

TaiShan Server SDS Solution(Mixed Deployment) User Case

Huawei ICT Day
Steven Wangjinchao
Director Intelligent Computing Development Dept

Security Level:




Contents

01 User Case

- Challenges
- Taishan SDS Solution

02 TaiShan Products

Introduction

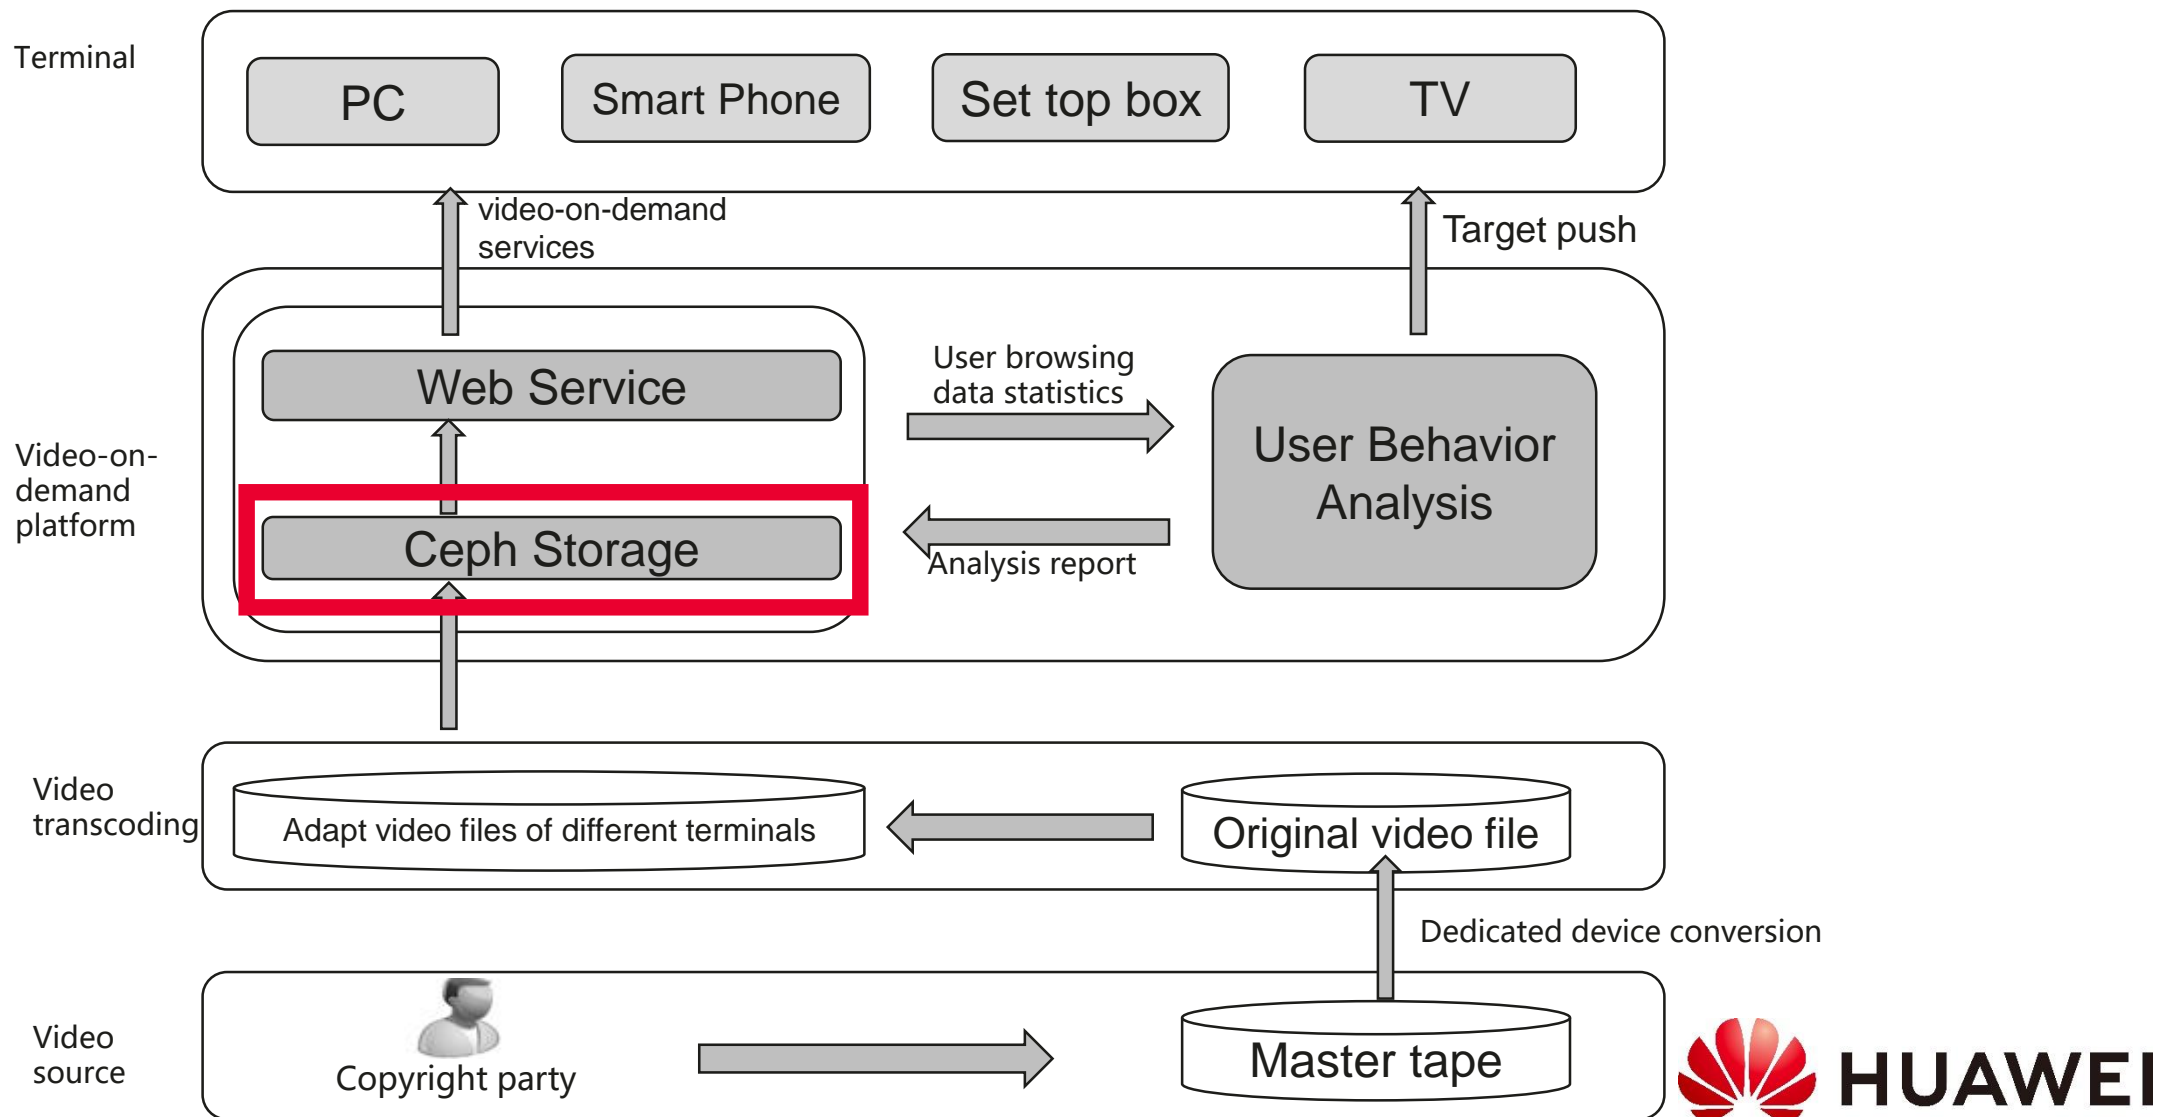


Why Ceph		
Community	High Scalability	High Availability
Active open source community and rich information	expansion with no interruption	Multiple copy mechanism, fault domain separation, can endure fault domains and repair itself

