

# **Tackling the ExaScale Data Challenge**

19th Workshop on high  
performance  
computing in  
meteorology

# EXAScale Challenges

The Next Generation isn't just about Brute Force Performance and Scale...

## Traditional View

- Managing very large scale of devices
- Single Application IO Volume
- Relaxed POSIX (and KV)
- Energy/Power consumption

## New View

- All those... PLUS
- 100% Software – No Hardware Dependencies
- Superior Data Protection
- Native Multi-Tenanted Security
- Ability to run MultiCloud
- Strong Support for IPU's and GPU's
- Intelligence around Data and System Management
- Elasticity, Online Everything





# DDN Intelligent Infrastructure for the **AI ERA**

- 100% Software, 100% portable
- Simple to install, simple to Manage, simple to Grow
- Uniquely Capable to serve all IT and At Scale Workloads with Performance
- Intelligence Built-in

**On Prem**

**In Cloud**

**IOPs**

**Throughput**

**Containers**

**GPUs**

**QoS**

**Tenants**

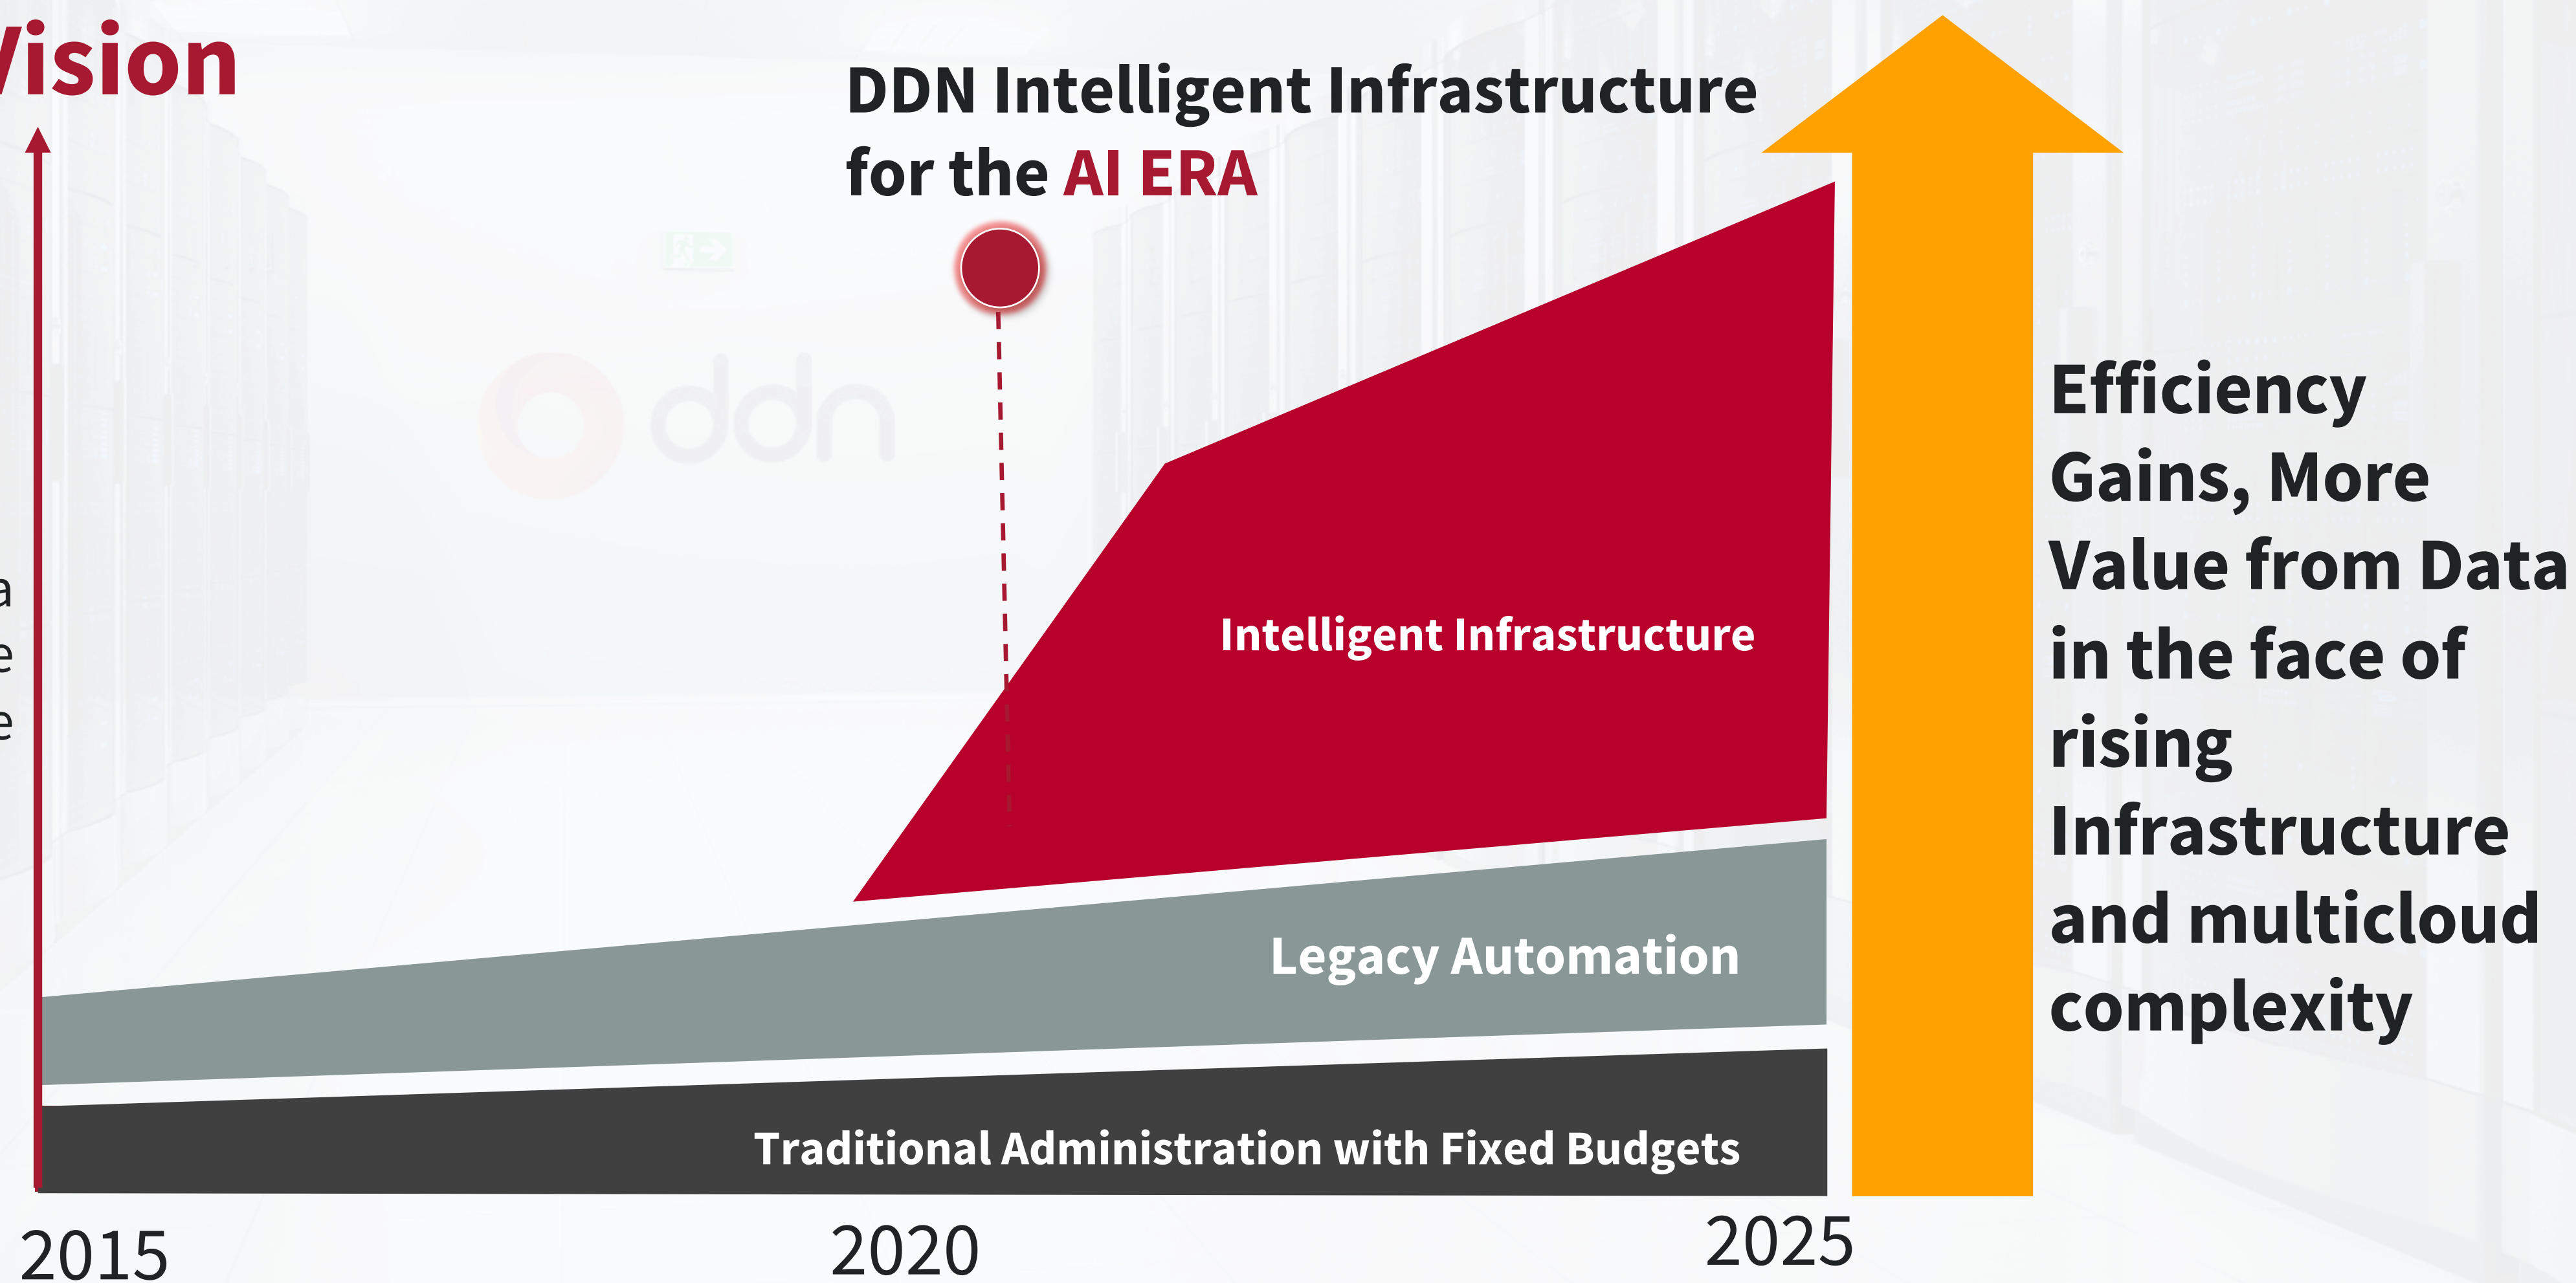
**Scale**

**Elasticity**



## DDN's Vision

\$ Derived from Data  
and Storage  
Infrastructure





To Power the **Greatest Data Challenges We Need to Leverage MetaData**

# Meta Data

Tag, Search, Enrich  
Insights that Matter  
Optimal Automation  
Adaptive Workflows  
On Prem, in the Cloud  
Application Aware

