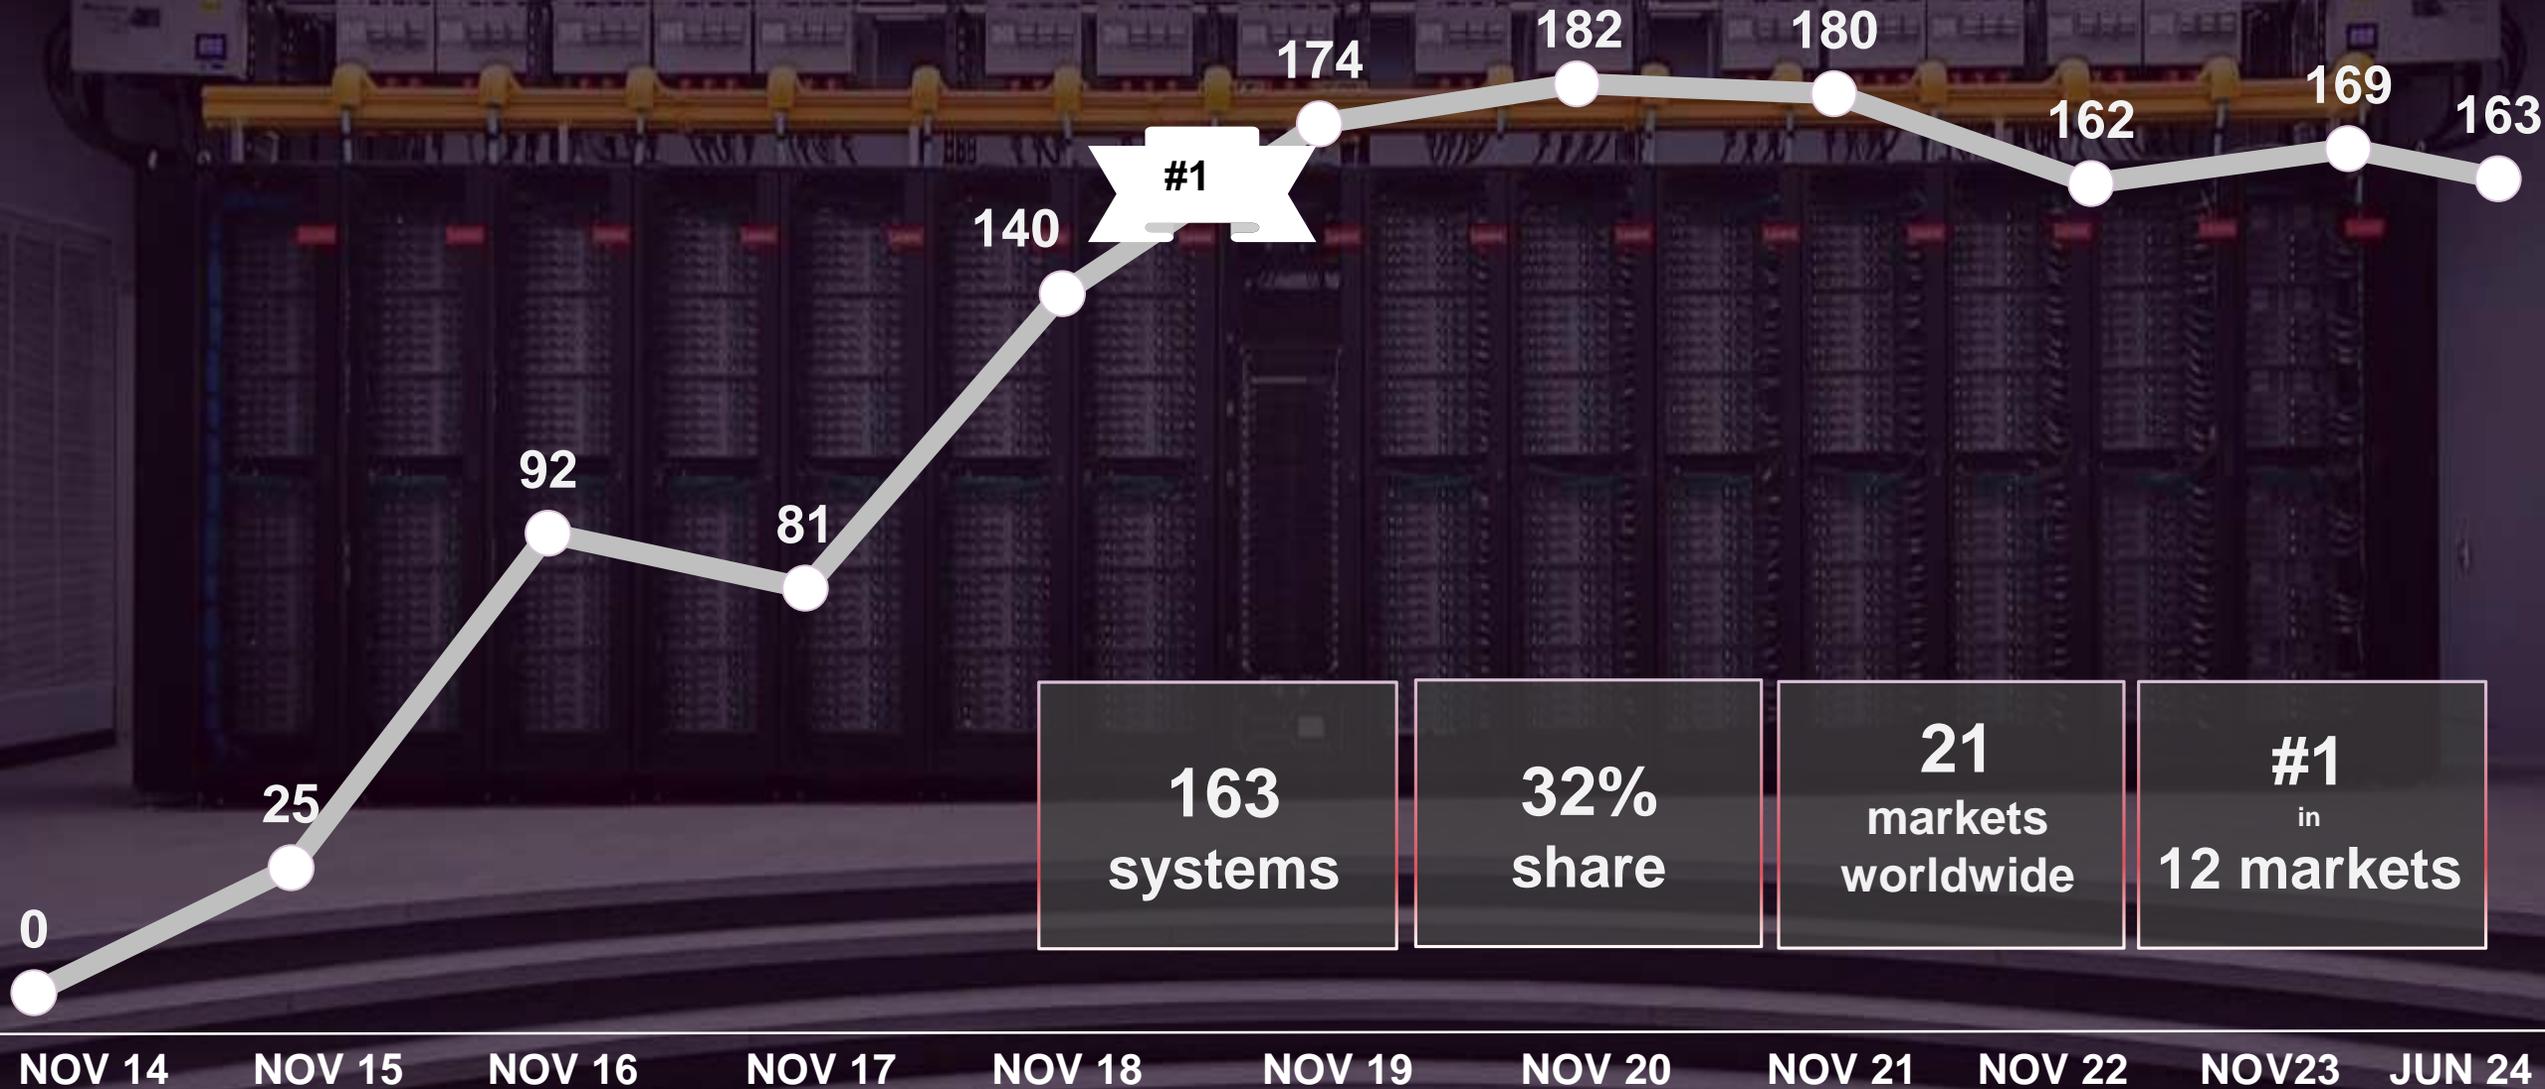
World records set by servers with 3rd Generation Intel® Xeon® Scalable Processors