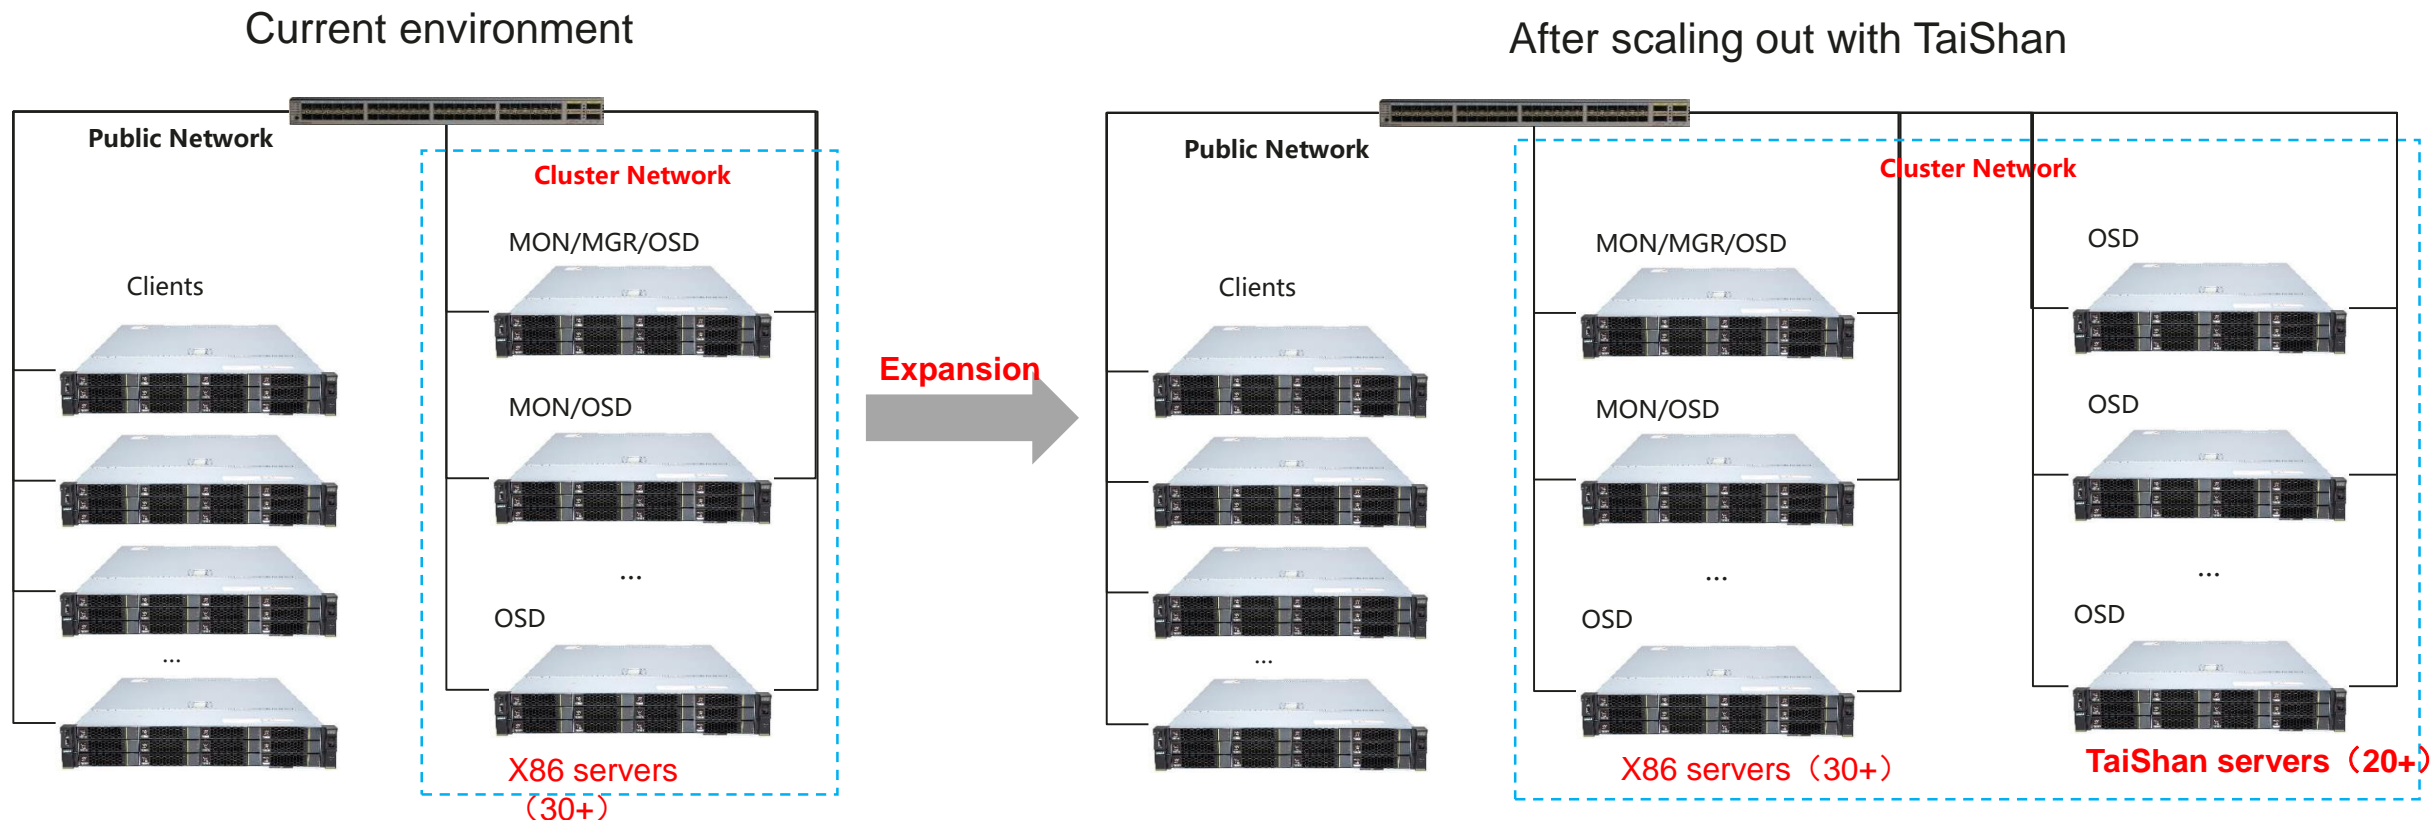
Why HUAWEI TaiShan		
Performance	Power Consumption	Compatibility
More CPU cores, , at least one core for one OSD	Same performance, less power consumption	Compatible with aarch64, support for hybrid deployment with x86 clusters



Customer's existing infrastructure



Mixed deployment with x86 clusters



Same operations for adding or removing OSD nodes with x86 cluster, smooth expansion experience.