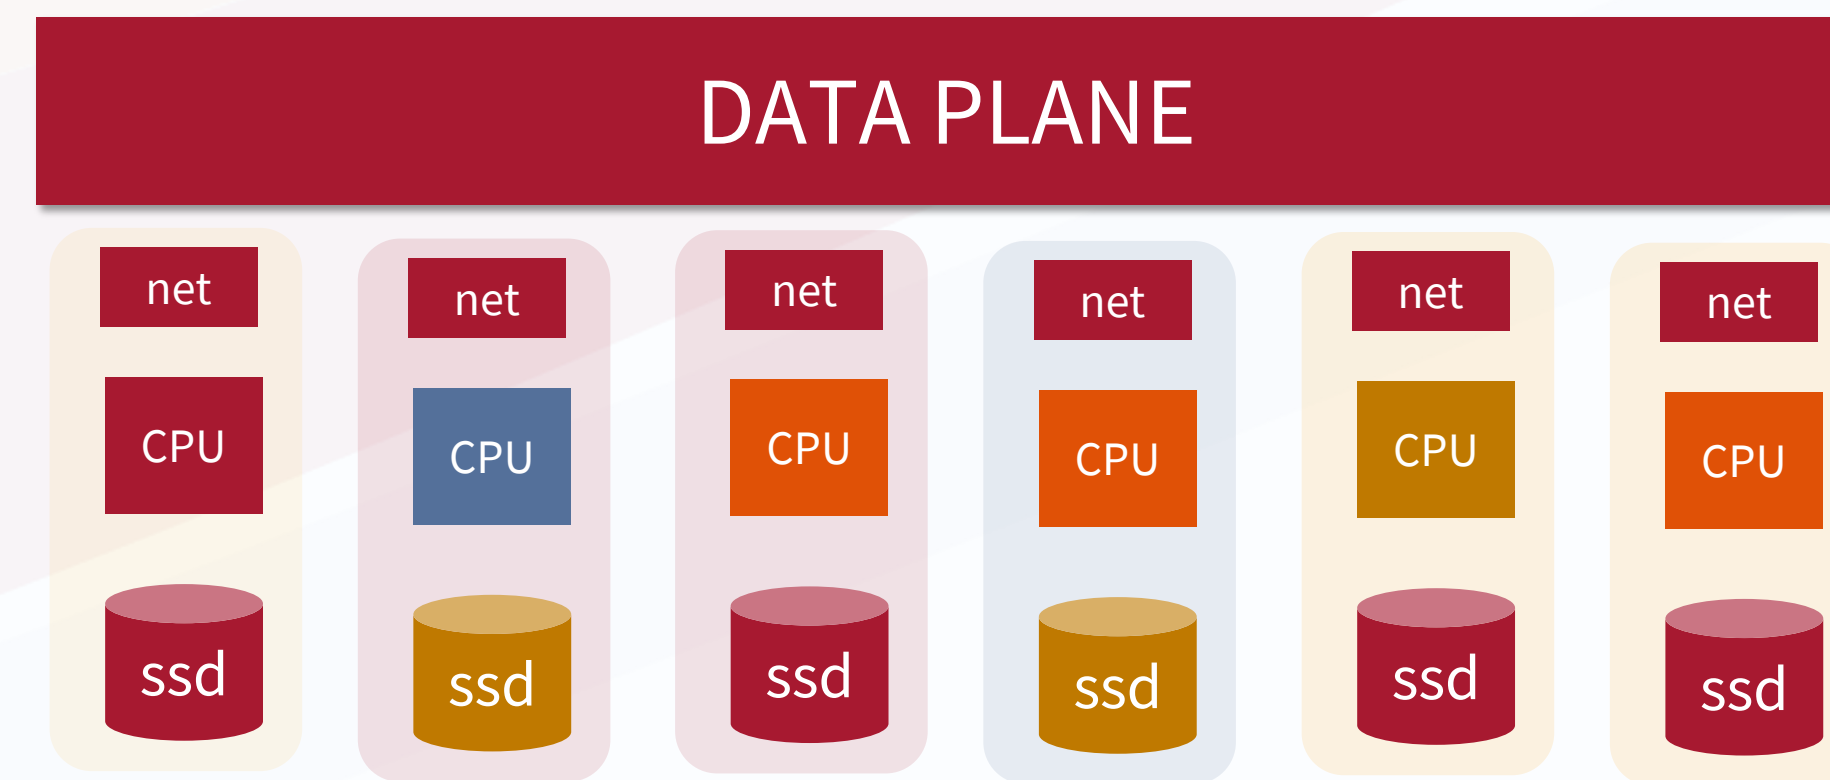
Quality of Service  
Compliance Management  
Security and Audit  
Data Placement  
Performance Consistency  
Endlessly Extensible



# Software Defined, RE-DEFINED

## The Ultimate Software Approach to Storage

- No Strong Hardware Requirements
- Need CPU, network, SSDs-that's it
- Different CPU architectures (ARM/ X86), VMs, Containers



Any CPU, Any Cloud, Any SSD, Any Network



# Native Elasticity

- Expand and Reduce your infrastructure online seamlessly
- Automatically handle Server, media failures and migrations
- Keep your services up to date and performant through generations of server storage





# Easily Match Project Demands with **Dial-In Resources**


**Tenant A**





**Tenant B**



**Tenant C**



**Data Plane**



- Tenants can manage their own policies and bring up their own data services within their given SLA's

