

# Huawei: Leading Global Provider of ICT Infrastructure and Smart Devices



Bring digital to every person, home and organization for a fully connected, intelligent world

Huawei's end-to-end portfolio of products, solutions and services are both competitive and secure. Through open collaboration with ecosystem partners, we create lasting value for our customers, working to empower people, enrich home life, and inspire innovation in organizations of all shapes and sizes.

At Huawei, innovation focuses on customer needs. We invest heavily in basic research, concentrating on technological breakthroughs that drive the world forward.



**196,000**  
Employees



**80,000+**  
R&D employees



**170+**  
Countries and regions



**36**  
Joint Innovation Centers



**68** in  
Interbrand's Top 100 Best Global Brands



**72** in  
Fortune Global 500

# Enhanced Investment in Cloud & AI

## Products & Solutions

### ICT Infrastructure

Network

Cloud & AI

### Smart Devices

## Fundamental Research

**Basic Research**

**Ascend 310/910**

Inclusive AI for industries

**Fundamental Engineering**

**Kunpeng 920**

ARM-based CPU industry's highest performance

**Joint Innovation**

**Solar 6.0**

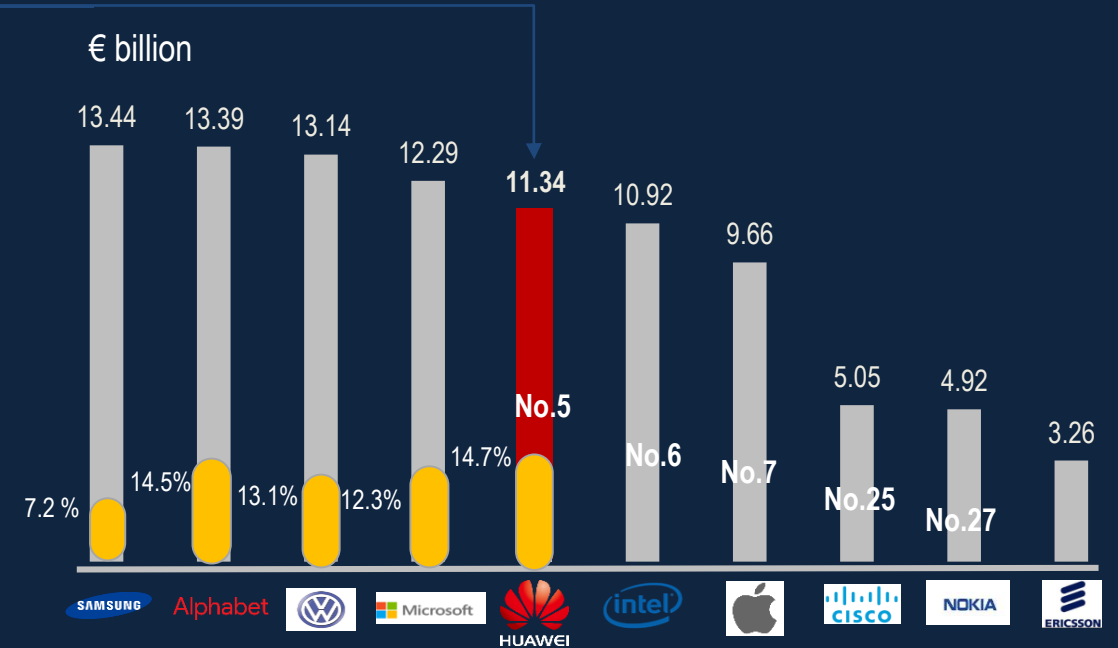
Industry's first 6.4T router chip 7nm, 0.05w/Gbit

**Graphene**

Graphene for heat spreading

## TOP 5 @ The 2017/18 Industrial R&D Investment

\*: As of Dec., 2018



Source: European Commission

- **12~15%** of Annual Revenue Invested in R&D
- **15%+** of Annual R&D Investment on Fundamental Research



The Cutting Edge of Storage Innovation



Commitment to Business Continuity



Extreme Performance Experience



Business Always-On

# The Cutting Edge of Storage Innovation

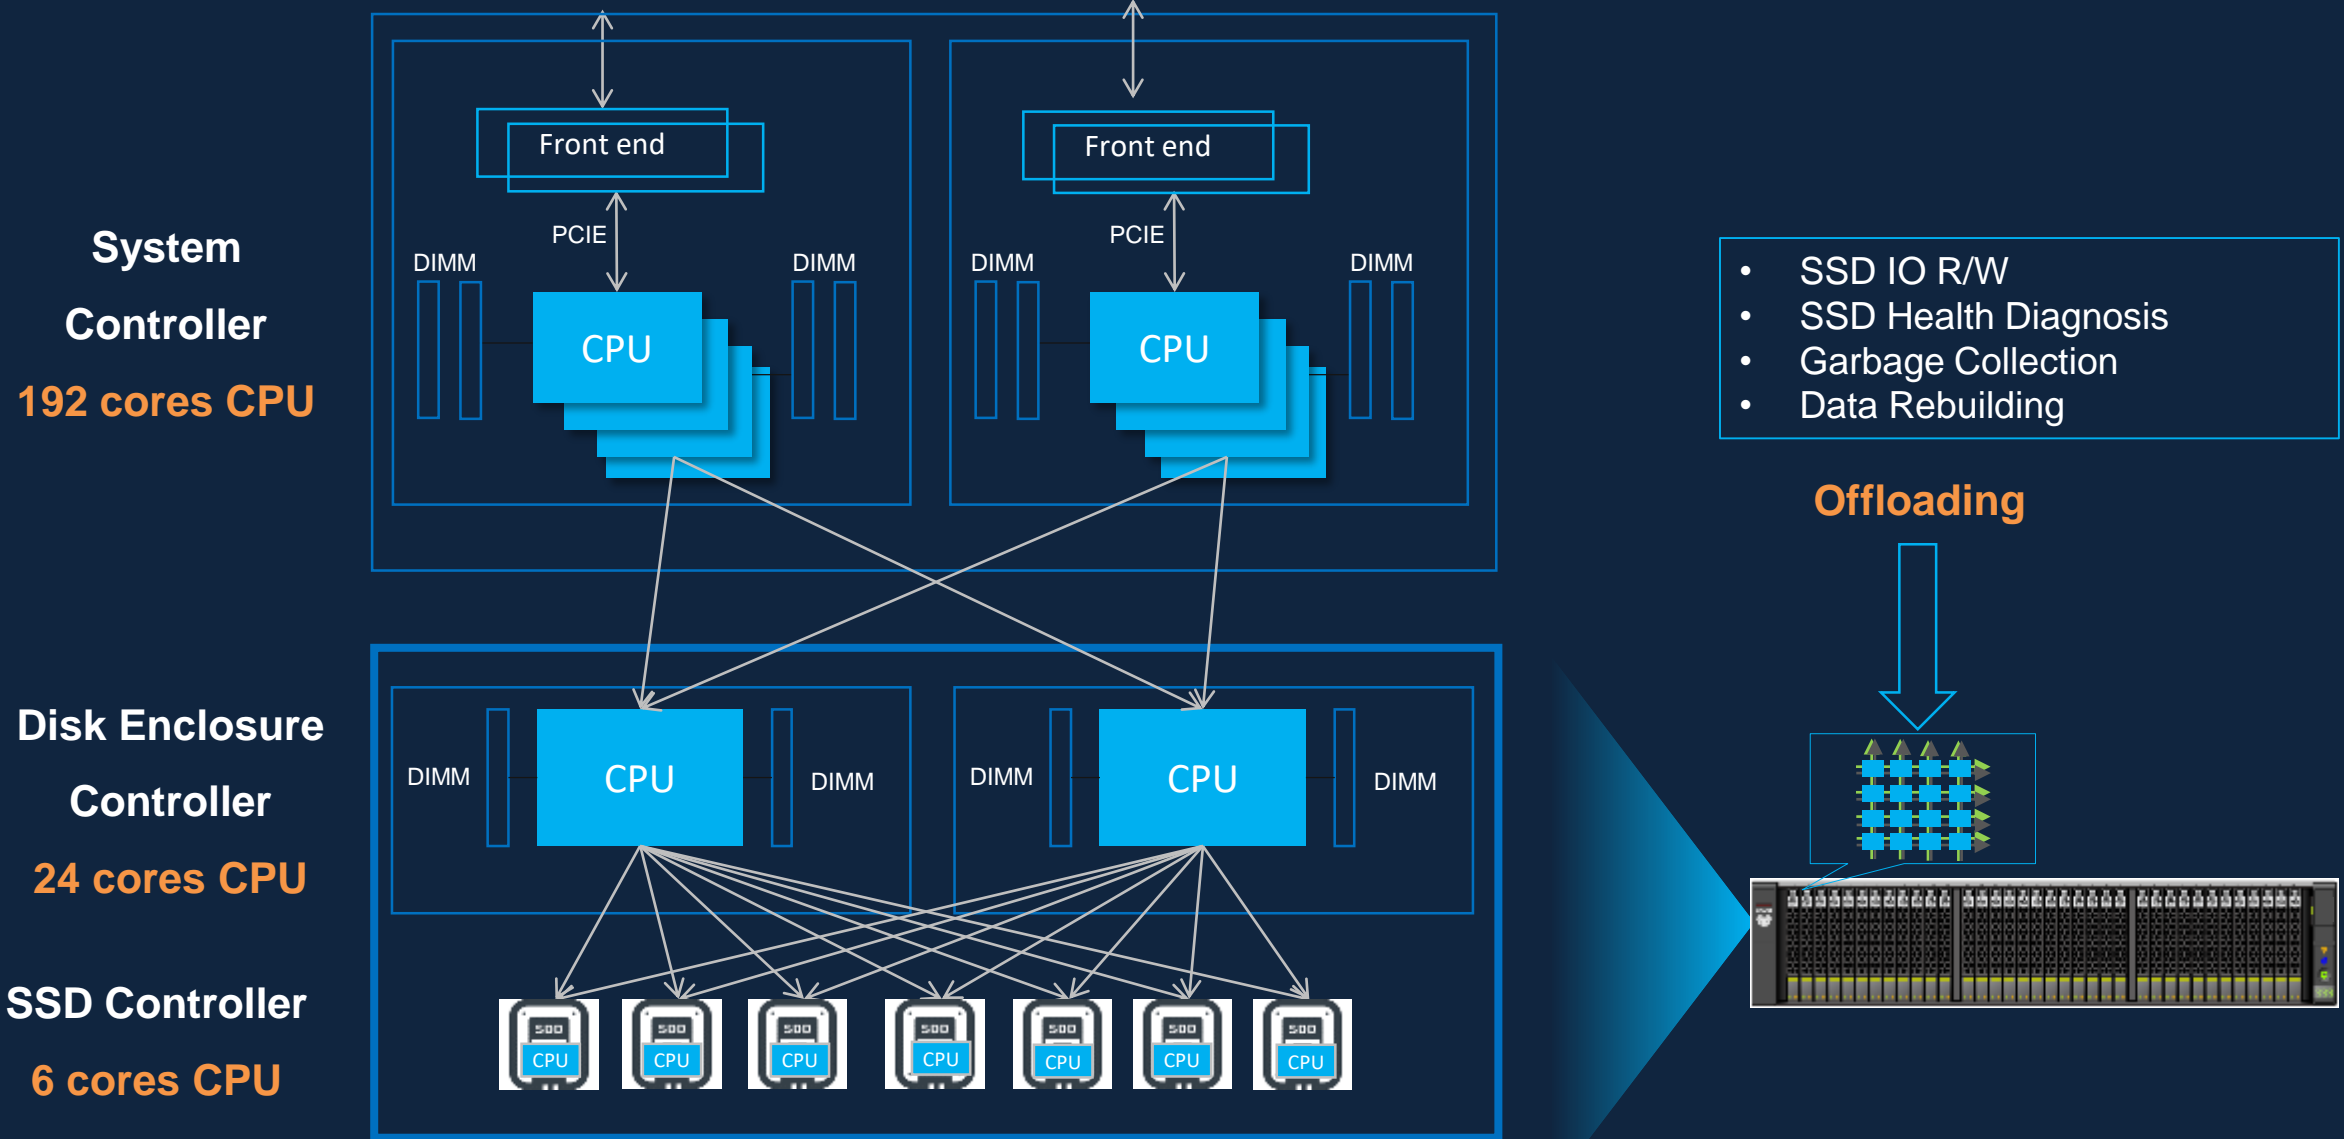
SMARTMATRIX  
GLOBAL CACHE  
**INTELLIGENT DAE**  
HIGH DENSITY DAE STORAGE FEDERATION  
HIGH DENSITY DAE RAID 2.0+ END-TO-END NVME  
ONE-SECOND NDU STORAGE ACTIVE/ACTIVE CLUSTER  
**ENGINE LEVEL FAULT TOLERANCE**  
FLASHEVER SELF-DEVELOPED CHIPS SMARTMATRIX  
SELF-DEVELOPED CHIPS ONE-SECOND NDU RAID 2.0+ GLOBAL CACHE  
PALM DRIVE ONE-SECOND NDU STORAGE FEDERATION END-TO-END NVME  
**SYMMETRIC ACTIVE/ACTIVE CONTROLLER**  
SMARTMATRIX STORAGE ACTIVE/ACTIVE CLUSTER HIGH DENSITY DAE  
PALM DRIVE FLASHEVER PALM DRIVE END-TO-END NVME  
SELF-DEVELOPED CHIPS **AI ENGINE**