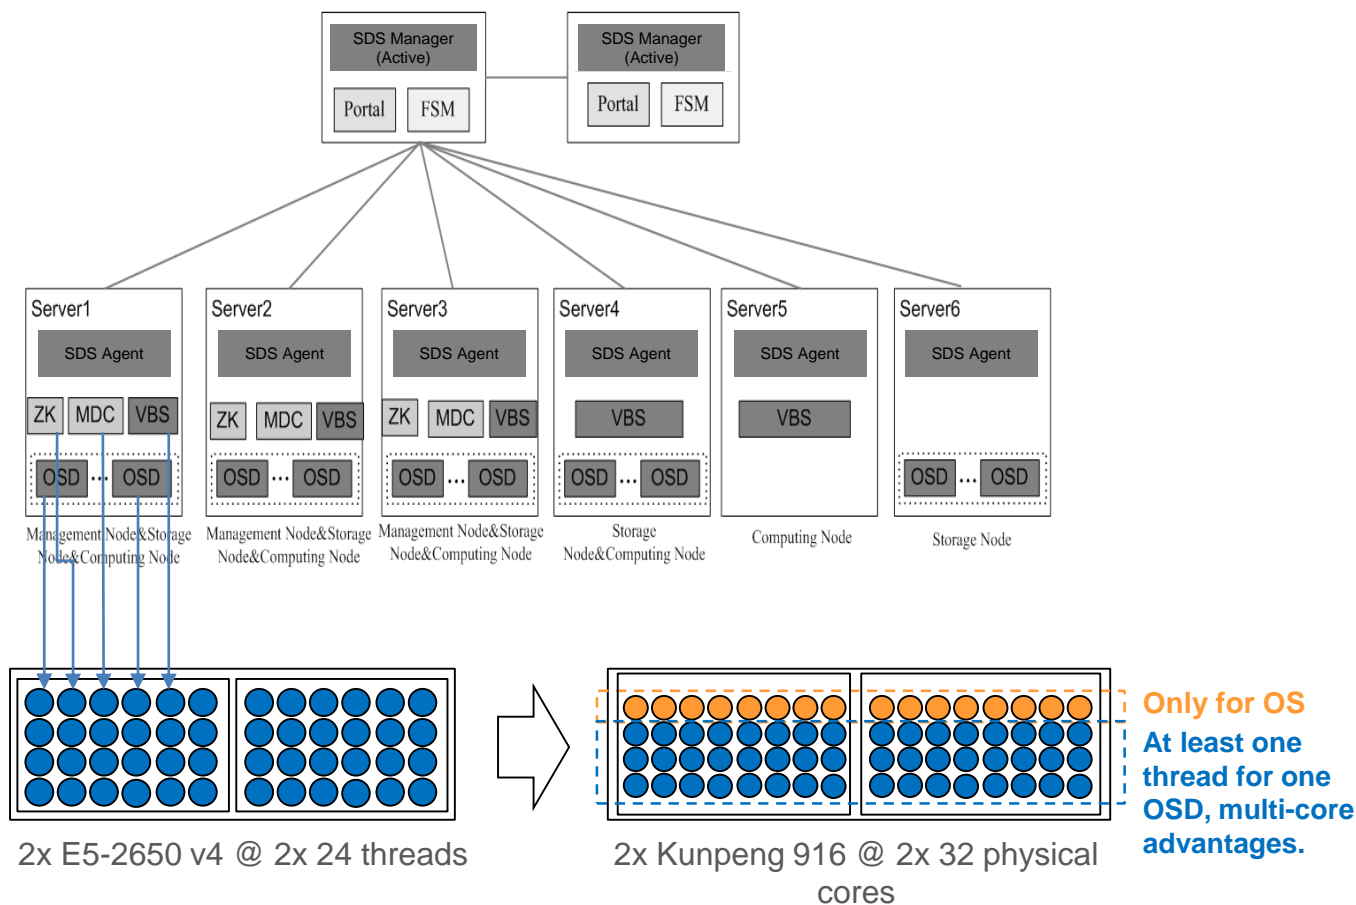
Expansion configuration - Storage



- CEPH storage solution –Expansion of OSD nodes
- TaiShan SDS solution configuration
 - 2 * Kunpeng 916 (2.4G/32C)
 - 2 * ES3620S V3 800G SAS SSD, RAID-1, for OS和Journal
 - 1 * ES3600C V3 800G NVMe, to accelerate metadata operations
 - 14 * 10TB NL-SAS HDD, for data storage

