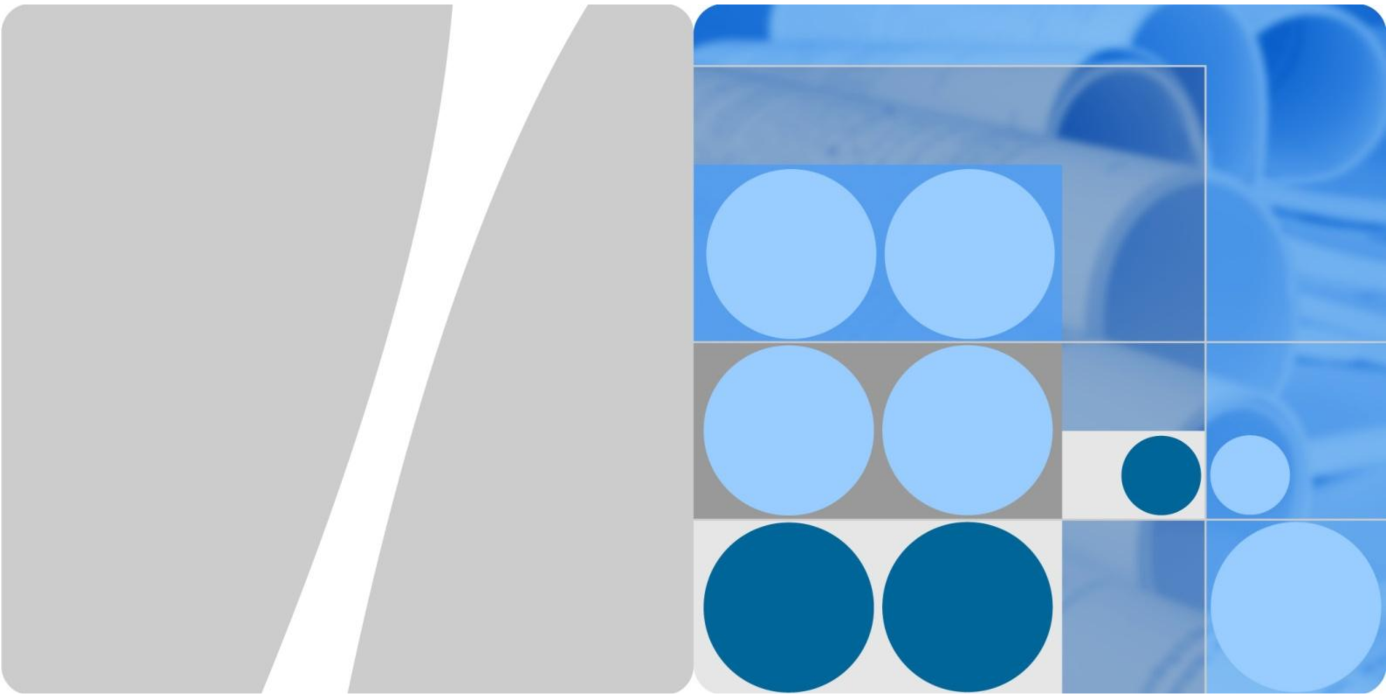

Technical White Paper on Huawei FusionCube v2.02 Database Infrastructure

Issue 1.0

Date 2014-07-17

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions

Notice

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: http://enterprise.huawei.com

About This Document

Change History

Date	Issue	Description	Author
2013-08-17	1.0	Completed the first draft.	Gu Qiang 00175014
2014-07-01	2.0	Updated on the basis of the first draft.	Xiong Wei 00116490

Contents

About This Document ii

1 Introduction 1

2 FusionCube Database Infrastructure Solution 2

2.1 System Architecture 2
2.1.1 Centralized Management 3
2.1.2 In-Depth Convergence 4
2.1.3 Elastic Capacity Expansion 5
2.1.4 Integration Before Delivery 5
2.2 Distributed Storage Software 6
2.2.1 Distributed Engine 7
2.2.2 Distributed Cache 7
2.2.3 Thin Provisioning 8
2.3 High-Performance SSD Storage 9
2.4 High-Speed IB Switching 10
2.5 Network Reliability Design 10