# OceanStor Dorado - New Gen of Mission Critical Storage

	OceanStor Dorado (2017)	OceanStor Dorado (2019)
Max. Performance	7M IOPS	20M IOPS
Max. Storage Controller	16	32
NVMe Support	Back-End	End-to-End
Backend Network	SAS/PCIe	100Gb RoCE v2
SSD Form Factor	25 Drive/2U Shelf	36 Drive/2U Shelf
SSD Shelf	Standard DAE (No CPU)	Intelligent DAE
Data Deduplication	Fixed-Length	Fixed & Variable Length
Controller Fault Tolerance	1 of 2	7 of 8
Engine Fault Tolerance	N/A	1 of 2
Artificial Intelligence (AI)	N/A	AI Module with Ascend Chip

# 3-Layer CPU Architecture




# FlashLink® - The Foundation of Evolution



## Multi-Protocol Network Chip: Hi1822

Support both FC and Ethernet



## AI Chip: Ascend 310

• AI SoC for training



## BMC Chip: Hi1710

• Troubleshooting accuracy 93%




## Array Controller Chip: Kunpeng 920

- #1 performance ARM processor, SPECint 930+
- + Embedded intelligent processor in disk-enclosures



## SSD Controller Chip: Hi1812e

• Half the latency of previous model



## Processing power requirement

- > 1 TeraFLOPS for real-time analytics

## Ascend 310 capability

- FP16 : 8 TeraFLOPS
- INT8 : 16 TeraOPS
- Max power : 8W

## Real-time analytics

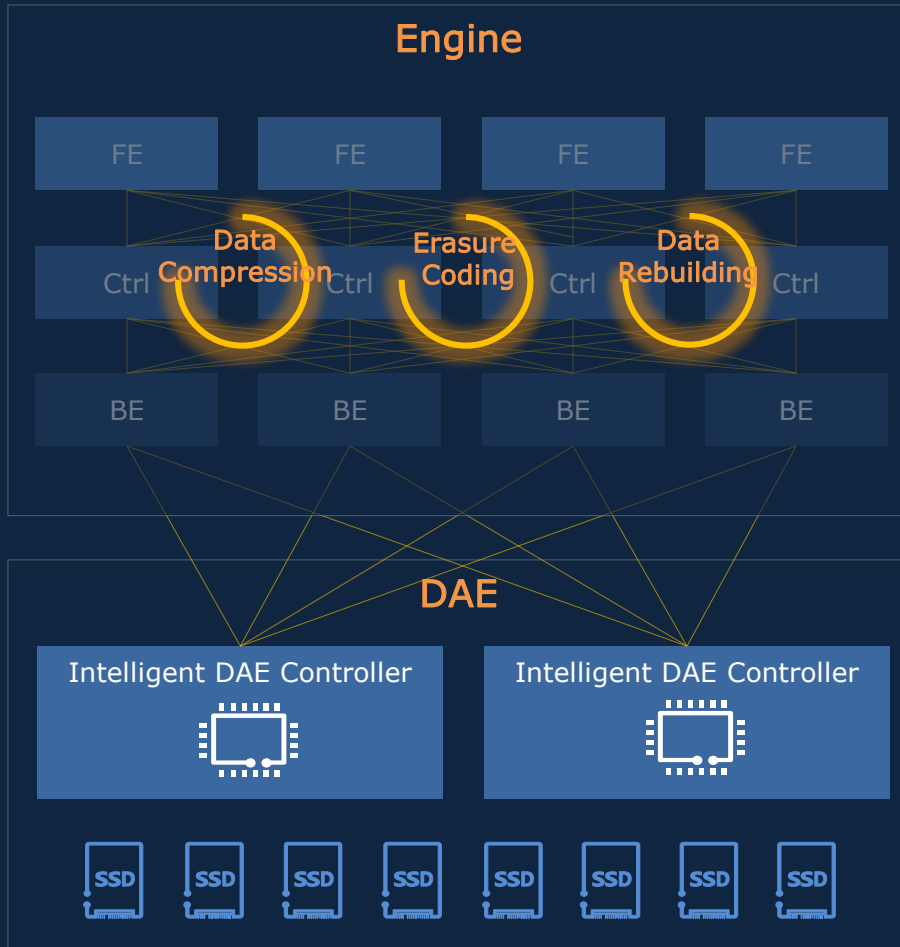
- Data Correlations ; Data Similarity ; Adaptive Optimization ; Health Analytics ; Data Temperature ; Failure Prediction

## Use case

- Intelligent Cache
- Smart QoS
- Intelligent Data Dedup

• .....