Customer's test result

Multi-core architecture




Low power consumption

	E5-2650 v4	Kunpeng 916
Vertical Segment	Server	Server
Processor Number	E5-2650 v4	Kunpeng 916
# of Cores	12	32
# of Threads	24	32
Processor Frequency	2.2GHz	2.4GHz
TDP	105W	75W
Process	14nm	16nm



Customer's test result



Contents

01 Ecosystem & Cases

02 TaiShan Products

- TaiShan Server
- TaiShan Solution



Kunpeng 920: Data center high-performance processor

Huawei Unveils Kunpeng 920 Processor & TaiShan Servers




Jan 7th, 2019, William Xu, Director of the Board and Chief Strategy Marketing Officer of Huawei, unveils Kunpeng 920 Processor

High Performance

SPECint®_rate_base2006 estimated score
930+

SPECint®_rate_base2006 estimated score




*Tested in Huawei lab. Results may vary in different environments.


High integration

4 chips in 1

**7nm
technology**



Kunpeng Processor Specifications



Kunpeng 916 Processor

- **32** cores, **2.4** GHz
- Mainstream SPECint performance, **75 W** Low TDP
- **4** memory channels
- PCIe 3.0 and SAS/SATA 3.0
- Integrated **GE/10GE LOM**
- Supports **2-socket** interconnects



Kunpeng 920 Processor

New

- **64** cores, **2.6** GHz
- **3x boost** in SPECint performance*
- **8** memory channels
- **PCIe 4.0** and CCIX
- Integrated **100GE LOM** and encryption and compression engines
- Supports **2- or 4-socket** interconnects