3 Computing and Storage Integrated Hardware 12

3.1 E9000 Subrack 12
3.2 E9000 Blade 13
3.3 10GE Switching 16
3.4 IB Switching 18

4 Configuration of FusionCube Database Infrastructure 19

4.1 FusionCube DX1 Specifications 20
4.2 FusionCube DX4 Specifications 21
4.3 Specifications of FusionCube for SAP HANA 22

1 Introduction

Today, more and more enterprises are using virtualization and cloud computing technologies to build their IT systems. Though being able to provide much higher operation efficiency than traditional IT systems, cloud computing still encounters the following challenges:

- Infrastructure platform deployment and management are complicated, which pushes operation and maintenance (O&M) costs up.
- The planning, deployment, and optimization skill requirements on operation personnel are high because hardware and software may be provided by different vendors.
- After-sales problems must be handled on operation portals of different vendors. Therefore, the response to problems may be slow.
- The maintenance system is huge because hardware and virtualization platforms from different vendors are involved in management and maintenance.
- Cloud computing has advantages only in large-scale database virtualization scenarios.
- Capacity expansion is not flexible. Most database appliances in the industry do not support configuration customization, on-demand deployment of computing and storage resources, and multiple databases.

Customers attach importance to cost control, rapid service provisioning, and risk management. They want a resource-scalable, reliable, and high-performance system with low total cost of ownership (TCO).

Huawei FusionCube Database Infrastructure solution integrates computing, storage, and network resources, featuring high system performance, automatic and rapid service deployment, reliable and secure service operation, unified O&M, and resource scaling. These features satisfy customers' long-term service development requirements and help them win rigorous market competitions.

This document describes the architecture, hardware and software, and product portfolios of FusionCube. This document is intended for sales engineers, channel sellers, senior service managers, and customers who intend to use FusionCube to deploy database-related infrastructures.

2 FusionCube Database Infrastructure Solution

2.1 System Architecture

Figure 2-1 shows the architecture of Huawei FusionCube Database Infrastructure system.

Figure 2-1 Architecture of Huawei FusionCube Database Infrastructure system

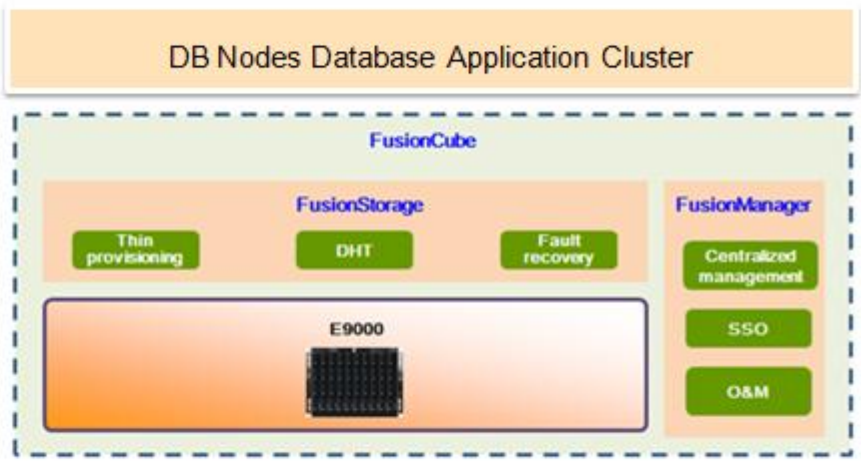

Huawei FusionCube is a flagship product of Huawei's cloud computing product line. Complying with an open architecture, FusionCube uses a 12 U subrack housing blade servers, distributed storage devices, and network switches, requiring no external storage device or switch. Virtual resources provided by FusionCube can be provisioned or expanded on demand.

FusionManager: provides a unified web-based user interface (UI) for centralized resource management and O&M, including user and cluster management, unified hardware monitoring, alarm reporting, storage resource configuration, and interconnection with northbound alarm interfaces.