# Intelligent DAE - The SSD Shelf with Processing Power



## Storage Controller Offloading

- Each DAE has two controllers, and each controller has its own processor, cache and adapter.
- DAE controller takes over some of workloads from array controller, including:
  - Data rebuilding
  - Erasure Coding (EC)\*
  - Data compression\*

The DAE is much more intelligent than ever, this distributed computing design **half** the data rebuilding time, and reduce the performance impact (max. IOPS) of controller from **15%** to **5%**, the bandwidth for data rebuilding is increased from **80MB/s** to **200MB/s** while the array controller's CPU Util% remains at 70%.

\* The workloads will be available in near future, including garbage collection.



# New Generation Innovative Hardware Platform

High-end controller enclosure



4U, 4 controllers, 28 shared interface slots

Mid-range controller enclosure



2U, 2 controllers per controller enclosure

Entry-level controller enclosure



2U, 2 controllers per controller enclosure

Intelligent DAE



2U, 2 controllers per controller enclosure

Intelligent DAEs

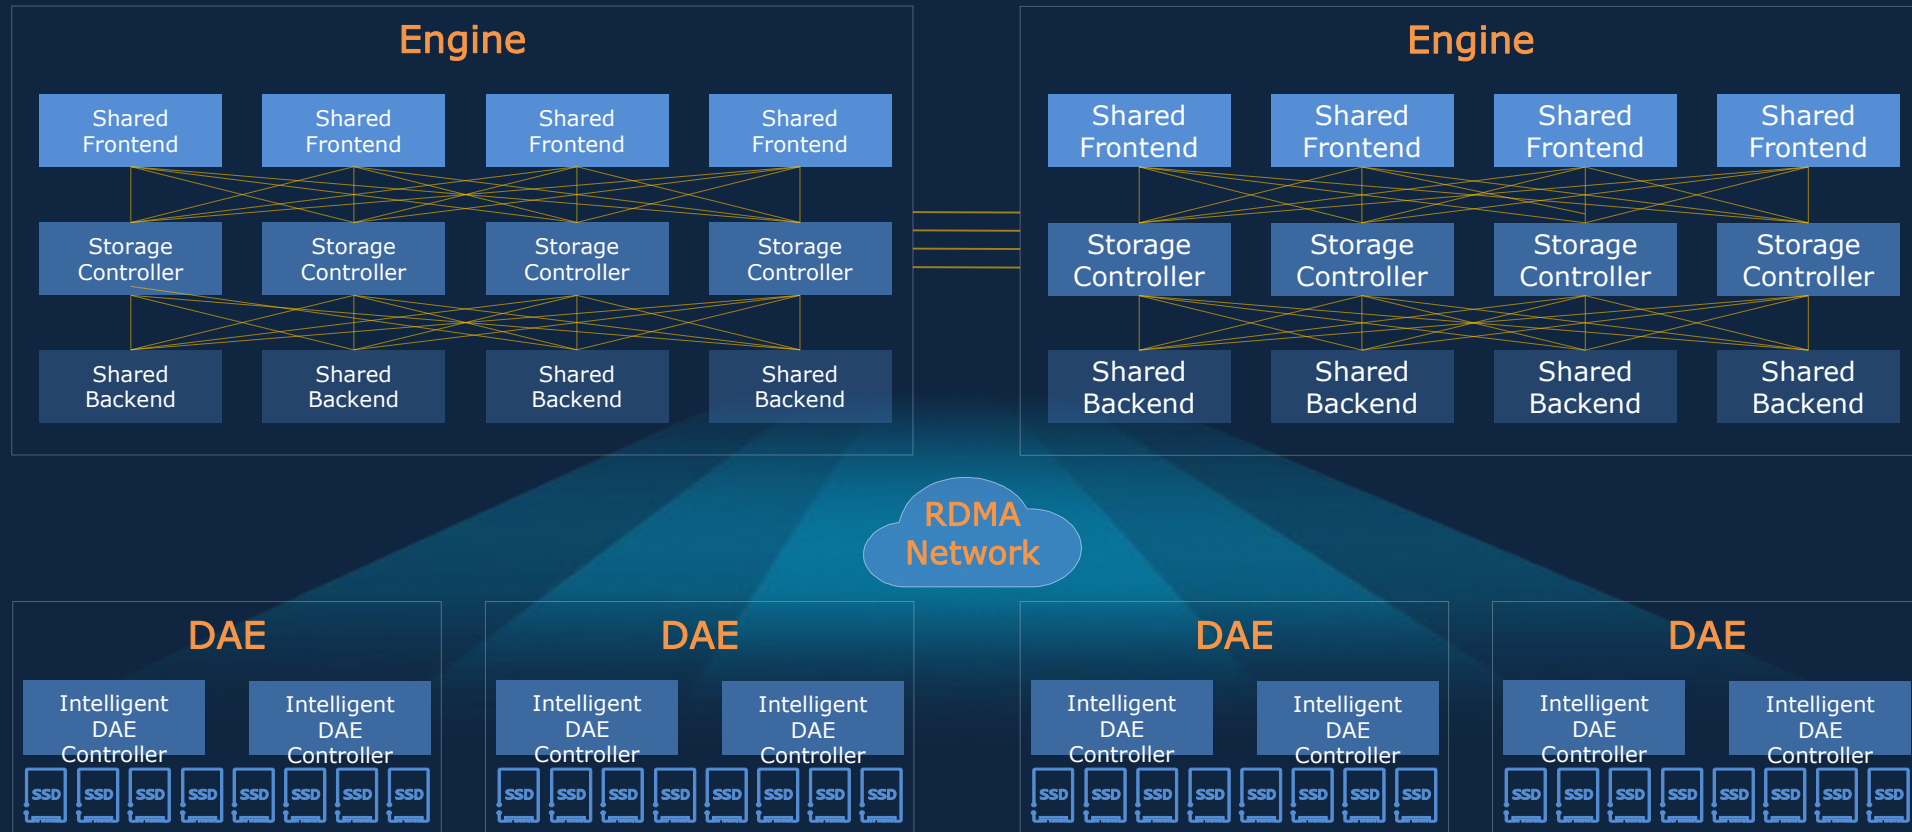


2U, 36 NVMe SSDs (high density, 44% increase)



2U, 25 SAS SSDs

# SmartMatrix - Symmetric A/A Controller Architecture



- Symmetric active/active controller with fully meshed topology
- Shared everything architecture from frontend, backend, to drive enclosure

- Persistent cache mirroring with max of 3 copies
- Non-disruptive firmware upgrade, IO hang-up time is limited within 1 second

- End to end NVMe support
- Backend RDMA network over 100Gb/s Ethernet.
- SCM support for read acceleration\*



The Cutting Edge of Storage Innovation



Commitment to Business Continuity



Extreme Performance Experience



Business Always-On

# Every Second is Valuable

Application

1XX Seconds

Timeout

Operating System

XX Seconds

Timeout

Host Bus Adapter (HBA)

XX Seconds

Timeout

Network

1X Seconds

Timeout

Storage

X Seconds

Timeout



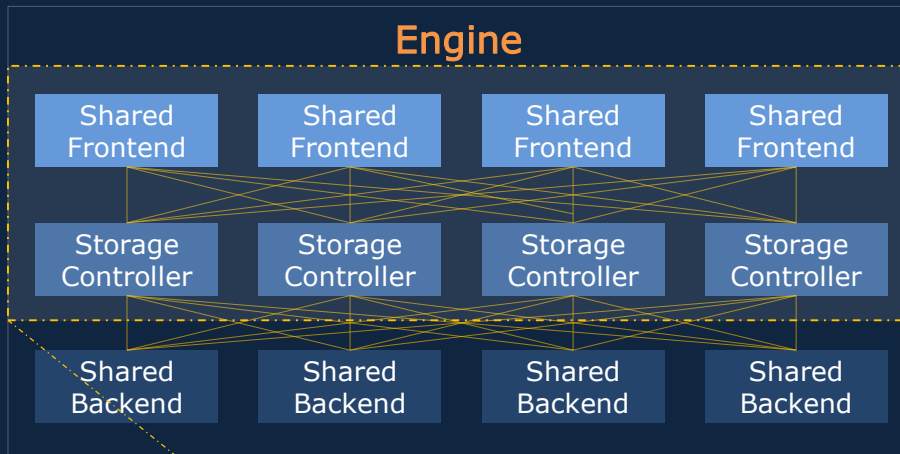
Each second of timeout to mission critical business could be:

- Tens percent of transaction lost
- Thousands of dollars profit lost
- Tens of thousands unsatisfied customer (specifically on black Friday)

Some of the big FSI enterprise require storage to minimize timeout to 1 SECOND, e.g.

- Industrial and Commercial Bank of China
- Itaú Unibanco
- China Construction Bank
- Agricultural Bank of China
- .....