85

World records set by servers using 1st and 2nd Generation Intel® Xeon® Scalable Processors

#1 Provider of Supercomputers in the world

1 in 3 run on Lenovo HPC



Based on count of listing on the TOP500.org June 2024 list.

Lenovo ThinkSystem Server Portfolio

Comprehensive compute capabilities from the edge to the enterprise

ThinkSystem ST series

Powered by Intel®



ST650 V3 ST650 V2



ST250 V2 ST50 V2



ST250 V2 ST550

ThinkSystem SR series

1U

2U

Powered by AMD



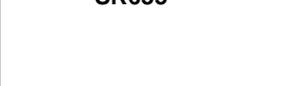
SR645 V3



SR635 V3



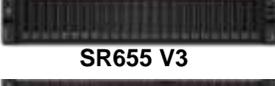
SR645



SR635



SR665 V3



SR655 V3



SR665



SR655

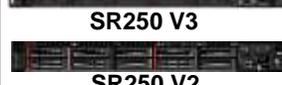
Powered by Intel®



SR630 V3



SR630 V2



SR250 V3



SR250 V2



SR250



SR650 V3



SR650 V2



SR650 V2



HS350X V3

ThinkSystem Modular

SD series (Multi-Node)

Powered by Intel®



SD530
2U 4N



SD630 V2
2U 4N



SD550 V3
2S2U



SD530 V3
2S1U

Powered by AMD



SD535 V3
1S1U

ThinkEdge



SE10



SE350



SE450



SE350 V2



SE360 V2



SE455 V3



SE30



SE50



SE70

ThinkSystem Mission Critical

Powered by Intel®



SR950 V3
8 CPUs



SR860 V3
4U up to 4 CPUs



SR850 V3
2U up to 4 CPUs



SR950
4U up to 8 CPUs



SR860 V2
4U up to 4 CPUs



SR850 V2
2U up to 4 CPUs

HR & MGX Systems

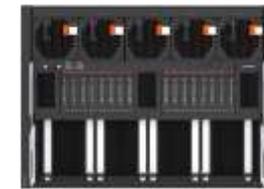
HR650
HR630
HG630N
HG650N
HG660X V3
HR650N



ThinkSystem Segment Optimized

8GPU AI Servers

AMD & Intel®



SR680a V3



SR685a V3



SR780a V3

HPC/AI Optimized

AMD & Intel®

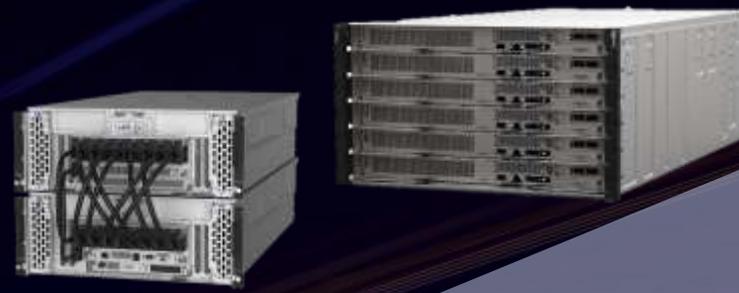


SR670 V2 (GPU Rich),
SR675 V3 (GPU Rich),
Closed loop liquid option)



SD650 V3, SD650-I V3,
SD665 V3,
Liquid Cooled:
SD650-N V3 SD665-N V3

New Lenovo ThinkSystem V3 Servers deliver smarter infrastructure solutions ideal for mission-critical, HPC, and cloud-based containerized workloads.