*Compared with Kunpeng 916 Processor


Kunpeng Processor Series for Different Application Scenarios

Kunpeng 920 64-core model

Ultimate
Performance



Industry's highest score at
SPECint®_rate_base2006




Kunpeng 920 32-/48-core model

Ultimate Energy
Efficiency



Higher SPECint performance
per unit power consumption




Kunpeng 916 32-core model


Ultimate Low Power




Lower power consumption
per unit SPECint performance




TaiShan Server Solutions




Big Data




SDS



ARM Native




HPC




Database

Kunpeng
Processor



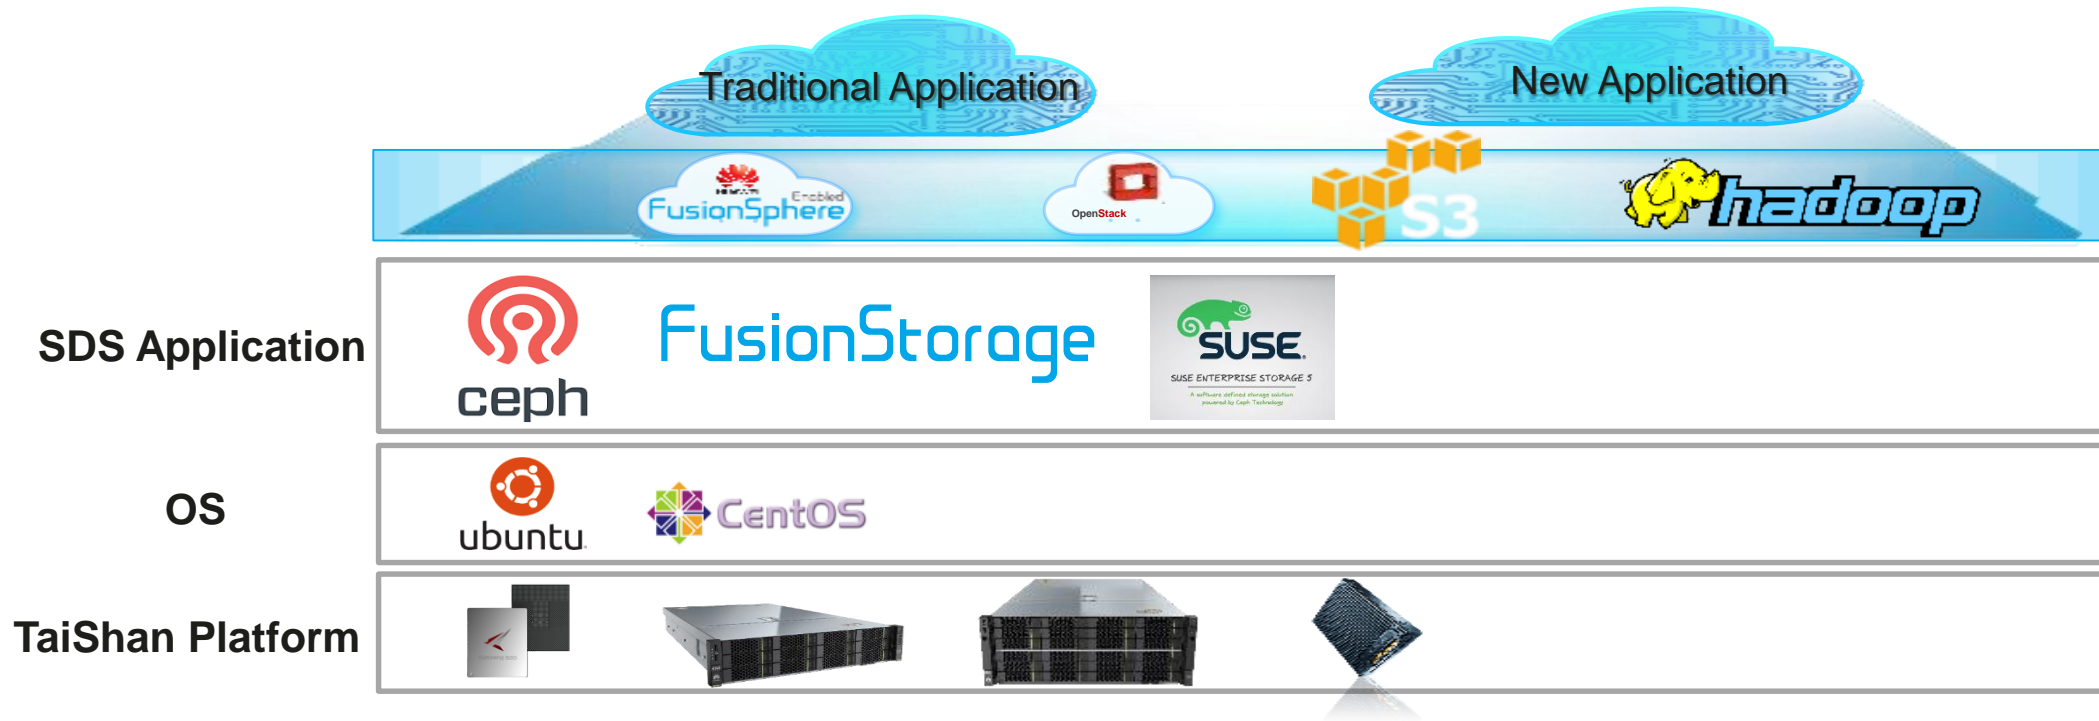
TaiShan
Server



OS





TaiShan Distributed Storage Solution



High Performance: Higher IOPS Performance for Hot Storage

Test Conditions & Hardware Configurations


Type	TaiShan + Ceph block storage	
Item	4 KB I/O random read/write, IOPS, and latency	
OS	CentOS 7.5 arch64	
Ceph	mimic 13.2.2	
Server	TaiShan 200 2280	x86 server
Processor	2*Kunpeng 920-4826	2*x86 6148
Memory	12 x 16 GB DDR4-2933 MHz	12 x 16 GB DDR4-2400 MHz
Network ports	25GE LOM	External 25GE NIC
OS drive	2 x 480 GB SATA SSDs	2 x 480 GB SATA SSDs
Data drive	12 x 3.2 TB NVMe SSDs	12 x 3.2 TB NVMe SSDs



* Tested in Huawei lab with TaiShan Server working on FusionStorage. Results may not be guaranteed with other test conditions

Dedicated Compression Acceleration Engine for Faster Data Compression

Kunpeng 920 with Built-In Data Compression Acceleration Engine



Test Conditions & Hardware Configurations

Algorithm	zlib	
Server	TaiShan 200 2280	x86 server
Processor	2*Kunpeng 920 -4826	2*x86 6148
Memory	12 x 16 GB DDR4-2933 MHz	12 x 16 GB DDR4-2400 MHz