# One-Second's Magic - Shared Frontend Adapter

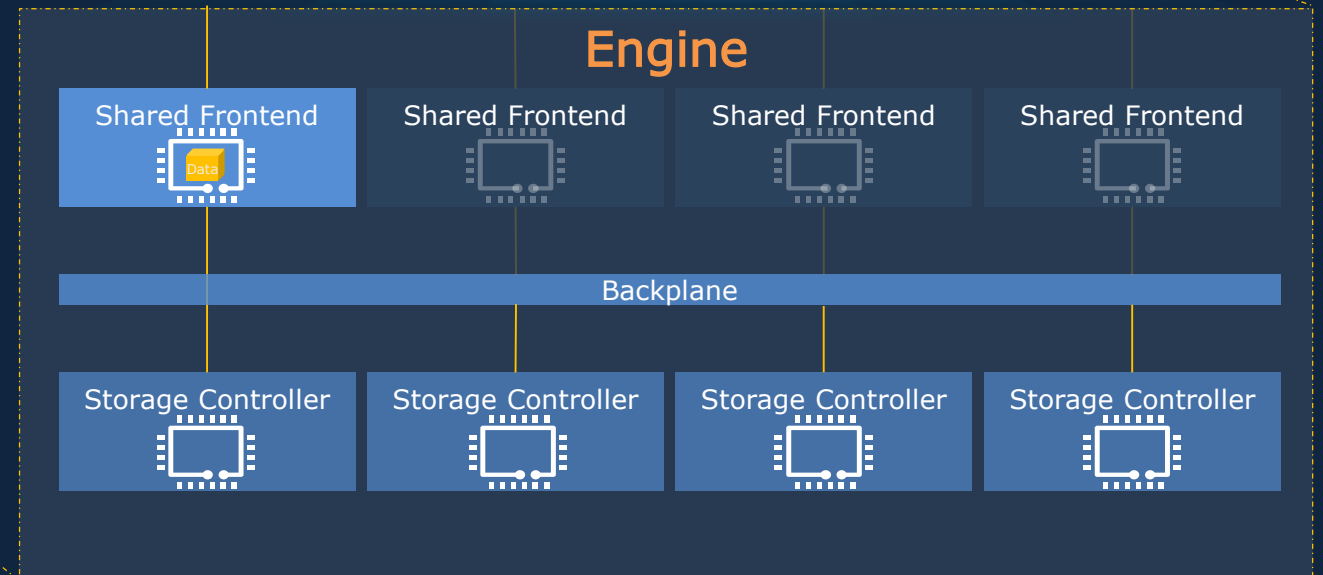


## Server



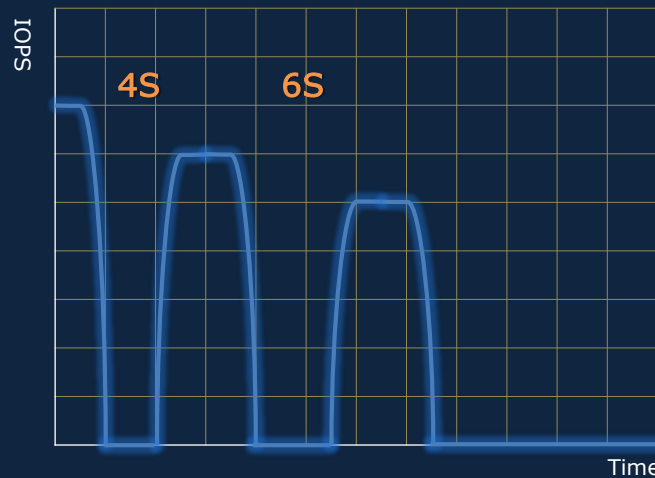
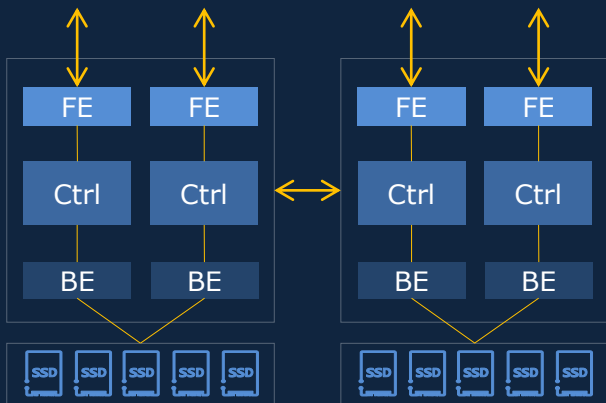
## FC/Eth Network

- Frontend adapter holds the connection with server & LUNs independently, storage controller is not involved.
- With intelligent frontend adapters, I/O can/will be load-balanced over storage controllers through backplane
- If the controller fails, the I/O will be redirected to other surviving controllers, while the connection between the frontend adapter and the server/LUN is maintained, the server is not aware of the failure.


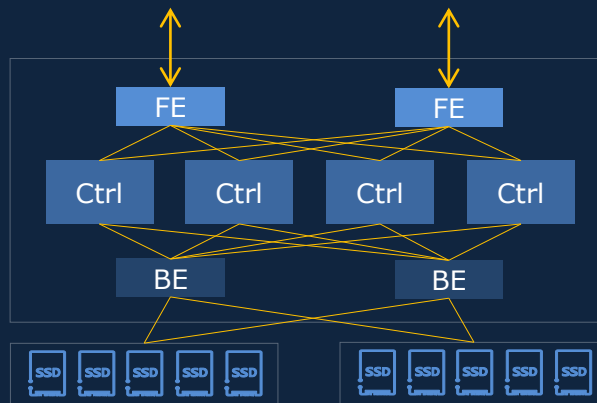


# One-Second Controller Failover

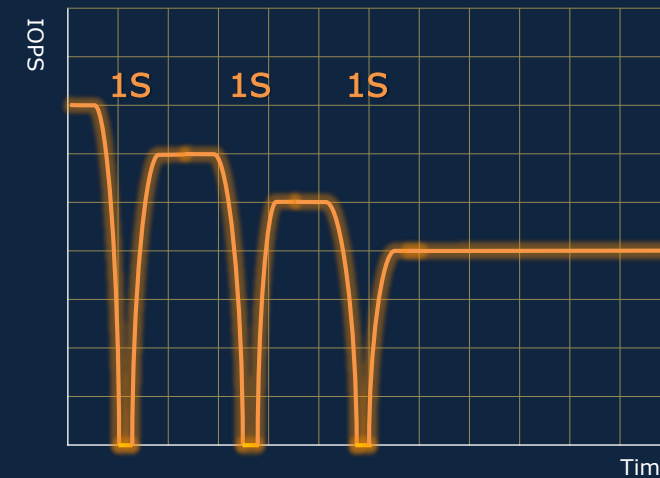
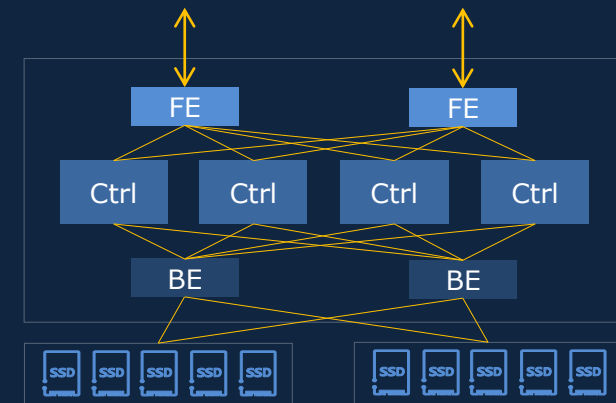
## Solution A



## Solution B



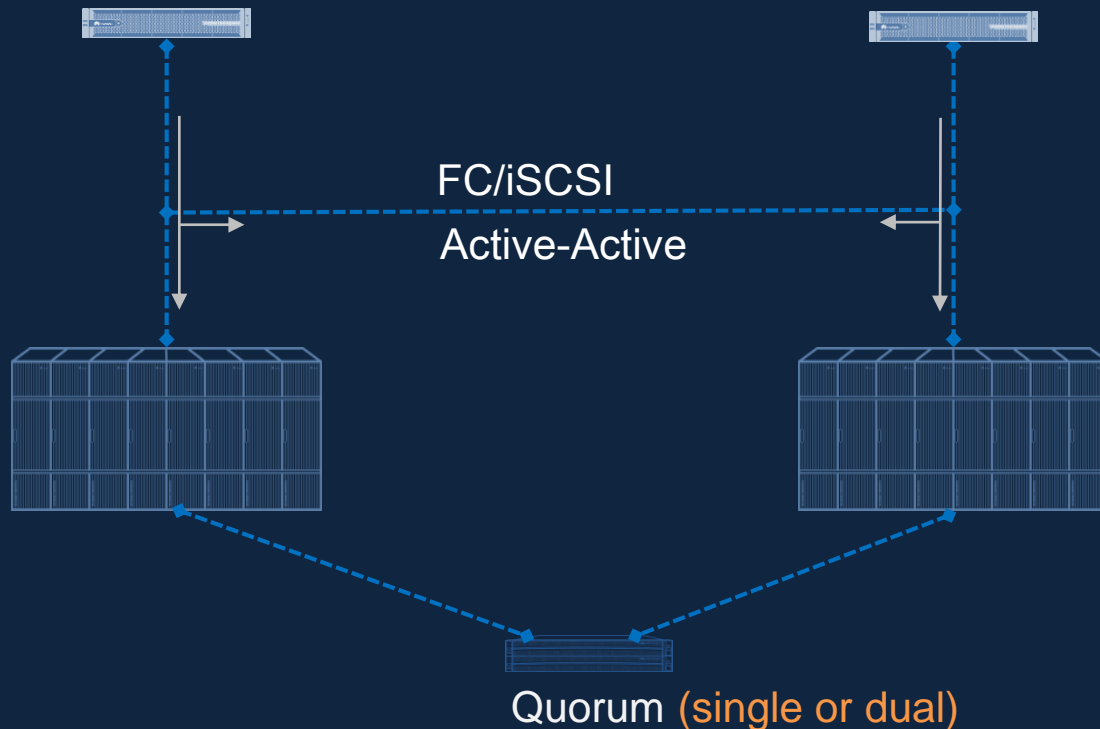
## HUAWEI



\* The above figures are referring to the testing result in Huawei lab.

