

A Global Company with a Strong European Presence

[ANY SCALE] [ANY DATA CENTER] [ANY CLOUD]

DDN Recognized as #1 in the U.S. for Gartner® Use Case: AI in the Gartner Critical Capabilities for Enterprise Storage Platforms

DDN is proud to share that DDN ranks #1 in the U.S. for Gartner Use Case artificial intelligence (AI) in its Critical



Jensen Huang

CEO and co-founder NVIDIA

Alex Bouzari,

CEO and co-founder

DDN

"NVIDIA is powered by DDN ... Without DDN, NVIDIA supercomputers wouldn't be possible."

NVIDIA Reference Supercomputer: single tier full flash EXAScaler



32 NVIDIA GB200 NVL72

2304 B200 GPUS

1 RACK OF DDN STORAGE

16 DDN DATA APPLIANCES: AI400X3

3 DDN METADATA APPLIANCES: AI400X3

A Presence at Every Stages of the Industrial Cycle


In Europe DDN is active at every stages of the Industrial cycle

- **Knowledge generation**
 - Mentorship, training, teaching and funding Ph.Ds
- **Research collaboration**
 - Participation to EU research programs
 - In-kind contributions to EU projects
 - Seed funding for partners research programs
- **Product definition**
 - Head of Product and the CTO based Europe
 - R&D team in France
 - Open-Source development + Industrial partnership
- **Customer feed-back**
 - Joint lab with strategic European partners



An Open-Source Driven Company

Lustre Contributions by Release




Lustre Open-Source Parallel File system (OpenSFS)

- Designed for HPC: data extension of the compute platform
- OpenSFS provides overall directions and a forum for discussion among users
- DDN is the lead contributor to Lustre
- User meetings in Europe / ASIA co-organized by DDN

Data courtesy of Dustin Leverman (ORNL)

Source: <https://git.whamcloud.com> 2.15.50 to 2.15.91



DDN European R&D 15% of WW headcount

- Less QA, more than 15% of R&D is current done in Europe (France, Germany)
 - Significant investments to be announced
- Quality of Service – 100% in Europe
- Multi-tenancy – 100% in Europe
- Linux Kernel IO Optimization – 100% in Europe (open source)
- EXAScaler monitoring – Europe + WW
- Client Compression – Europe + WW
- CCI in Lustre – under assessment

DDN, EuroHPC & Europe – 3 pillars of development

- **A – EuroHPC Infrastructures powered by DDN**



Leonardo



Melluxina



Discoverer



Vega



Deucalion



- **B – EuroHPC Research Programs and DDN**

- Currently Open - Admire, DaFab, EoCoE-3
- Recently Closed - EVOLVE, EoCoe, Esiwace
- EuroHPC design of next-Gen IO system
- AI-automated features extractions from Satellite Images
- Nuclear Fusion code optimization

- **C – DDN R&D Spending in Europe**

- A significant portion of our WW turnover is spent on R&D
- France focused on Software Platforms



ABOVE THE STACK

DATA PLATFORMS VISION

DDN's Data Intelligence Platform is Powering the World's AI Transformation



- ✓ Maximum Efficiency AI Training at Any Scale: Extreme throughput file platform




- ✓ Edge-Cloud-Core, Inference and Training, Model and Data Governance


DDN DATA INTELLIGENCE

EXAScaler: Industrialization of Lustre with Specific Features


GPU Direct Storage (GDS) enables a direct data path for direct memory access (DMA) transfers between GPU memory and external storage bypassing CPU. GDS Increases throughput, reduces latency, eliminates memory copies, bypass internal architecture bottlenecks.




Hot Nodes converts local SSDs in application servers into read-only cache to improve re-read performance. AI applications benefit significantly by avoiding network hops to fetch data. Cached files persist through client node downtime and are available again once the client node reboots. Files get detached from cache automatically if modified.



Multi-rail network enable aggregation of bandwidth of multiple network interfaces into one single fat pipe. It simplifies network configuration, increases bandwidth and increases resilience, by allowing alternative transmission paths.




CONTAINER SUPPORT




EXAScaler supports all popular container frameworks including Docker, Kubernetes with Container Storage Interface (CSI) with automated mounting of filesystem without needing root access.

MULTI TENANCY



EXAScaler supports tenant isolation via subdirectory mount & node-maps to identify client nodes/VM/container used by a tenant. It also supports VLAN based isolation with LNET Routers.

MULTI PROTOCOL ACCESS



Along with default access via Lustre native clients, EXAScaler also supports NFS, SMB, S3 and CSI interfaces for data sharing. Ingest via one protocol and egest via another is supported. Commodity servers are required to run the gateway protocol.

exab

DDN Infinia: A unified, cloud-optimized platform for high-performance data and AI-driven insights at scale.

Unlock the power of structured, semi-structured and unstructured data with seamless access, management, and analytics across edge, core and cloud.



**Structured AND
Unstructured data**



**Storage Delivered
as a Service**



**Hyper Simple
Operation**

DDN Cloudification

ORACLE Cloud Marketplace

Oracle Cloud Home Publishers ▾ Resources ▾ Sign In English ▾



DDN Infinia

Get App >

The Data Intelligence Platform for AI & Analytics

Oracle Cloud Infrastructure

Big Data , Data & Analytics , Artificial Intelligence

💰 Software Price: BYOL ⓘ



Google Cloud

Documentation

Technology areas ▾

Cross-product tools ▾

Related sites ▾

🔍 Search



🌐 English ▾

Console



Managed Lustre

Overview

Guides

Reference

Resources

Contact Us

Start free

Filter

Discover

Product overview

Supported locations

Get started

Access control with IAM

Enable the Managed Lustre API

Configure a VPC network

Create and manage instances

Create a Managed Lustre instance

Increase storage capacity

Manage instances

Storage > Managed Lustre > Guides

Was this helpful? 👍 🗨

Managed Lustre

Send feedback

Google Cloud Managed Lustre delivers a high-performance, fully managed parallel file system optimized for AI and HPC applications. With multi-petabyte-scale capacity and up to 1 TBps throughput, Managed Lustre facilitates the migration of demanding workloads to the cloud.

Highlights

- Developed in collaboration with DDN, the leading provider and primary maintainer of the Lustre file system.
- Supports extreme IOPs and low latency.

On this page

Highlights

Specifications

Use Cases

Pricing

AI Factories: NVIDIA + DDN: End-to-End AI Solution

TRAINING

- ✓ Maximum GPU Productivity
- ✓ NVIDIA DGX Platforms



DIGITAL TWIN

- ✓ Omniverse with Cosmos
- ✓ Accelerate NIMs and Nemo

EDGE

- ✓ Edge-Cloud-Core
- ✓ Collection, Curation, Transfer



THANKS!