Most Efficient Servers

Fastest SAP Platform



World's Highest Performance System



World's Most Reliable X86

World's Most Secure X86

60%

More VDI Users

2x

VM per System

3.2x

Performance Increase Per Rack

2x

AI Modeling

1S/2S Rack Servers



Flexible

- Open architecture servers
- Highly versatile storage
- Multiple GPU support
- Wide selection of I/O cards
- As-a service options

Reliable

- Industry's most-reliable x86 server for 10 years in a row*
- Easy to deploy, update and service

Secure

- #1 Ranking for security**
- First Tier 1 manufacturer to offer Intel Transparent Supply Chain (TSC)***

1S/2S Tower Servers



Flexible

- Open architecture servers
- Highly versatile storage
- Flexible support with PCIe Gen4 and Gen5 adapters
- As-a service options

Reliable

- Industry's most-reliable x86 server for 10 years in a row*
- ST650 V2 rated by CRN as Product of the Year for 2022**
- Easy to deploy, update and service

Secure

- #1 Ranking for security***
- First Tier 1 manufacturer to offer Intel Transparent Supply Chain (TSC)****

Mission Critical Servers



Performance

- The latest Intel Xeon CPUs
- Ultra-fast DDR5 memory
- Direct NVMe connectivity speeds database response times
- Pre-validated solutions for system stability and reliability

Agility

- Next-gen workload ready - AI/ML, 3D modeling, and GPU intensive
- Versatile storage and I/O configuration options
- Flexibility to manage multiple workloads on one server

Scalable

- Highly scalable, reliable, and secure servers that can be configured with two, four, or eight processors
- Versatile storage and I/O configuration options

High-performance Data Center Infrastructure Solutions optimized for AI

ML and DL Analytics



- SR850 V3 Intel
- SR860 V3 Intel
- SR650 V2 Intel
- SR650 V3 Intel
- SR655 AMD
- SR655 V3 AMD
- SR665 AMD
- SR665 V3 AMD

AI Model Training



- ST650 V2 Intel
- ST650 V3 Intel
- SR670 V2 4x PCIe Intel
- SR670 V2 Intel
- SR670 V2 8x PCIe Intel
- SR670 V2 4-GPU HGX Intel
- SR675 V3 4x PCIe AMD
- SR675 V3 8x PCIe AMD
- SR675 V3 4-GPU HGX AMD

NEW SR680a V3 8 GPU Intel
NEW SR780a V3 5U 8 GPU Intel
NEW SR685a V3 8 GPU AMD

Liquid-Cooling Training



- SD650-N V2 Intel
- SD650-I V3 Intel
- SD665-N V3 AMD

NEW SR780a V3 8 GPU

Lenovo leverages best-in-class technologies

ThinkSystem NVIDIA H100, H200, L40S GPUs

NVIDIA NVLink

AMD Instinct™ MI210 and MI300X GPU

Intel GPUs

AI-optimized Servers for Data Centers

- ✓ Introducing two 8U, 2-socket servers with 8 GPUs and high-speed interconnects
- ✓ Highly accelerated systems delivering massive computational capability to handle data and compute-demanding AI and HPC workloads
- ✓ Offering flexible choices in CPUs (Intel or AMD) and GPUs (AMD, NVIDIA) on the industry-leading ThinkSystem platform – secure, reliable, and performant
- ✓ Uses industry-standard 19-inch server racks
- ✓ Includes XClarity Systems Management plus a choice of AI Software Stack
- ✓ Supported by world-class engineering, industry-leading supply chain, in-house manufacturing, and worldwide services

ThinkSystem SR680a V3
ThinkSystem SR685a V3



Ideal for Model development, training, and retraining, machine learning, deep learning, LLMs, HPC simulations

Introducing ThinkSystem SR680a V3

An 8U, two-socket system optimized for massive computational performance with Intel processors and choice of GPUs



Ideal for both AI and HPC

- AI Workloads: Model development, training and retraining, Machine Learning, Deep Learning, LLMs
- HPC workloads: Modeling and simulation, rendering, financial tech, scientific research

Performance

- Supports 5th generation of Intel® Xeon® Scalable Processors
- Maximum GPU Performance with NVIDIA, and AMD's latest GPUs & accelerators
- Fastest transfer rate with NVIDIA NVLink and AMD's Infinity Fabric

Flexible

- Support for AMD MI300X, NVIDIA H100/H200 and future GPUs
- Sufficient Thermal Headroom for future higher-power GPUs

Ease of Deployment

- Lenovo XClarity System Management software improves the efficiency of your infrastructure deployment
- Fits into industry-standard 19-inch racks

Introducing ThinkSystem SR685a V3

A 2-socket, 8U system workload optimized for performance with AMD processors and GPUs



Ideal for both AI and HPC

- Industries: Financial services, healthcare, energy, climate science, and transportation
- AI Workloads: Model development, training and retraining, Machine Learning, Deep Learning, LLMs
- HPC workloads: Modeling and simulation, rendering, financial tech, scientific research

Performance

- Industry-leading 4th Gen AMD EPYC™ processors
- Maximum GPU Performance w/ AMD Instinct™ MI300x
- Interconnects for an ultra-fast transfer rate with AMD Infinity Fabric™ at 800GB/sec

Flexible

- Support for AMD MI300X, NVIDIA H100 / H200/B100 GPUs
- Drop-in upgrade to AMD next-gen processor
- Easy to manage and simplified deployment with Lenovo XClarity System Management Software
- Fits into industry-standard 19-inch racks

Advanced Thermals

- Designed with additional thermal headroom to accommodate future, high powered, GPUs
- Advanced Air-Cooled Design

Lenovo Liquid Cooling Technologies

A leader in innovative technology pioneering liquid cooling for the enterprise

Direct Water Cooling

- Warm water cooling inside the server
- Higher performance than air-cooled servers
- Up to 40% reduction in power costs