# HyperMetro Gateway Free Active-Active

Unique storage architecture with 1 ms latency over 100km when active-active



## ❖ RPO=0, RTO=0

Active-Active Metro-Cluster + Replication




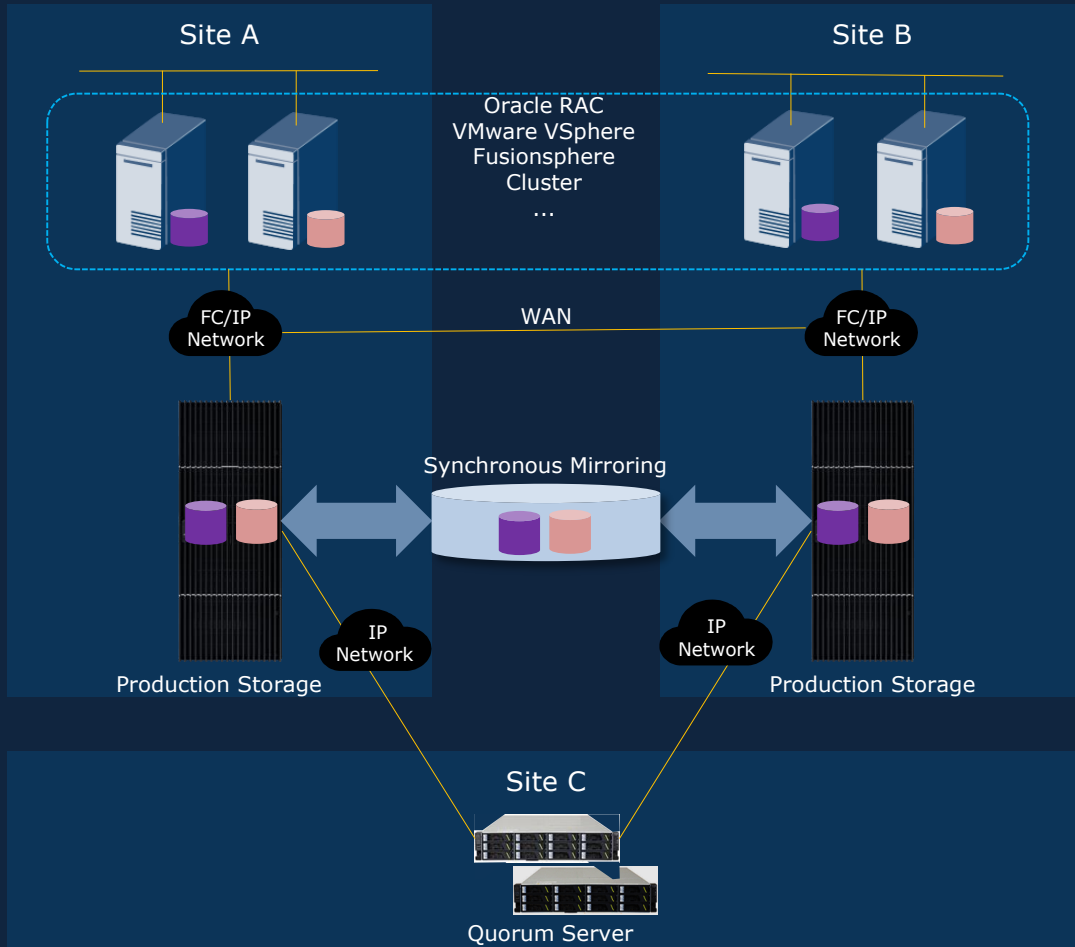
## ❖ Non gateway

Embed active-active in system, no need gateway or extra software, minimal DR cost.

## ❖ Minimal performance loss:

- 1) Non-gateway Active-Active;
- 2) Fast write technology combine metro-cluster commands and data, reduce interactions between two systems;
- 3) Two systems can read/write LUN simultaneously.

# Comprehensive HA/DR Solutions: HyperMetro & HyperReplication





# SmartVirtualization Overview



Storage – Resource – Consolidation

Heterogeneous  
Storage  
Virtualization



## SmartVirtualization Brief

- Simplified management of heterogeneous arrays and reuse of legacy storage systems thanks to the heterogeneous takeover function
- Non-disruptive migration among LUNs of heterogeneous storage arrays
- Investment protection, reuse legacy array for their life-span
- Snapshot function implements rapid backup for LUNs of heterogeneous storage arrays\*



The Cutting Edge of Storage Innovation



Commitment to Business Continuity

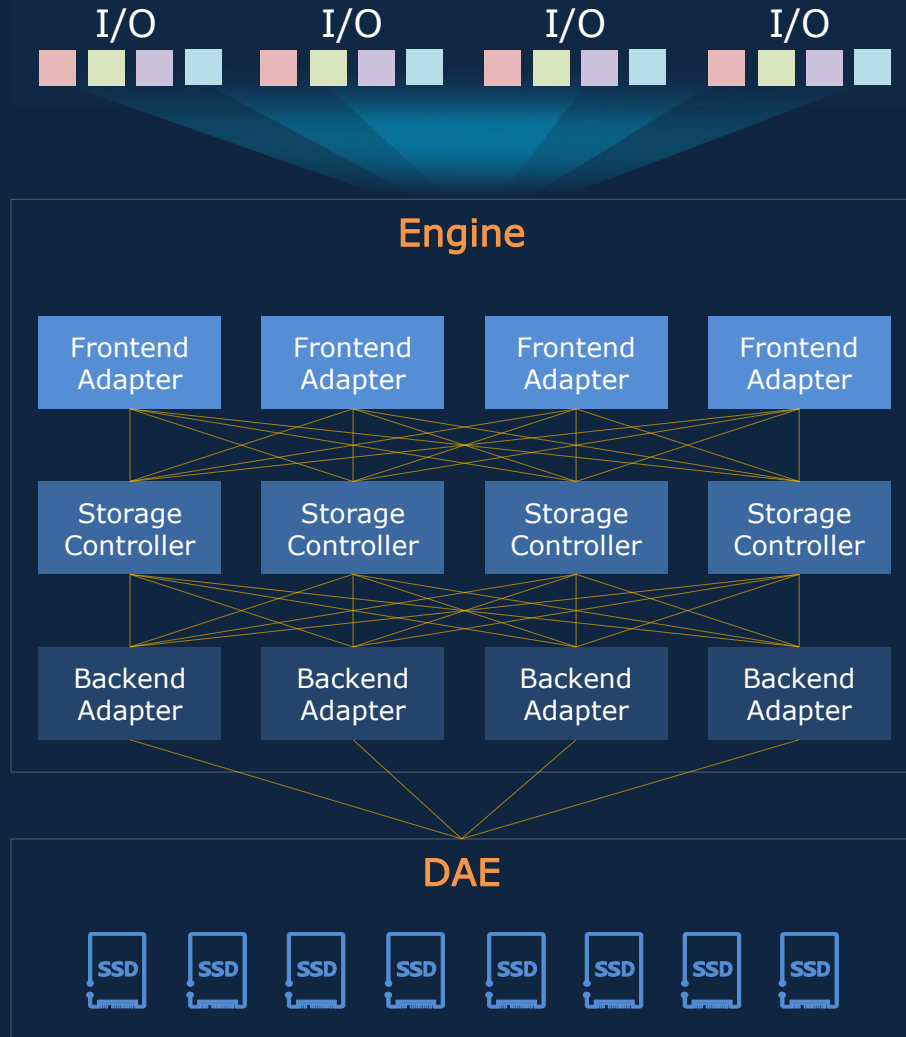


Extreme Performance Experience



Business Always-On

# End-to-End Load Balancing



## Shared Front-end Adapter

- Requests from host can be **evenly distributed on every front-end link**
- LUNs are **shared by all controllers** (aka no controller LUN ownership). **Managed by intelligent frontend adapters.**