Network ports	25GE LOM	External 25GE NIC
OS drive	2 x 480 GB SATA SSDs	2 x 480 GB SATA SSDs
Data drive	12 x 3.2 TB NVMe SSDs	12 x 3.2 TB NVMe SSDs




* Tested in Huawei lab with TaiShan Server working on FusionStorage. Results may not be guaranteed with other test conditions

Efficient Storage EC: Huawei FlexEC Algorithm Boosts Storage Utilization and Accelerates Data Rebuild


High Storage Utilization Rate

- Doubled storage utilization rate (from 33% to 66%) with 50% fewer nodes

3-Copy Solution (6+6+6)




Erasure Code Solution (6+3)



Fast Data Rebuild

- Huawei FlexEC algorithm improves rebuild speed by 35% and encoding speed by 10%

Erasure Code Solution (6+3) Reduces Data to Be Read for Rebuild



Bad or lost data



Rebuilding requested data with the FlexEC algorithm



Rebuilding requested data with a normal RS algorithm



HUAWEI

* Tested in Huawei lab with TaiShan Server working on FusionStorage. Results may not be guaranteed with other test conditions

TaiShan Ecosystem in Big Data Domain

Platform
Application

FusionInsight



Middleware



OS



Hardware



Kunpeng
Processor



TaiShan
Server



ES3000
NVMe SSD



Ascend
AI Chip



Atlas AI
Accelerator



Fusion Server Promotion

- Promotion Period: 01/09/2019 to 31/12/2019 with extendable ability
- Participant: Tier-1 Distributors
- Key Promotion products: 1U, 2U rack server,
- Promotion ID: PXXXXXXXXXX
- PCI modules, supplied from Huawei EU Supply Center from Hungary, 7-10 days typical lead time.
- PCI will be integrated into SCT with common code for each configuration to simplify ordering.
- Details refer to attachment.