Smart Control Software

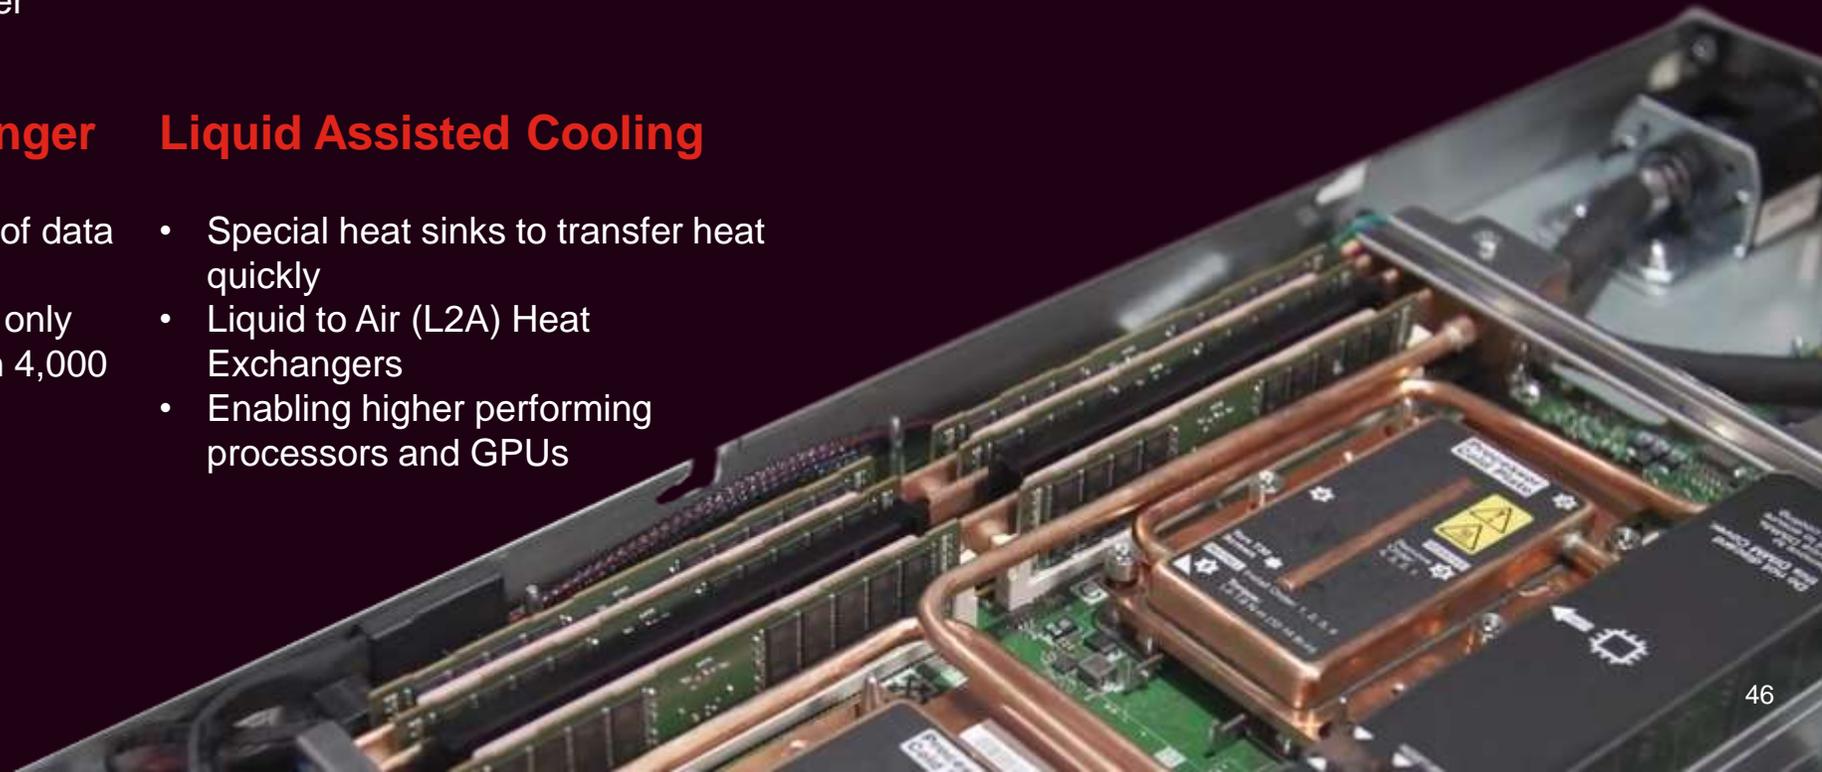
- Minimizes power consumption for each workload using software

Rear Door Heat Exchanger

- Rack option takes heat out of data center in chilled water
- 3.5X more efficient than air only
- Saves enough power to run 4,000 LED bulbs

Liquid Assisted Cooling

- Special heat sinks to transfer heat quickly
- Liquid to Air (L2A) Heat Exchangers
- Enabling higher performing processors and GPUs



Sustainable Data Center Solutions

Lenovo Energy- Efficient Technologies

Energy Management Tooling *Measure, manage, optimize power*

xClarity Energy Manager (LXEM)



Energy Aware Runtime (EAR)



Efficient Component Selection *Same work, less power*

Titanium Power Supplies



Idle Power State Controls

Dynamically
optimize the
frequency and
power control

Low Loss Materials



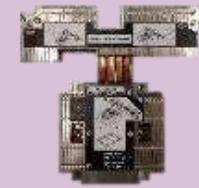
Liquid-Cooling with Lenovo Neptune™

Heat Mitigation Innovation *Lower fan speed, less power*

Liquid to Air Exchanger (L2A)



Thermal Transfer Module (TTM)



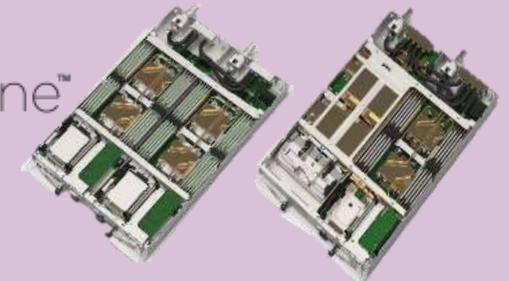
Liquid Cooling

Rethink how cooling gets done

Rear Door Heat Exchanger



Direct Liquid Cooled Systems



Neptune™

New Compute Intensive Portfolio

ThinkSystem SD520 V4



The ThinkSystem SD520 V4 is a 1U, half-width server powered by Intel® Xeon® 6 maximizes rack density and efficiency for compute-intense, transactional data.

- 1x Intel® Xeon® 6 processor with up to 144 cores
- Up to 1x CPU up to 350W
- Up to 6x 2.5" or 2x E3.S 1TB drives - max storage of 92TB
- Supports 1x 75W GPU
- Up to 8x TruDDR5 RDIMMs up to 1TB
- Supports CXL 2.0 memory expansion
- 2x Gen 5 slots (1x PCIe and 1x OCP 3.0)
- SW RAID (VROC) with both Simple Swap & Hot Swap configs
- Optional Neptune® open-loop liquid-cooling