## Global Cache

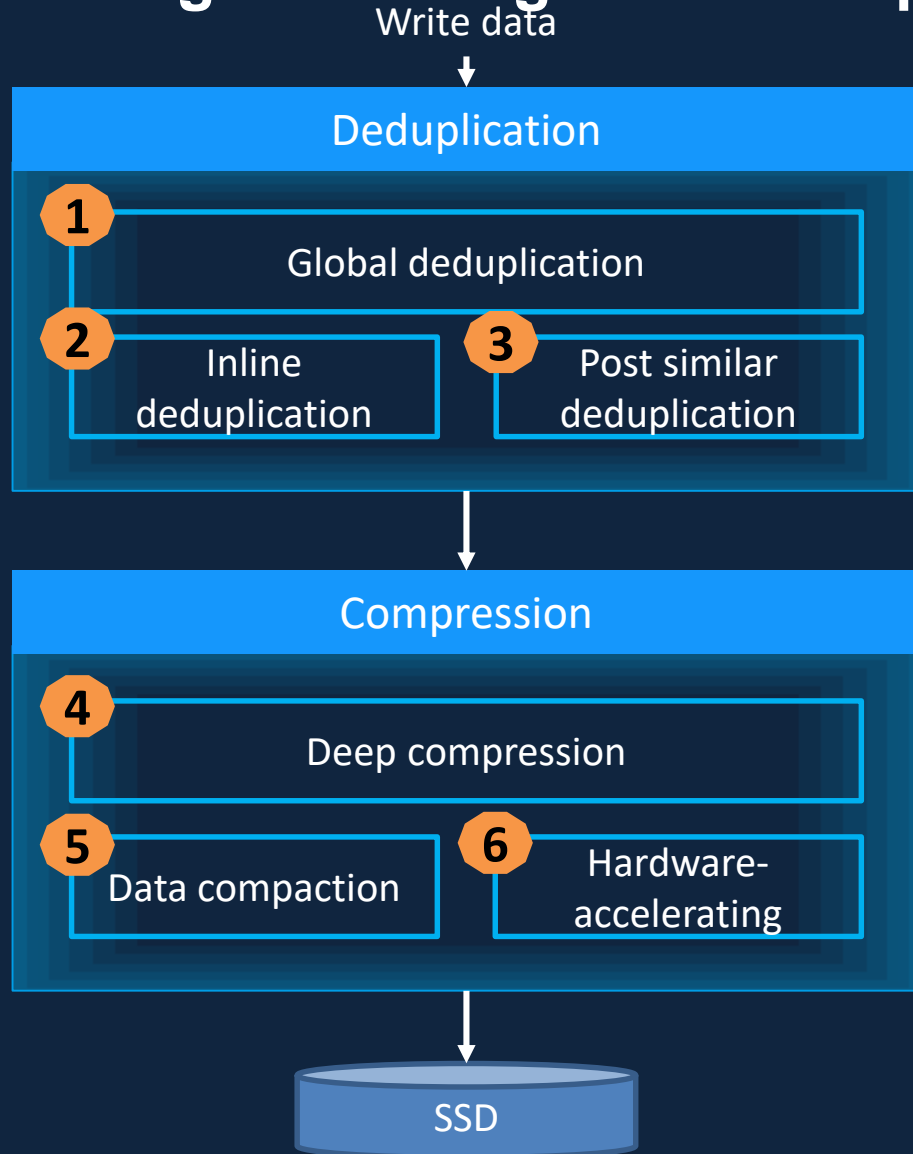
- Write I/O requests for single LUN can be placed into cache space from multiple storage controllers.
- For better cache read hit, storage controller can place the

## Global Storage Pool

- **Global storage pool** can be accessed by multiple storage controllers.
- With RAID 2.0+, multiple LUNs are distributed over multiple SSDs naturally.

# Data Reduction:

## Leading Technologies of Deduplication and Compression, average 5:1



### ① Global deduplication

Global fingerprint table for deduplication of data from all LUNs and engines in a diskdomain.

### ② Inline deduplication

Always inline, 4K~32K granularity configurable, LUN-level configurable.

### ③ Post similar deduplication

Identification and deduplication of Inner-block 64Bytes aligned similar data in the post process.

### ④ Deep compression

Self-developed deep compression algorithm.

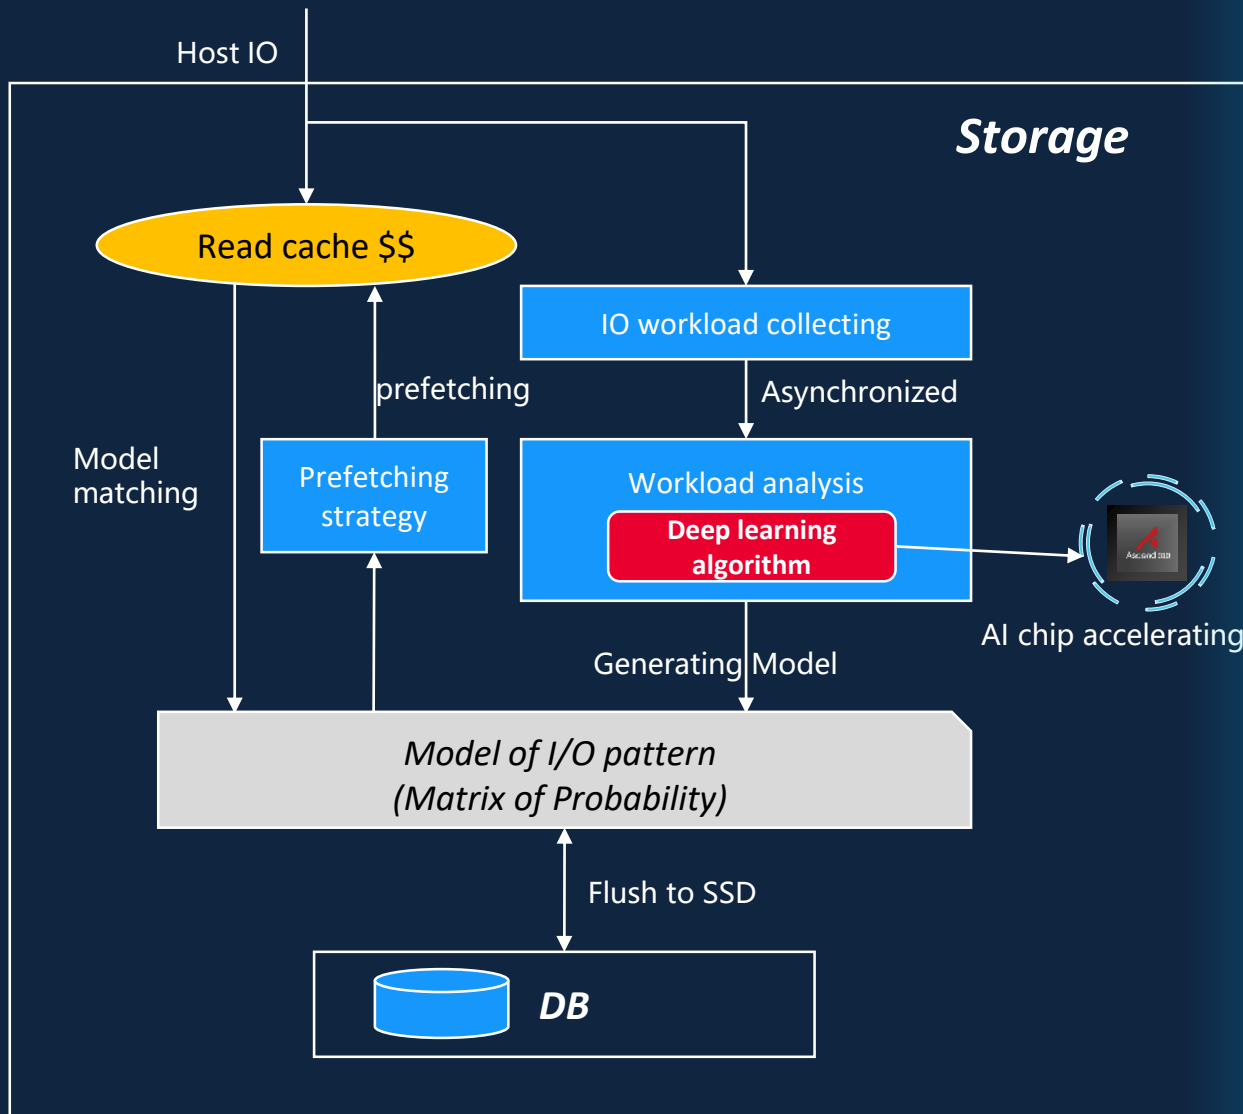
### ⑤ Data compaction

Byte level compaction, space wasted less than 1%.





### ⑥ Hardware-accelerating

Specialized compression engine in Huawei Kunpeng CPU for the deep compression algorithm.