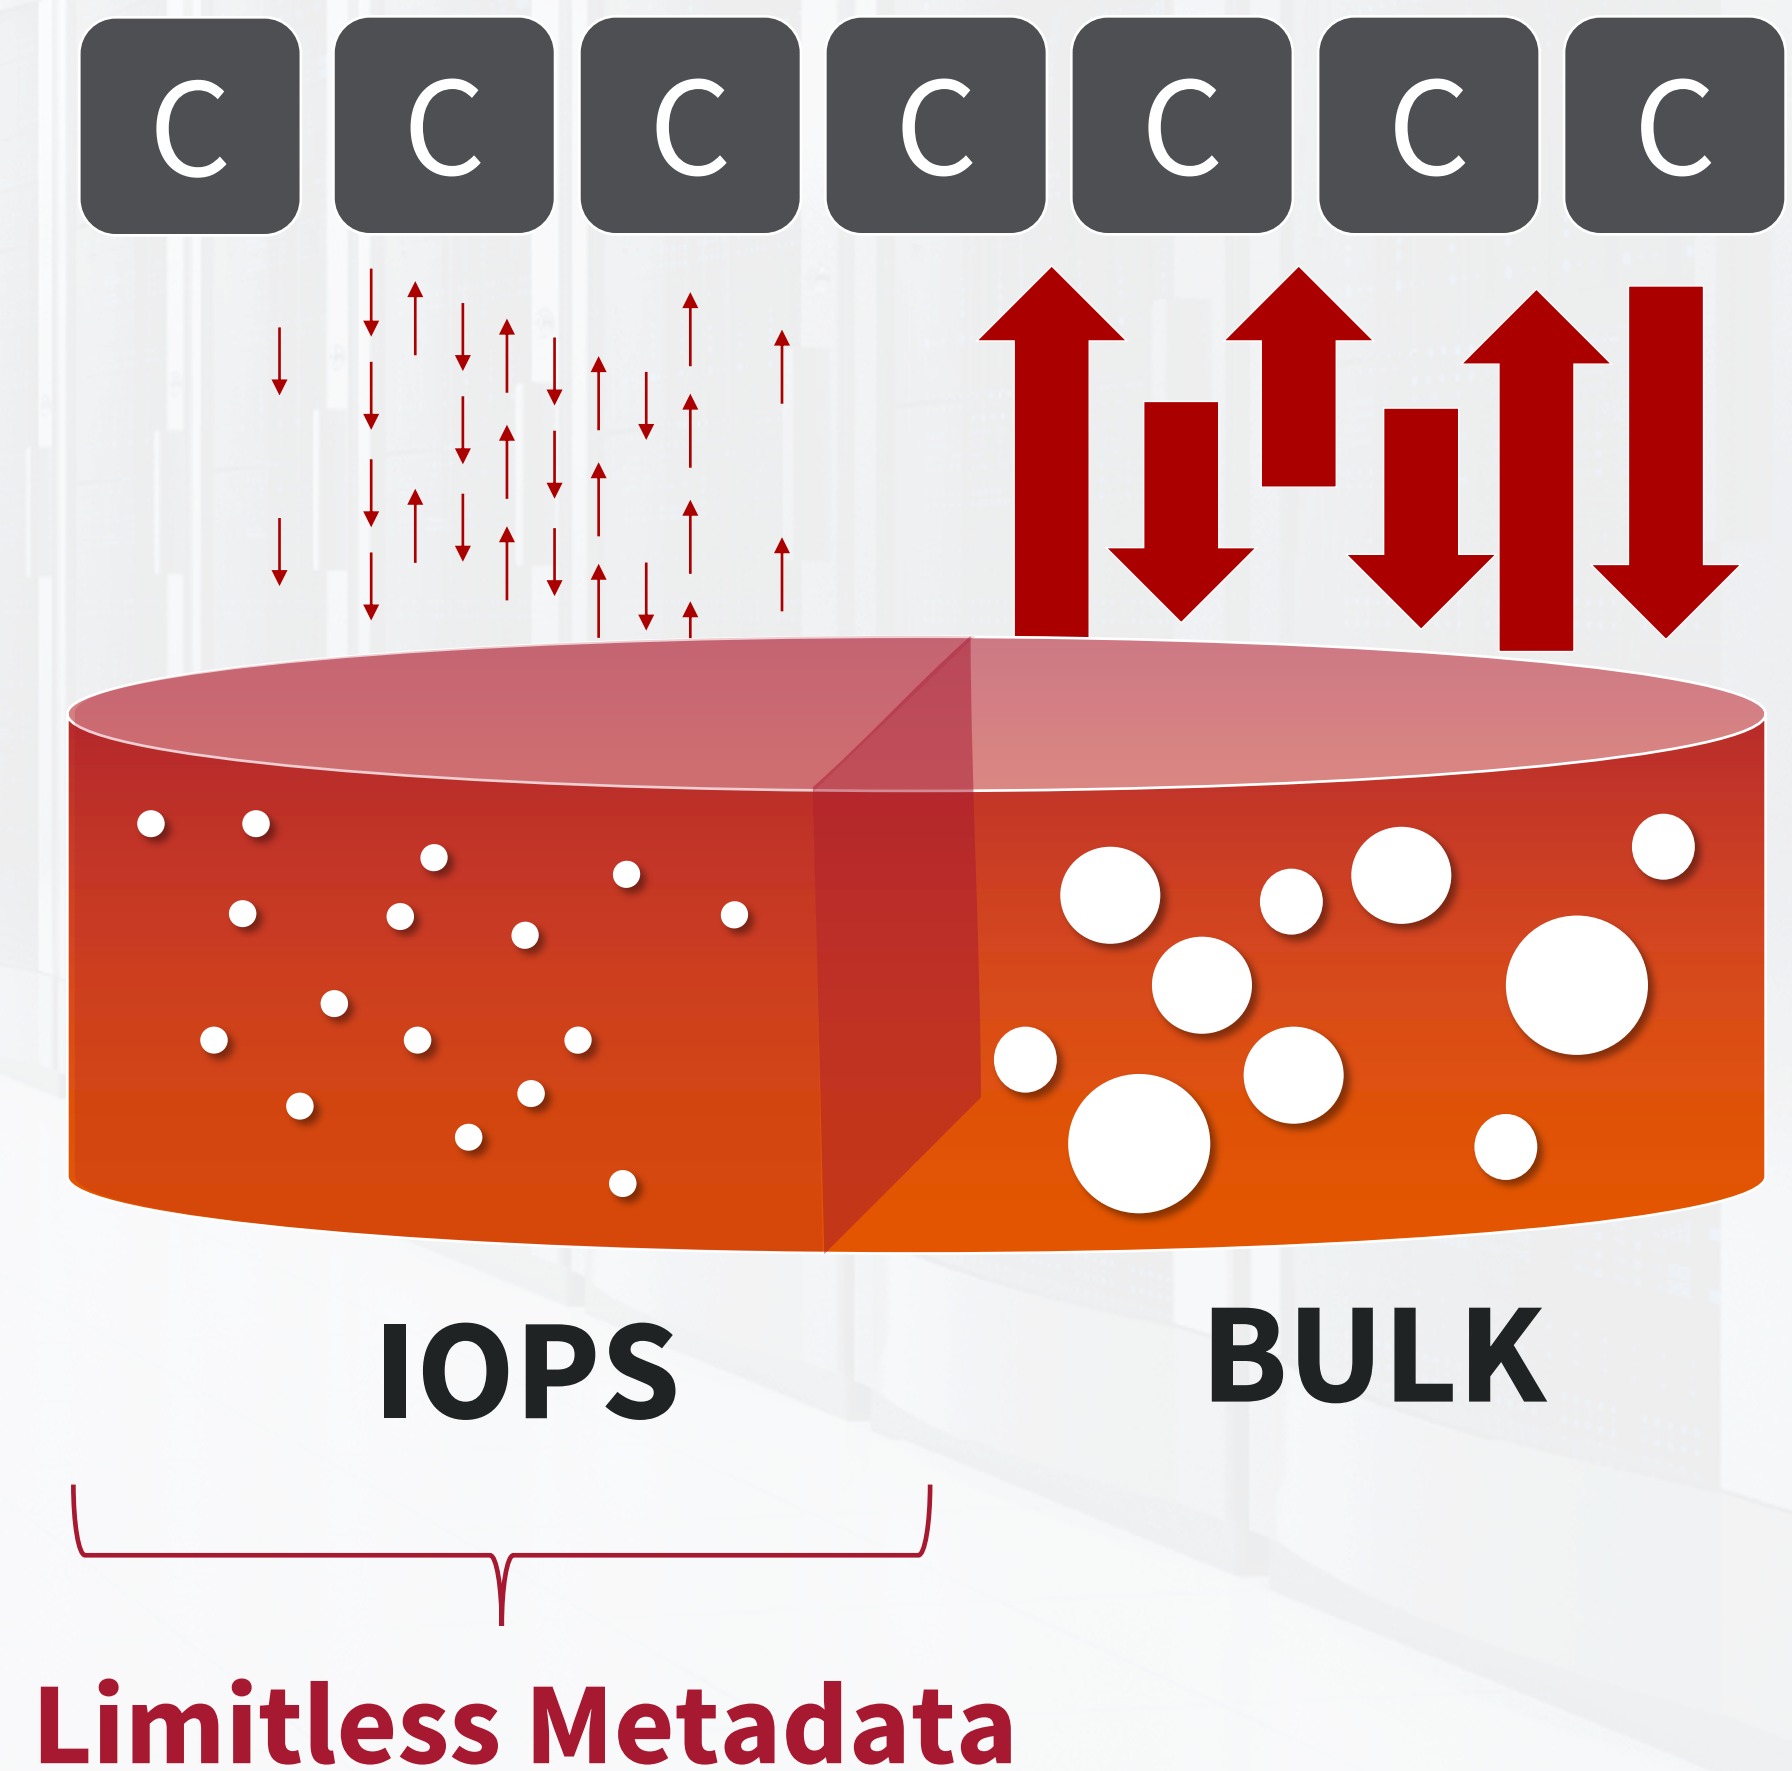
## **Tenant Self Service**

- Create and manage Data services
- Apply SLAs
- Apply security policies
- Request more resources



# Breakthrough Performance for **ANY Workload**

- Hybrid elastic IO Engines
- **Automatically** accelerates the **right** data with the **right** engine
- Data Placement for each engine is optimized for IO pattern and size