FusionStorage: uses distributed cloud storage technologies to virtualize local hard disks on multiple blade servers into a resources pool, uses the scale-out architecture to eliminate the bottlenecks that may be caused by traditional engines, and provides high-performance, high-reliability block storage services.

Hardware platform: uses E9000 servers that support computing, storage, switch, and power modules deployed as separated modules in a subrack. All computing and storage modules can be flexibly configured based on service requirements. This platform also supports SAS hard disks, PCIe solid-state drive (SSD) cards, and 10GE and InfiniBand (IB) switching networks.

Integrating Huawei-developed hardware platform and distributed storage, FusionCube has the following characteristics:

- High performance
  - Hardware platform: provides high-density computing nodes and high-speed intra-subrack switch modules. The 12 U subrack supports 32 CPUs and 6 TB memory, the backplane provides 15.6 Tbit/s switching capability, and the outgoing bandwidth for a single subrack reaches 160 Gbit/s.
  - High-speed IB network: exchanges data between computing nodes and storage nodes over a 56 Gbit/s IB network, ensuring the high-speed cache synchronization between database nodes and high-speed I/O access on the storage plane network.
  - High-performance SSD storage: uses traditional SSD cards to replace traditional mechanic hard disks to provide 30 times higher IOPS.
  - Distributed storage: uses Huawei-developed FusionStorage distributed storage software that has the scale-out architecture, is equipped with non-volatile dual in-line memory modules (NVDIMMs) or SSDs as the high-speed cache, and closely interworks with hardware to thoroughly eliminate I/O bottlenecks.
- Excellent reliability
  - Computing reliability: Database applications are clustered. Therefore, the fault of a single computing node does not interrupt services.
  - Storage reliability: The FusionStorage system stores two or three copies for all data so that the data can be always accessible regardless of any node fault. The system automatically rebuilds data and restores multiple copies if a node or storage disk becomes faulty. It takes less than 30 minutes to rebuild 1 TB data. The strong-consistency replication protocol ensures that data copies are the same and protects data against power failures.
  - Network reliability: Both the external 10GE network and internal IB network work in dual-plane redundancy mode.
- Centralized management
  - FusionCube provides a unified portal for centrally configuring and managing computing, storage, and network resources.
  - FusionCube automatically monitors IT resources and system running status, displays alarms for all components, and supports alarm notification by email.
  - Computing and storage nodes can be expanded on demand by adding disks, blades, and subracks.

2.1.1 Centralized Management

The centralized O&M system of FusionCube is implemented by FusionManager. FusionManager manages hardware devices and resources pools. It also provides a unified portal for the lifecycle management of various resources in the system.

FusionManager provides the following functions:

Unified resource management: FusionManager consolidates resources of different systems and establishes a unified resource management model to provide unified resource management, such as monitoring and alarm reporting on various resources.

Unified alarm monitoring: Compared with traditional reference architectures in which alarms generated for servers, storage devices, and switches are presented by the monitoring systems of their own vendors, FusionCube centrally monitors the status and alarms of all components and allows alarm notification by email or alarm reporting to the customer's management platform using simple network management protocol (SNMP) northbound interfaces.

Role-based Access Control (RBAC): RBAC applies to enterprise security management. RBAC creates a role collection between users and rights. Each role has a group of rights. Once a role is defined for a user, the user can perform all the operations authorized to the role. The feature simplifies user management because roles change less frequently than rights.

Automatic hardware discovery and capacity expansion: FusionManager automatically discovers all hardware in FusionCube and scans the hardware specifications. Users only need to select the discovered hardware when creating a resource pool. FusionCube automatically installs software for the selected hardware, configures the network, and adds the hardware to the resource pool. During capacity expansion, FusionCube can also automatically discover the newly added hardware and complete configuration when you add the hardware to the resource pool.

2.1.2 In-Depth Convergence

Figure 2-2 shows the internally integrated hardware of FusionCube Database Infrastructure.