Threshold (USD)	<50K	50K-0.5M	>0.5M
Rebate	0	1.50%	2%

Product code	Product	Description
02312LRA	PCI-1288H V5-4110	1 x Intel Xeon Silver 4110,1 x 1U 8*2.5 inch HDD Chassis, 1 x 2*GE & 2*10GE, .1 x DDR4 16GB 2666,1 x Platinum 550W, .1 xSR450C-M 2G Cache 12Gb/s,1 x Static rail
02312LRC	PCI-1288H V5-5118	1 x Intel Xeon Gold 5118,1 x 1U 8*2.5 inch HDD Chassis, 1 x 2*GE & 2*10GE, .1 x DDR4 32GB 2666,1 x Platinum 900W, .1 xSR450C-M 2G Cache 12Gb/s,1 x Static rail
02312LRN	PCI-1288H V5-3106-Upgrade	1 x Intel Xeon Bronze 3106 , 1 x 8*2.5 inch HDD Chassis, 1 x 2*GE & 2*10GE, . 1 x DDR4 16GB 2666, 2 x 240GB Sata SSD, 1 x Platinum 550W, , 1 x SR450C-M 2G(Avago3508) 12Gb/s, 1 x Static rail
02312LRP	PCI-1288H V5-4110-Upgrade	1 x Intel Xeon Silver 4110, 1 x 1U 8*2.5 inch HDD Chassis, 1 x 2*GE & 2*10GE.1 x DDR4 16GB 2666, 2 x 240Gb SATA SSD, 3 x 1.2TB SAS 10K, 1 x Platinum 550W, , 1 xSR450C-M 2G Cache 12Gb/s, 1 x Static rail
02312LRG	PCI-2288H V5-3106	1 x Intel Xeon Silver 3106,1 x 2U 8*2.5 inch HDD Chassis,1 x 2*GE & 2*10GE, .1x DDR4 16GB 2666,1 x Platinum 550W,1 xSR450C-M 2G Cache 12Gb/s,1 x Static rail
02312LRH	PCI-2288H V5-4110	1 x Intel Xeon Silver 4110,1 x 2U 8*2.5 inch HDD Chassis,1 x 2*GE & 2*10GE, .1x DDR4 16GB 2666,1 x Platinum 550W,1 xSR450C-M 2G Cache 12Gb, 1 x Static rail
02312LRK	PCI-2288H V5-5118	1 x Intel Xeon Gold 5118,1 x 2U 8*2.5 inch HDD Chassis,1 x 2*GE & 2*10GE, .2x DDR4 32GB 2666,2 x Platinum 900-W, .1 xSR450C-M 2G Cache 12Gb/s,1 x Static rail
02312LRM	PCI-2288H V5-4110-NVMe	2 x Intel Xeon Silver 4110,1 x 2U 8*2.5 inch HDD Chassis (support8*SAS/SATA&12*NVMe HDDs), 1 x 2*GE & 2*10GE, .2x DDR4 32GB 2666, 2 x Platinum 900W, .1 x SR450C-M 2G Cache 12Gb/s, 1 x Static rail
02312LRS	PCI-2288H V5-3106-Upgrade	1 x Intel Xeon Silver 3106,1 x 2U 8*2.5 inch HDD Chassis,1 x 2*GE & 2*10GE,.2 x DDR4 16GB 2666,2 x 240GB SATA SSD, 1 x Platinum 550W,1 xSR450C-M 2G Cache 12Gb/s,1 x Static rail
02312SUF	Promotion_1288H V5_4214_CLX_LFF	1 x Intel Xeon Silver 4214, 1* DDR4 16GB,2933MT/s, 1*SR430C-M 2G Cache 12Gb/s, LOM:2 x 10GE + 2 x GE ports, 1 x Platinum 550W
02312SUG	Promotion_1288H V5_4214_CLX_SFF	1 x Intel Xeon Silver 4214, 1* DDR4 16GB,2933MT/s, 1*SR430C-M 2G Cache 12Gb/s, LOM:2 x 10GE + 2 x GE ports, 1 x Platinum 550W
02312SUH	Promotion_1288H V5_5218_CLX_SFF	1 x Intel Xeon Gold 5218, 1*DDR4 32GB,2933MT/s, 1*SR430C-M 2G Cache 12Gb/s, LOM:2 x 10GE + 2 x GE ports, 1 x Platinum 550W
02312SUJ	Promotion_2288H V5_4214_CLX_LFF	1 x Intel Xeon Silver 4214, 1* DDR4 16GB,2933MT/s, 1*SR430C-M 2G Cache 12Gb/s, LOM:2 x 10GE + 2 x GE ports, 1 x Platinum 900W
02312SUK	Promotion_2288H V5_4214_CLX_SFF	1 x Intel Xeon Silver 4214, 1* DDR4 16GB,2933MT/s, 1*SR430C-M 2G Cache 12Gb/s, LOM:2 x 10GE + 2 x GE ports, 1 x Platinum 900W
02312SUM	Promotion_2288H V5_5218_CLX_LFF	1 x Intel Xeon Gold 5218, 2* DDR4 32GB,2933MT/s, 1*SR430C-M 2G Cache 12Gb/s, LOM:2 x 10GE + 2 x GE ports, 2 x Platinum 900W