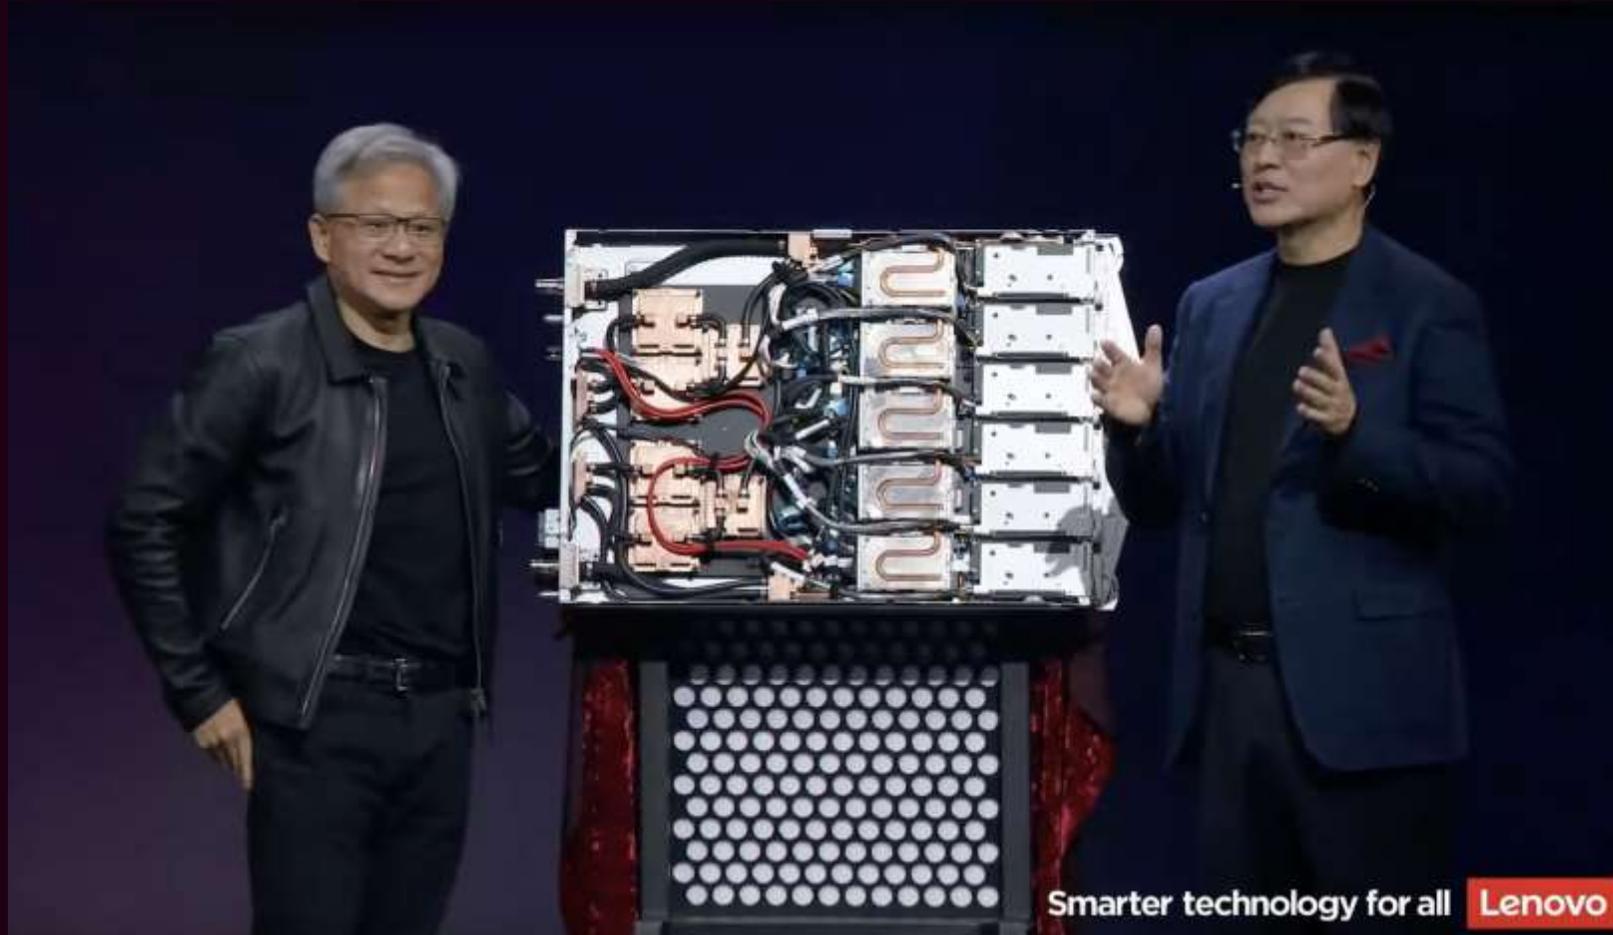
ThinkSystem SR630 V4



The SR630 V4 with the Intel® Xeon® 6 processors provides speed, reliability, and efficiency to maximize performance for versatile workloads.

- Up to 2x Intel® Xeon® 6 processors up to 288 cores
- Up to 32x TruDDR5 6400MHz memory up to 8TB
- Up to 12x low latency NVMe SSD high performance storage - 737TB max
- Support 16x EDSFF E3.S NVMe drives
- Supports CXL 2.0 memory expansion
- Maximum 5x PCIe Gen5 slots + 2x OCP 3.0 for I/O flexibility
- Up to 3x 75W GPUs for graphics-intensive workloads
- Optional Neptune® liquid-to-air or open-loop cooling technology