Figure 2-2 Hardware used in FusionCube Database Infrastructure

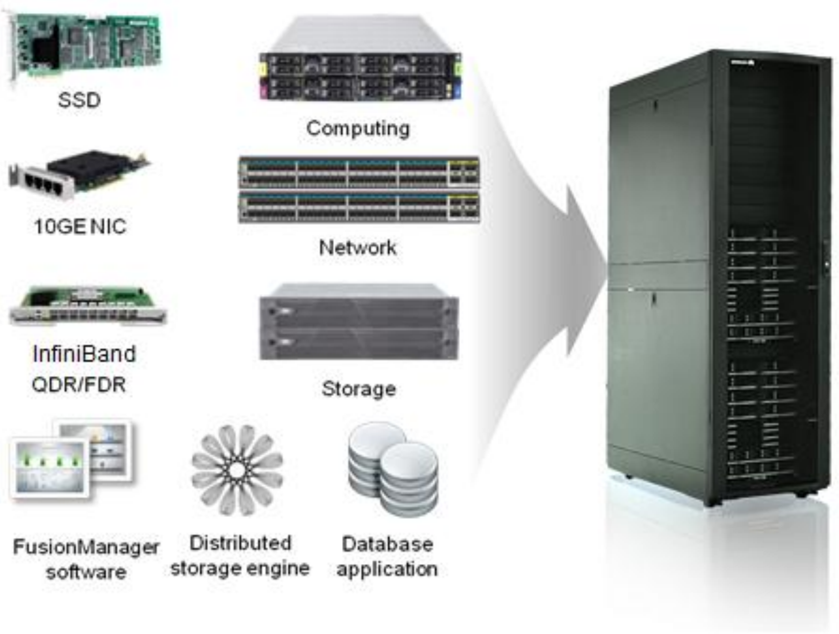

Huawei FusionCube Database Infrastructure uses a module-based design so that users can customize product profiles based on service requirements. The computing, storage, and network devices integrated in a FusionCube system are all developed by Huawei, providing optimized infrastructure and application performance. The converged platform also provides high-bandwidth, short-delay, multi-protocol-supported switching capabilities, NoSAN, and built-in distributed storage software. Not only hardware, the FusionCube system also integrates all the required software before delivery.

2.1.3 Elastic Capacity Expansion

FusionCube supports elastic capacity expansion based on service requirements without interrupting running services.

- If servers in a cluster are insufficient, users can add computing server blades to the cluster.
- If the storage resources in a cluster are insufficient, users can add storage server blades or independent hard disks to the cluster.
- If the resources in one subrack are insufficient, users can simply add subracks to the cluster and route cables.

FusionCube supports linear cluster capacity expansion without service interruption. For each cluster, the computing and storage resources are independent, but the O&M and service provisioning are implemented through one FusionManager system in a centralized manner.

2.1.4 Integration Before Delivery

Huawei FusionCube has all its components installed before delivery, thereby simplifying onsite installation and commissioning, and shortening the commissioning duration from several weeks or even months to only days.

The following measures implement FusionCube integration before delivery:

- Hardware installation: Devices are installed in the cabinets, and cables are bundled properly.
- Integrated device commissioning: System configuration integrity and connectivity are checked to ensure that system configurations meet contract requirements and all internal components can communicate with each other.
- Cabinet shipment: The cabinet is assembled and shipped for delivery.

After FusionCube is delivered onsite, users can implement a hardware check, ordered-specifications compliance check, connectivity check, and service configuration and commissioning for the system. Then, FusionCube is ready to work.

Figure 2-3 shows the preintegration and delivery process of FusionCube.