# New Paradigm of Secure AI

## Who Needs Secure Data at Scale?

Private and AI as a Service Cloud  
Operations

Secure government

Life Sciences with Medical Data

## What are the Benefits we are Offering?

Eliminate almost All common threat  
models arising from malicious entry



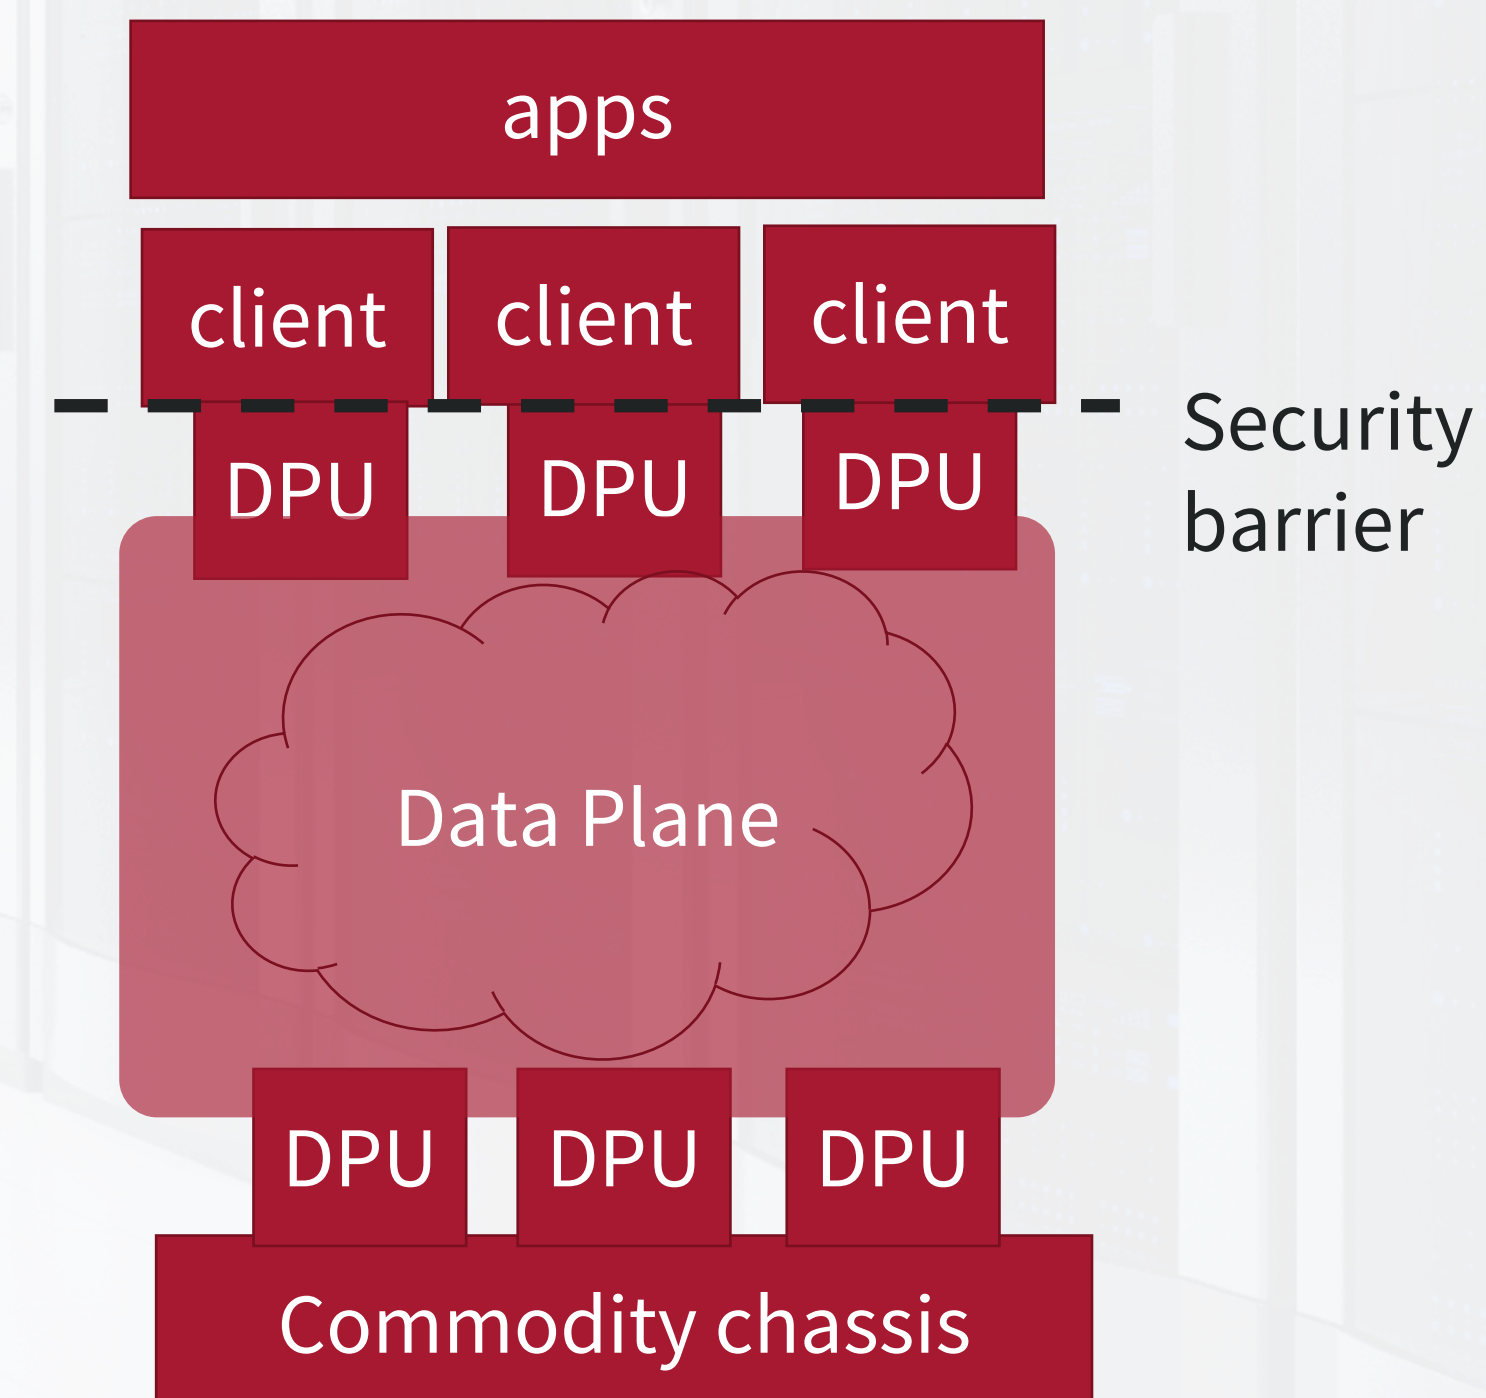

# Why Use Data Processing Units?

- **Eliminate Network Management Complexity for Data**

- SW stack manages congestion/end to end network efficiency
- Data Network reduced to physical connectivity layer - no need for smarts or knobs

- **Attack vectors are significantly reduced**


- Introduce strict security barrier that separates applications from the storage environment
- Reduce the API between app and storage to that ONLY of VirtIO
- Much stronger security than any storage and Network in isolation has today



# DDN Expanding Data Services

- One Core Data Plane
- Multiple Data services

## Multi Tenanted Data Services







# **A Platform for EXAScale that goes beyond Performance**

**Come and Talk to DDN to find out more!**  
**[jcoomer@ddn.com](mailto:jcoomer@ddn.com)**