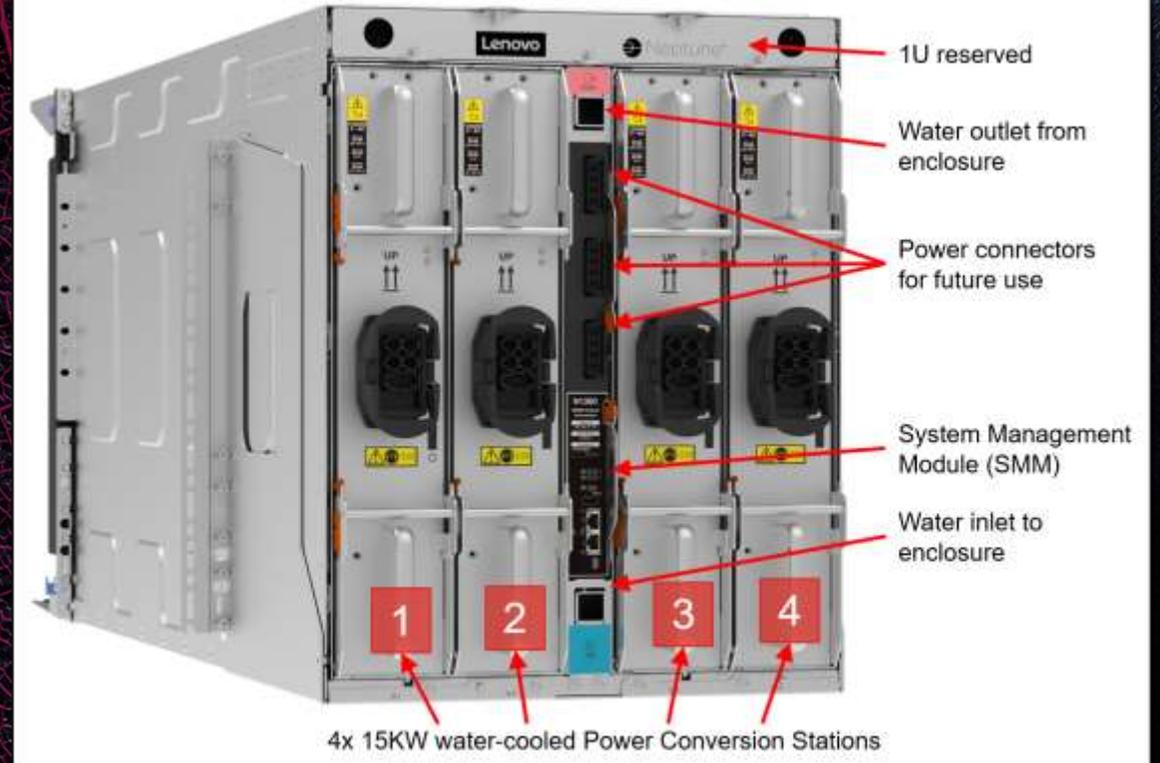
Lenovo Tech World'24 - ThinkSystem SC777 with NVIDIA Blackwell



ThinkSystem SC750 V4 Neptune

ThinkSystem N1380 Neptune

8 trays (16 servers)
in a 13U enclosure:
32 processors



Lenovo ThinkAgile HCI Portfolio

Engineered and validated designs for confident deployment



ThinkAgile V3 MX and SX Series



- Integrated Appliance with Premier Support
- Integrated Microsoft or Azure Stack with Arc
- Hybrid cloud for Azure with VDI and Microsoft SQL optimization
- Rack scale solutions



ThinkAgile V3 HX Series



- Integrated Appliance with Premier Support
- Integrated Nutanix Stack
- Datacenter and Edge solutions
- VDI Solutions, SAP HANA, Oracle, and SQL



by Broadcom

ThinkAgile V3 VX Series



- Integrated Appliance with Premier Support
- Integrated VMware Stack
- Datacenter and Edge solutions
- VDI Solutions, SAP HANA, Kubernetes, and VCF

Lenovo Management & Automation Software



Centralized systems management suite to monitor, manage and automate Lenovo's infrastructure



An automated platform to deploy, optimize and manage cloud infrastructure (bare metal servers, containers, VMs).



Lenovo intelligent computing orchestration (LiCO) brings several open-source components into a single GUI to make the access of computing resources easier for AI development and training and HPC cluster management

XClarity Software Suite

Cloud automation and IT service management processes

4

Lenovo

XClarity Integrators

- Integration into leading virtualization management consoles and IT service management tools
- VMware, Microsoft, Cloudforms, Chef, Puppet, ServiceNow, MSFT WAC

Innovative Centralized Management

3

Lenovo

XClarity Orchestrator

Lenovo

XClarity Administrator

- Centralized software-based delivery and management for ThinkSystem servers, storage, and networking
- Mobile app for anywhere management
- REST APIs for ease of integration into software-defined environments

Platform Management Tools

2

Lenovo

XClarity Essentials

- Collection of one-to-one management tools
- Scripting tools
- Bootable Media Creator

Embedded Platform Management

1

Lenovo

XClarity Controller

- Embedded management engine common in all ThinkSystem servers
- Fresh, uncluttered graphical user interface
- Redfish-compliant web-based REST APIs for ease of inter-operability