Figure 2-3 Preintegration and delivery process of FusionCube

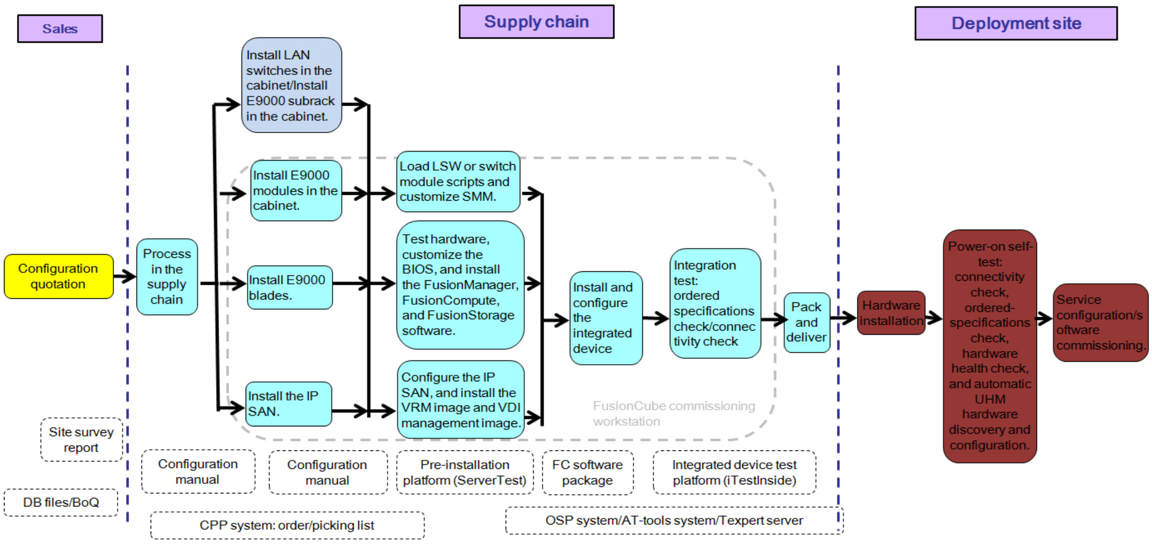

2.2 Distributed Storage Software

High-performance becomes a key requirement on enterprise storage systems. However, in recent years, ensuring long-term high-performance for all applications becomes especially difficult. It has been tested that traditional storage systems can no longer keep up with the demanding storage requirements and increasing data volume. Moreover, high-end storage devices that serve as independent NEs are not only expensive but also need dedicated management and maintenance.

Huawei-developed distributed storage software, FusionStorage, virtualizes local disks on servers or blades into a virtual storage resource pool and implements computing and storage integration. FusionStorage can completely take the place of external SAN devices in some scenarios. The integration of computing and storage in FusionStorage achieves high performance, high reliability, and high cost-effectiveness.

- **Advanced distributed architecture:** FusionStorage uses a distributed structure that features the distributed management clusters, distributed hash routing algorithm, distributed and stateless engines, and distributed intelligent caching. This structure can effectively prevent single points of failure (SPOF) in the entire storage system.
- **High-performance and high-reliability**: FusionStorage balances loads among all disks and dispersedly stores data, which helps ensure that no data hotspot exists in the system. Effective data routing algorithm and distributed caching mechanism help ensure high performance. Server memory can be used for cache read/write. FusionStorage supports multiple data copies. The data copies are stored on different servers/disks. Failures in a single hardware device have no adverse impact on services. When three data copies are stored, data availability reaches 99.999999999%. In addition, FusionStorage uses the strong-consistency replication protocol to ensure data consistency between data copies. Users can write data to only one of the copies but read the same data from any data copy. Cache data can be protected when a power failure occurs.

- **Rapid parallel data restoration upon failures:** Data fragments are dispersed in the resource pool. Upon a hard disk failure, these data fragments are automatically restored within the resource pool in parallel. For example, in a resource pool, it takes less than 30 minutes to restore 1 TB data and the CPU usage is less than 10%.
- **Easy expansion and ultra-large capacity:** The distributed stateless engines in the FusionStorage system support horizontal expansion, smooth and parallel capacity expansion of storage and computing resources, and non-chimney ultra-large capacity expansion.
- **Deep integration of computing and storage resources:** FusionStorage is deployed on servers that have local hard disks to virtualize all the local disks on the servers as a virtual resource pool. This resource pool integrates computing and storage resources of the servers and can function as an external storage device for the servers, saving costs for external storage devices and facilitating unified O&M.

2.2.1 Distributed Engine

FusionStorage uses distributed stateless engines. These engines are deployed on each server that requires access to FusionStorage, thereby preventing performance bottlenecks that may be caused by traditional, centrally-deployed engines. Moreover, these distributed engines deployed on standalone servers consume much fewer CPU resources and provide more I/O than centrally-deployed engines do. The engines are also expanded when storage blades are added, which totally eliminates I/O bottlenecks.