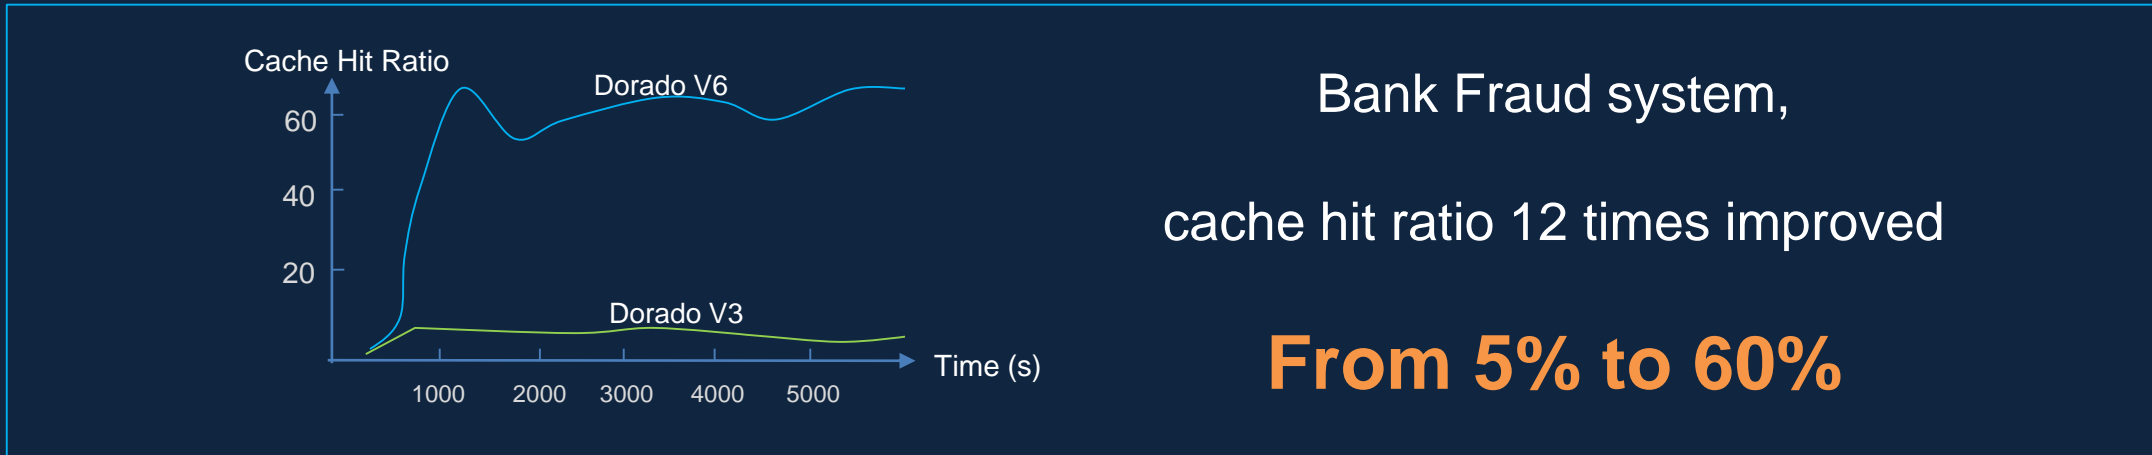
# Intelligent Cache by Deep Learning in Storage