VMware by Broadcom Updates – New Bundles

Previous Products

vSphere Enterprise Plus

vSphere Enterprise

vSphere for Desktop

vSphere Scale-Out

vSphere

vSAN

NSX

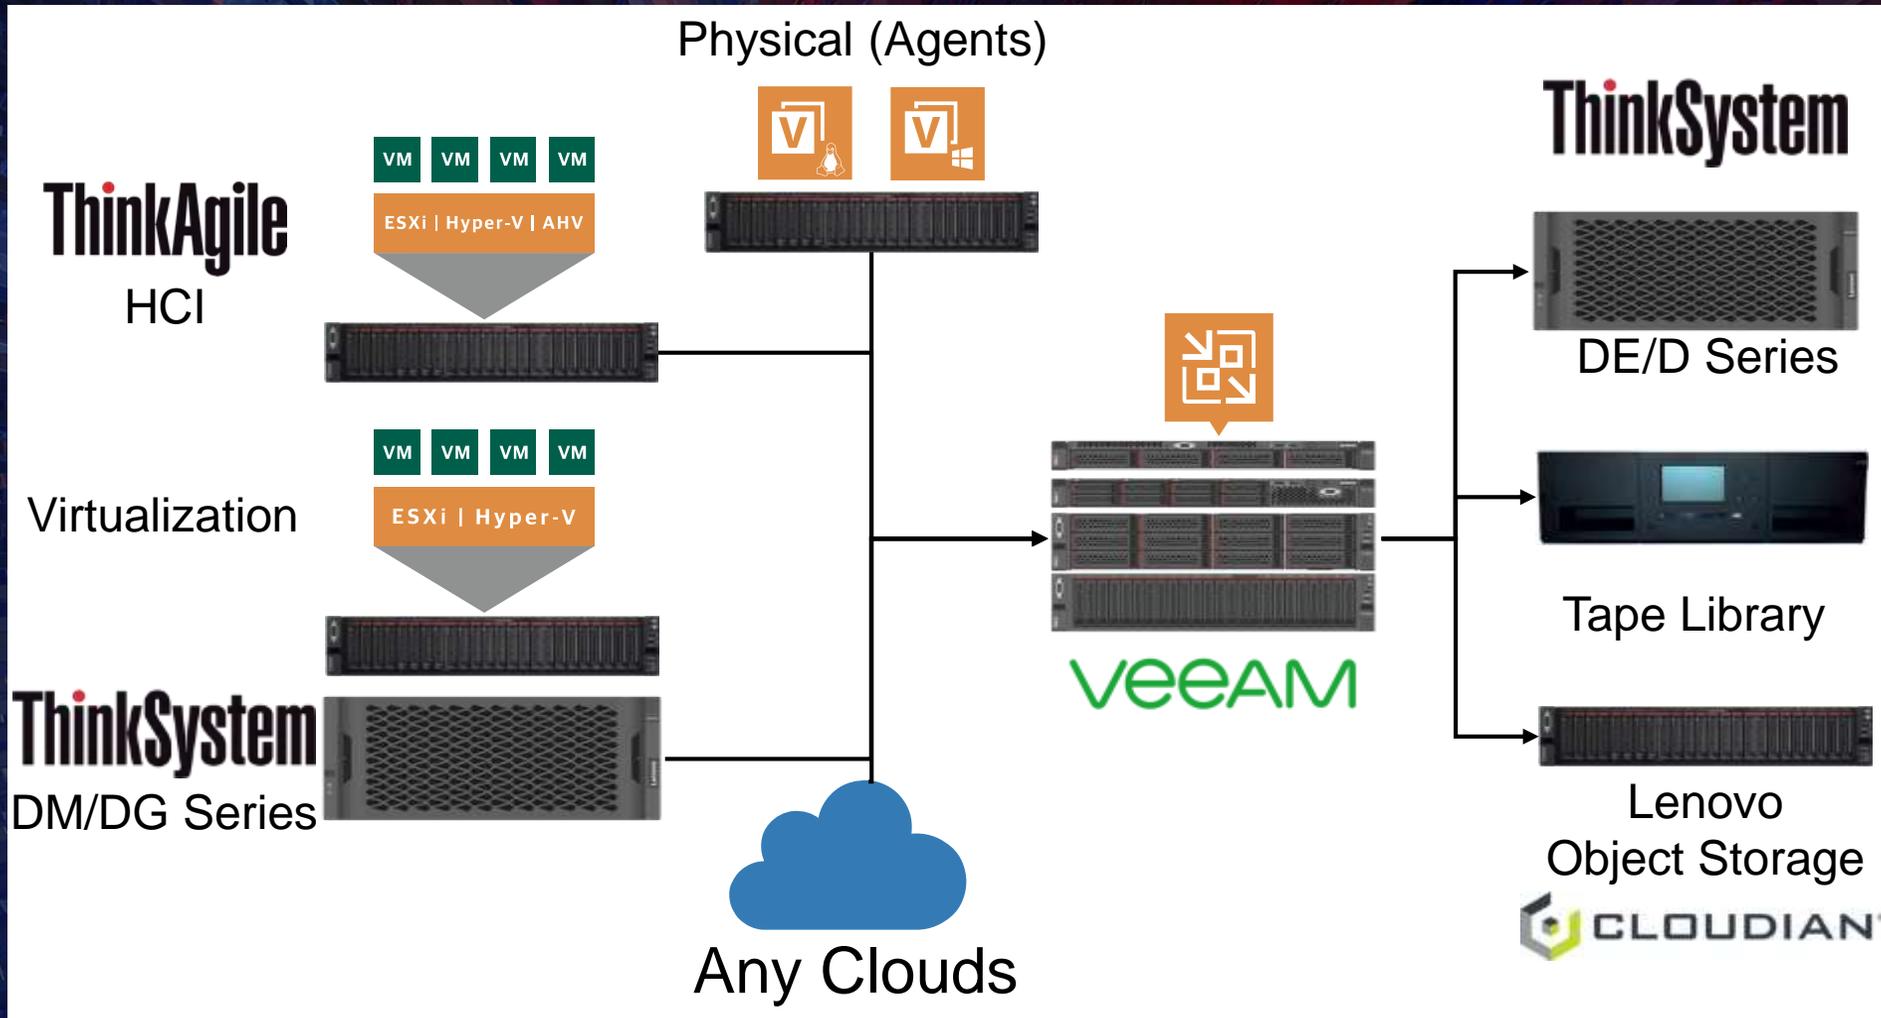
Aria

Recommendation

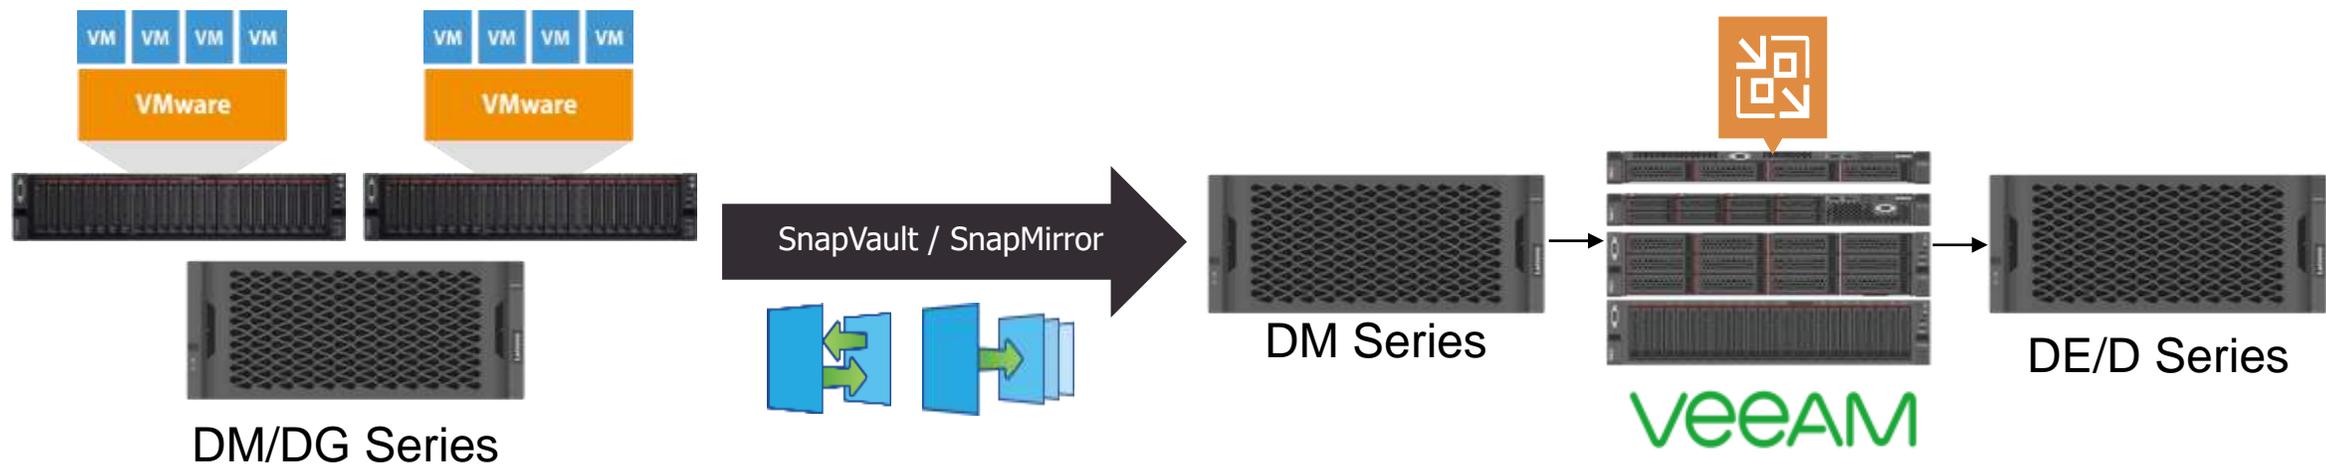
vSphere Foundation &
vSAN Add-on*

VMware Cloud
Foundation &
Firewall*

Lenovo & Veeam - Typical System Design



Lenovo & Veeam – DM/DG Series Integration



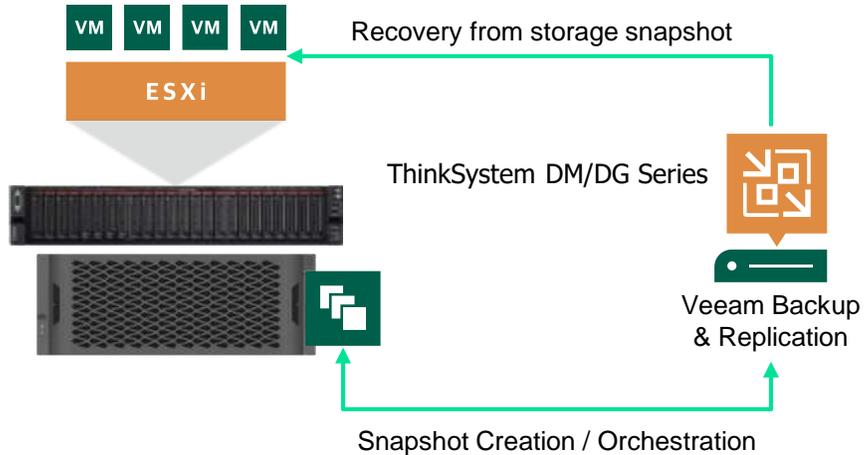
Faster backups: Back up VMware virtual machines (VMs) directly from Lenovo storage snapshot copies.

Quick recovery: Recover entire VMs or individual items from Lenovo storage snapshot, and SnapMirror software.

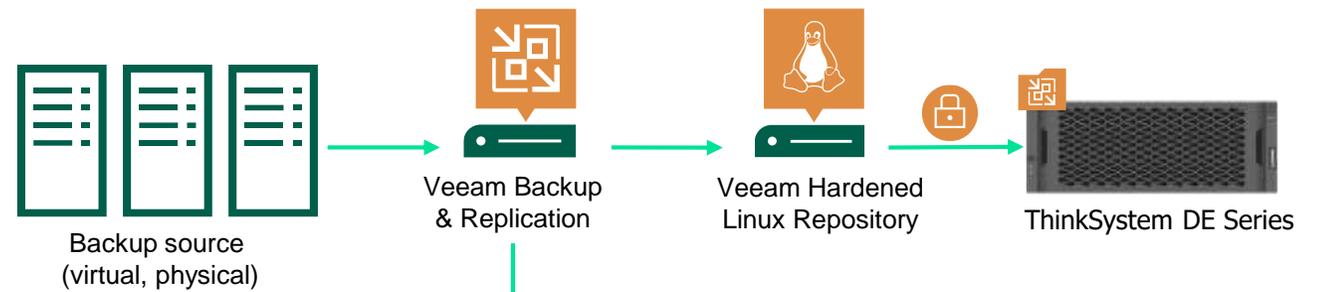
Improved protection: Create instant secondary backups from Lenovo storage snapshot copies.

Lenovo & Veeam – Ransomware Protection

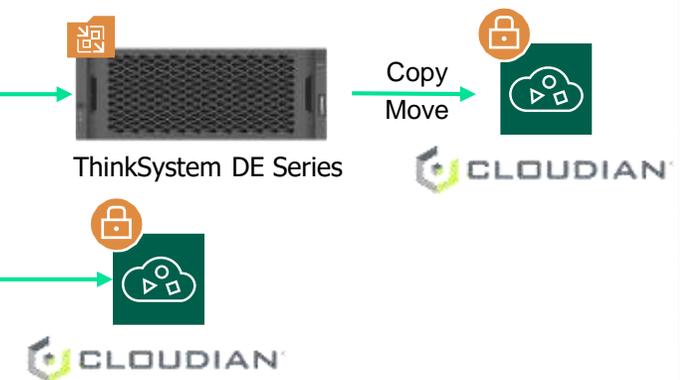
Immutability Lenovo DM Series SnapLock Snapshots



Immutability Lenovo DE Series or DAS Hardened Linux Repository



Immutability Object Storage Cloudian



Leveraging AI to diagnose cancer up to 10x faster

Ministry of Healthcare of the Republic of Kazakhstan

To help accelerate breast cancer screening at 120 hospitals in Kazakhstan, the Ministry of Healthcare deployed a hyperconverged infrastructure solution powered by 14 Lenovo ThinkAgile HX5520 nodes—enabling rapid, AI-powered processing of radiological images.

Customer Problem

Cancer is one of the leading causes of mortality in Kazakhstan, but a shortage of radiologists made it difficult for hospitals to screen for the disease effectively.

Lenovo Solution

The Ministry of Healthcare deployed an HCI solution powered by 14 Lenovo ThinkAgile HX5520 nodes, enabling it to create a central repository to store and analyze radiological images.

Business Impact

The new solution accelerates screening programs by a factor of three, enabling 10x faster breast cancer diagnosis and a 35% increase in early detection.



Powered by Intel®



“Our achievements include a 375% increase in early detection of pathologies and a 35% increase in early detection of cancer, resulting in improved quality of life for patients and increased treatment effectiveness.”

Spokesperson

Ministry of Healthcare of the Republic of Kazakhstan



3 million radiological images analyzed



10x faster breast cancer diagnosis



35% increase in early cancer detection



Smarter
technology
for all

Lenovo

thanks.