02312SUN	Promotion_2288H V5_5218_CLX_SFF	1 x Intel Xeon Gold 5218, 1* DDR4 32GB,2933MT/s, 1*SR430C-M 2G Cache 12Gb/s, LOM:2 x 10GE + 2 x GE ports, 1 x Platinum 900W
	Components	Description(Has to be bundled with above typical configuration)
	Power Supply	PAC550S12-BE AC power supply unit
	Power Supply	PAC900S12-BE AC power supply unit (Wall Plug and PDU Style,Europe,Europe(PDU))
	Base Configuration	SR230-M,Avago3004,M.2 RAID PCIE card-RAID0,1,JBOD,Support Sideband Management,PCIE 3.0 X8,No Cache,HH/HL
	Base Configuration	M.2 SSD-SATA 6Gb/s-64GB,Hot-Swappable
	CPU	Intel Xeon Bronze 3106(1.7GHz/8-core/11MB/85W) Processor (with heatsink)
	CPU	Intel Xeon Silver 4110(2.1GHz/8-core/11MB/85W) Processor (with heatsink)
	CPU	Intel Xeon Gold 5118(2.3GHz/12-core/16.5MB/105W) Processor (with heatsink)
	CPU	Intel Xeon Silver 4214(2.2GHz/12-Core/16.5MB/85W)Cascade lake Processor (with heatsink)
	CPU	Intel Xeon Gold 5218(2.3GHz/16-Core/22MB/125W)Cascade lake Processor (with heatsink)
	MEMORY	DDR4 RDIMM Memory,16GB,2666MT/s,2Rank(1G*8bit),1.2V,ECC
	MEMORY	DDR4 RDIMM Memory,32GB,2666MT/s,2Rank(2G*4bit),1.2V,ECC
	MEMORY	DDR4 RDIMM Memory,16GB,2933MT/s,2Rank(1G*8bit),1.2V,ECC
	MEMORY	DDR4 RDIMM Memory,32GB,2933MT/s,2Rank(2G*4bit),1.2V,ECC
	HDD	HDD,4000GB,NL SAS 12Gb/s,7.2K rpm,128MB,3.5inch(3.5inch Drive Bay)
	HDD	HDD,4000GB,SATA 6Gb/s,7.2K rpm,128MB cache or above,3.5inch(3.5inch Drive Bay)
	SSD	SSD,240GB,SATA 6Gb/s,Read Intensive,PM883 Series,2.5inch(3.5inch Drive Bay)
	SSD	SSD,480GB,SATA 6Gb/s,Read Intensive,PM883 Series,2.5inch(3.5inch Drive Bay)
	SSD	SSD,240GB,SATA 6Gb/s,Read Intensive,PM883 Series,2.5inch(2.5inch Drive Bay)
	SSD	SSD,480GB,SATA 6Gb/s,Read Intensive,PM883 Series,2.5inch(2.5inch Drive Bay)
	SSD	ES3600S V5 SSD,800GB SAS 12Gb/s,Mixed Use,2.5inch(2.5inch Drive Bay)
	HDD	HDD,1800GB,SAS 12Gb/s,10K rpm,128MB,2.5inch(2.5inch Drive Bay)
	Flexible NIC	SM212 Onboard NIC,4xGE Electrical Interface(I350),RJ45
	PCle Card-NIC	Ethernet Adapter,1Gb Electrical Interface(Intel I350),2-Port,RJ45,PCle 2.0 x4
	PCle Card-NIC	Ethernet Adapter,1Gb Electrical Interface(Intel I350),4-Port,RJ45,PCle 2.0 x4

Thank you.

把数字世界带入每个人、每个家庭、
每个组织，构建万物互联的智能世界。

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

**Copyright©2018 Huawei Technologies Co., Ltd.
All Rights Reserved.**

The information in this document may contain predictive statements including, without limitation, statements regarding the future financial and operating results, future product portfolio, new technology, etc. There are a number of factors that could cause actual results and developments to differ materially from those expressed or implied in the predictive statements. Therefore, such information is provided for reference purpose only and constitutes neither an offer nor an acceptance. Huawei may change the information at any time without notice.