# AI Promotes massive Performance increase



Hot data pre-fetch to relative data pre-fetch



Bank Fraud system,  
cache hit ratio 12 times improved

**From 5% to 60%**



The Cutting Edge of Storage Innovation



Commitment to Business Continuity

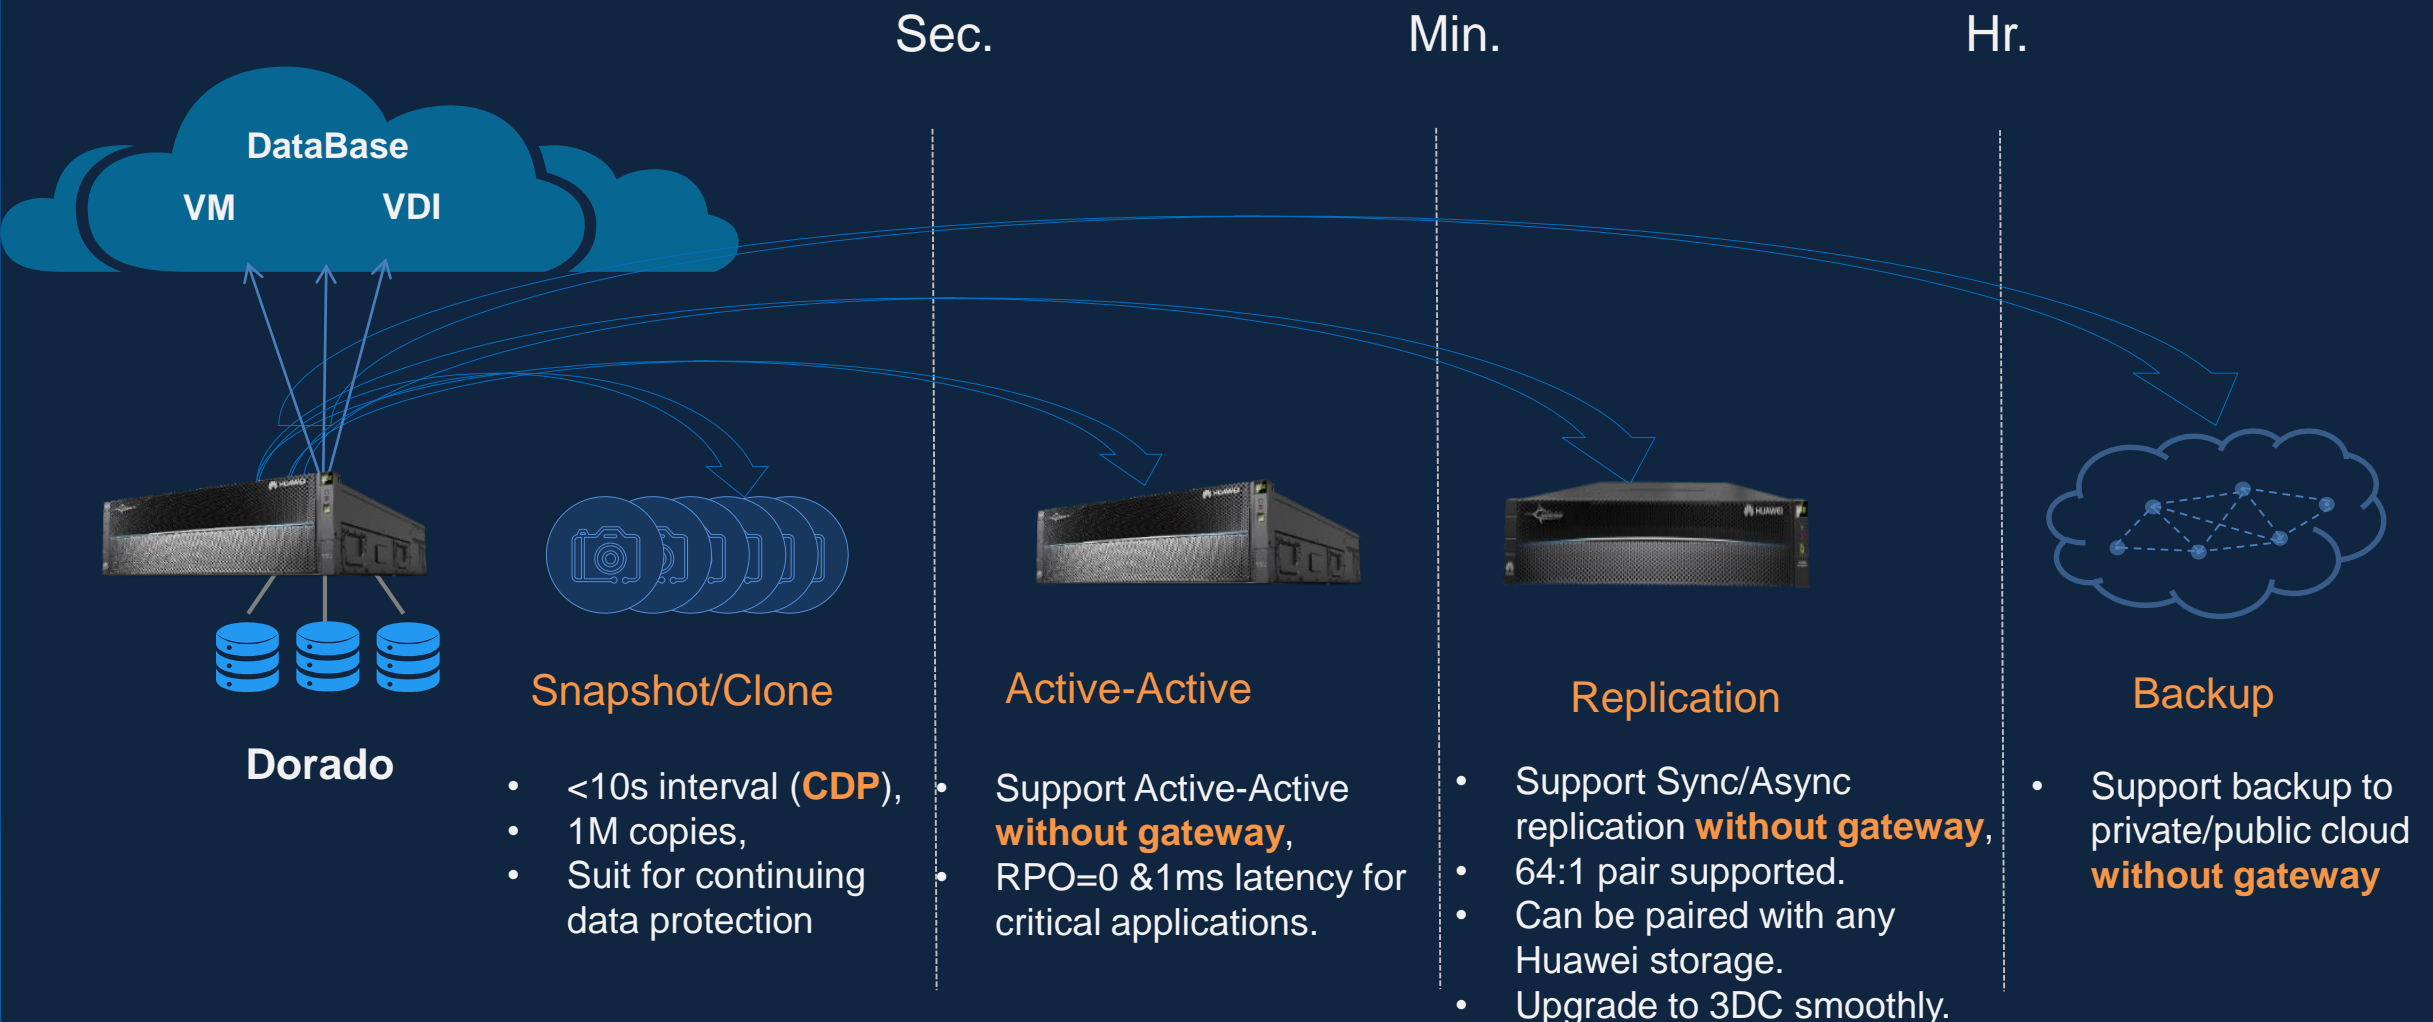


Extreme Performance Experience



Business Always-On

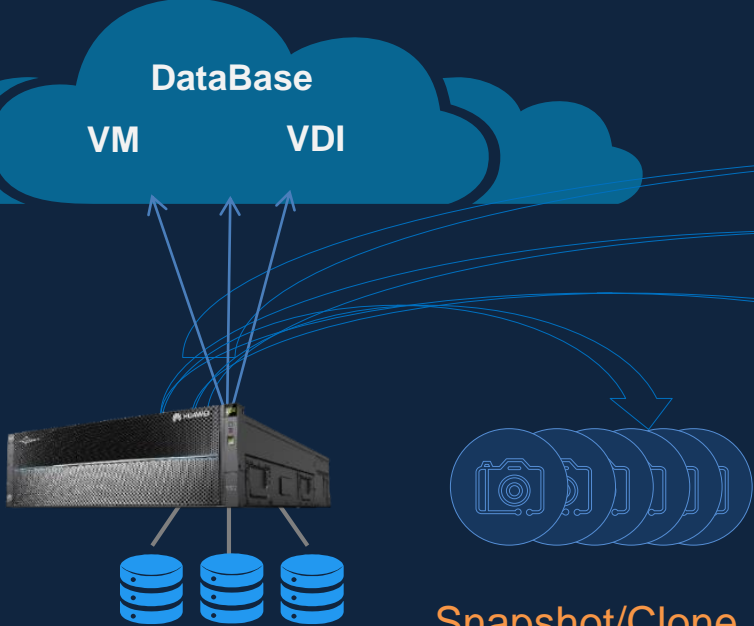
# Simple O&M: 4 Tiers Data Recovery without Gateway



Sec.

Min.

Hr.



### Snapshot/Clone

Dorado

- <10s interval (**CDP**),
- 1M copies,
- Suit for continuing data protection




### Active-Active

- Support Active-Active **without gateway**,
- RPO=0 & 1ms latency for critical applications.



### Replication

- Support Sync/Async replication **without gateway**,
- 64:1 pair supported.
- Can be paired with any Huawei storage.
- Upgrade to 3DC smoothly.



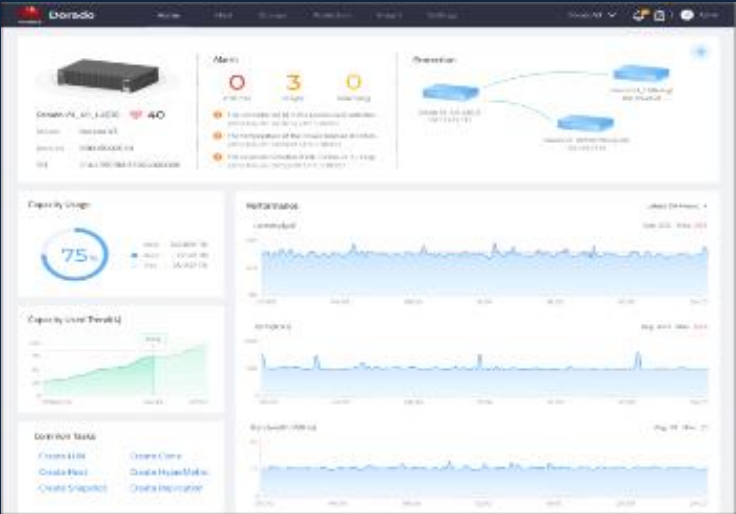
### Backup

- Support backup to private/public cloud **without gateway**

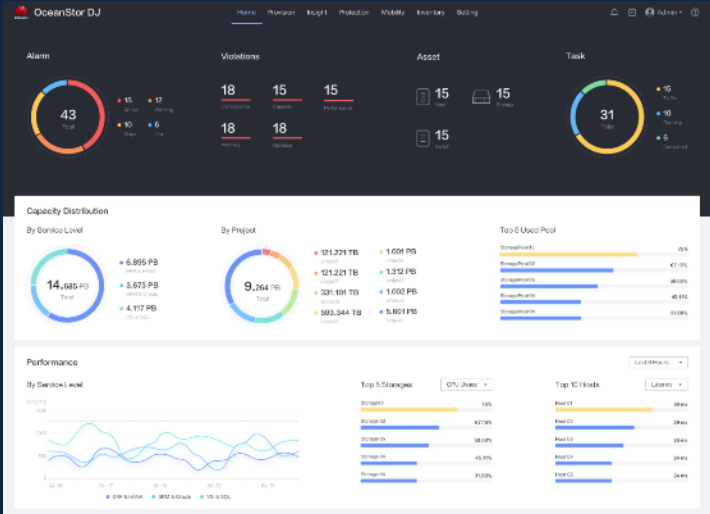


# Management Software: DeviceManager , DJ and eService

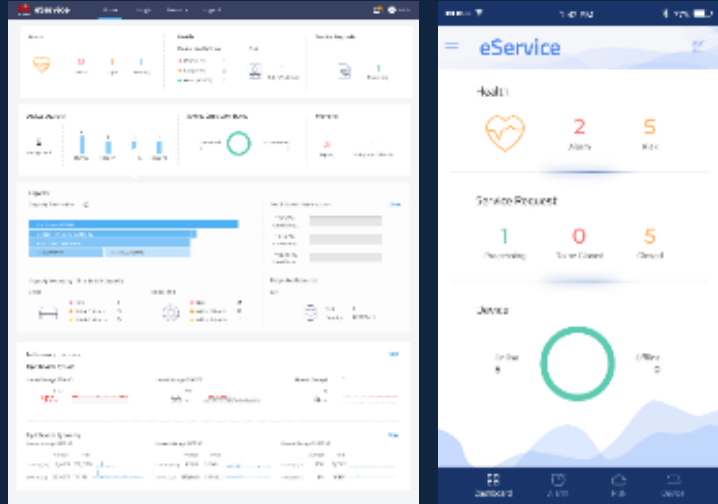
## Device Manager



## DJ - SDS



## eService – AI based




- Simplified configuration experience, facilitating service provisioning, service protection, and system configuration.
- Provides the object relationship topology visualization capability.
- Comprehensive performance monitoring and analysis capabilities

- SDS Platform
- Provides the central management of multiple storage arrays
- Full life-cycle resource management
- SLA-Based automatic provisioning template-based provisioning
- Policy-based risk and compliance management
- AI-enabled prediction on capacity/performance

- Massive data-based AI training on Cloud
- Advanced services such as fault prediction, capacity management, and performance exception analysis.
- Collaboration with Huawei Support teams and solve customer system hidden trouble ahead of time.

# FlashEver Program for 10-Year + Always-On Business Non-Disruptive Tech. Refresh



## FlashEver Program

- Support for Non-Disruptive controller upgrade, even include next several generations, minimum 10 years
- Tech refresh within the existing assets to obtain the advantages of the latest technology
- No data migration

## Storage Federation

- Up to 128 controllers
- Support OceanStor Dorado and the following generations
- Can mix different gens of OceanStor Dorado in one federation cluster
- Support data mobility non-disruptively, and online node reorganization

## Smart Virtualization

- Virtualize third party storage by taking over the access path
- Reuse old storage to protect customer investment
- Smoothly cutover the business to run in OceanStor Dorado and the following generations

# Take Away

Growing with Your Business  
Migration free life-cycle-management

Highest Availability

AI-Chipset  
Smart Data Management

NVMe, F1 Race Track

**7 of 8**  
Controllers failure tolerance

**5:1**  
Data reduction

**10**  
Years minimum EverNew

# OceanStor Dorado V6 — Ever Fast, Ever Solid, Ever Available

## Building a Fully Connected, Intelligent World

**AI-Powered**

Build the foundation  
of convergent  
processing

**Automated**

Greatly improve  
efficiency  
and save OPEX

**Affordable**

Replace HDDs with  
SSDs  
for high performance  
and reliability

