



The World's Leading  
Data Intelligence Platform

[ ANY DATA ] [ ANY APPLICATION ] [ ANYWHERE ]

© 2025 DDN

DEPRECIATION  
**SUNK COST**

**ACQUISITION COST**

**TCO**

DATA CENTER COSTS

FACILITY COSTS

**CAPEX**

**HIDDEN COST**

SOFTWARE LICENSES

INFRASTRUCTURE COST

**OPEX**

INTEREST COST

NETWORKING COSTS

SCALING COSTS

**MAINTENANCE COSTS**

POWER USAGE EFFECTIVENESS (PUE)

OPPORTUNITY COST

DEPLOYMENT COSTS

SUPPORT & SERVICES COSTS

**ROI**



COOLING COSTS



HARDWARE COSTS

INSTALLATION RATE

**RESALE VALUE**

RENEWAL COSTS

**RACK SPACE COSTS**

\$ FINANCING COSTS

INSTALLATION & SETUP COSTS

TRAINING INEFFICIENCY

VENDOR LOCK-IN 

**E-WASTE DISPOSAL COSTS**

CONFIGURATION COSTS

COMPLIANCE COSTS

PERFORMANCE COST PER TOKEN

**SLAS**

CARBON FOOTPRINT COSTS 

IDLE CAPACITY



DEPRECIATION  
**SUNK COST**

**ACQUISITION COST**

**TCO**

DATA CENTER COSTS

FACILITY COSTS

**CAPEX**

**HIDDEN COST**

SOFTWARE LICENSES

INFRASTRUCTURE COST

**OPEX**

INTEREST COST

NETWORKING COSTS

SCALING COSTS

**MAINTENANCE COSTS**

POWER USAGE EFFECTIVENESS (PUE)

OPPORTUNITY COST

DEPLOYMENT COSTS

SUPPORT & SERVICES COSTS

**ROI**



COOLING COSTS



HARDWARE COSTS

INSTALLATION RATE

**RESALE VALUE**

RENEWAL COSTS

**RACK SPACE COSTS**

\$ FINANCING COSTS

INSTALLATION & SETUP COSTS

TRAINING INEFFICIENCY

VENDOR LOCK-IN 

**E-WASTE DISPOSAL COSTS**

CONFIGURATION COSTS

COMPLIANCE COSTS

PERFORMANCE COST PER TOKEN

**SLAS**

CARBON FOOTPRINT COSTS 

IDLE CAPACITY



# Tokenomics:

## The Economics of AI at Scale

# What is a Token?

# Why Tokenomics Matters

Every AI interaction has costs:

- GPU compute
- Storage throughput
- Data movement
- Networking
- Power/cooling
- Time-to-results

# Why Tokenomics Matters

## Cost per useful token

Not just generating tokens... but generating them:

- faster
- cheaper
- more accurately
- at scale

# DDN is to Data What NVIDIA is to Compute

## COMPUTE



Full-Stack Provider of **HPC & AI COMPUTE** Platforms in Data Centers and Multi-Clouds

## DATA



Full-Stack Provider of **HPC & AI DATA** Platforms in Data Centers and Multi-Clouds

## You Cannot Succeed in AI Without **DDN**



“NVIDIA Is  
**Powered By DDN**”

...Without DDN, NVIDIA  
Supercomputers  
**Wouldn't Be Possible.**”

Jensen Huang, NVIDIA CEO



*Jensen Huang, CEO of Nvidia, and Alex Bouzari, CEO of DDN,  
at DDN Beyond Artificial, February 20, 2025*

# DDN Powers More HPC & AI Than Anyone Else

ACROSS 35 INDUSTRIES

OVER 1 EXABYTE  
DEPLOYED AT SINGLE SITE

11X FASTER MODEL TRAINING

OPERATING IN OVER 60 COUNTRIES

75X QUICKER DATA DISCOVERY

SUPERCHARGING OVER 1 MILLION GPUS

10X REDUCTION IN GPU IDLE TIME


11,000+ CUSTOMERS

¼ THE DATA CENTER FOOTPRINT



PERFORMANCE AT ANY SCALE

# DDN Delivers the World's Most Performing Storage At Any Scale



## Best IOPS & Throughput per rack

- ✓ Up to 80M IOPS in a single rack
- ✓ 3.0 TB/s of Read throughput
- ✓ 1.9 TB/s of Write throughput
- ✓ 20PB of Flash

## Best Price per Flash TB per rack

- ✓ Up to 14M IOPS in a single rack
- ✓ 360 GB/s of read throughput
- ✓ 240 GB/s or write throughput
- ✓ Up to 26PB of Flash


## Best Price per TB

- ✓ 20PB per Rack
- ✓ 90 GB/s of read throughput
- ✓ 65 GB/s of write throughput

## FLEXIBLE TIERING

# DDN OFFERS ALL MEDIA SO YOU CAN CHOOSE CAPACITY-PERFORMANCE-COST

- ✓ Flexible Tiered Storage offerings from Super-Fast through Archive
- ✓ All easily managed through a single SW interface
- ✓ All accessible through a single mount







**THE AI DATA COMPANY**

Wyatt DiVall | Account Executive

Mobile: 713.499.9870

Email: [wdivall@ddn.com](mailto:wdivall@ddn.com)