2.2.2 Distributed Cache

FusionStorage integrates computing and storage resources and evenly distributes caches and bandwidths to each server node.

Each disk on FusionStorage servers uses independent I/O bandwidth, preventing a large number of disks competing limited bandwidths between computing devices and storage devices in an independent storage system.

FusionStorage uses some memory of each server for read cache and non-volatile dual in-line memory module (NVDIMM) or solid-state drive (SSD) for write cache. Caches are evenly distributed to all nodes. The total cache size on all servers is greatly larger than that provided by external storage devices. Even when using the large-capacity and low-cost SATA disks, FusionStorage can still provide one to three times higher I/O performance and larger effective capacity.

Figure 2-4 shows the comparison between FusionStorage and external SAN storage.

Figure 2-4 Comparison between FusionStorage and external SAN storage

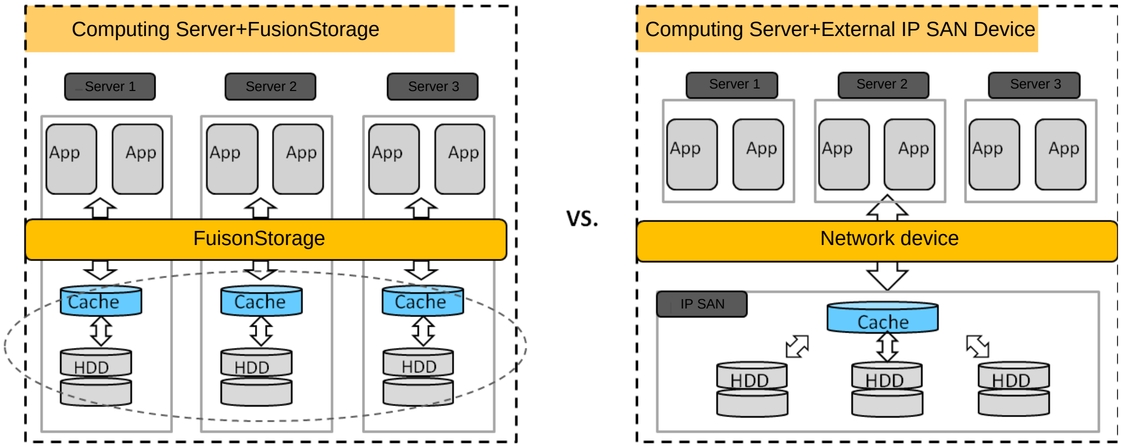

Table 2-1 lists the performance comparison between FusionStorage and external SAN storage.

Table 2-1 Performance comparison between FusionStorage and external SAN storage

Item	FusionStorage	External Storage System
Bandwidth (between disks and cache or between flash and apps)	Each disk provides 100 Mbit/s local bandwidth.	Each disk provides 20 Mbit/s bandwidth (back-end SAS disks: 6 Gbit/s; usage: 80%).
I/O latency	Access to the data stored on a local disk: direct access to the local disk	Access to the data stored on a local disk: A TCP/IP switching is required.
	Access to the data stored on a remote disk: A TCP/IP switching is required.	Access to the data stored on a remote disk: A TCP/IP switching is required. A PCIe switching is also required when multiple engines are configured.
Distributed large-capacity cache	Each disk provides more than 250 MB cache, which can be expanded by adding blades.	All disks share the cache.

2.2.3 Thin Provisioning

FusionStorage provides the thin provisioning feature, which allows users to use much more physical storage resources than real physical storage resources. In the data center serving multiple users, this feature helps to save initial investment on storage devices, because users do not consume all the allocated storage resources initially. This feature significantly improves storage utilization. FusionStorage uses the distributed hash technology to support thin provisioning, without pre-allocating space. Thin provisioning effectively prevents performance deterioration that may be caused on IP SAN devices due to capacity expansion. In Figure 2-5, dotted lines indicate that those volumes are not allocated.

Figure 2-5 Thin provisioning principles

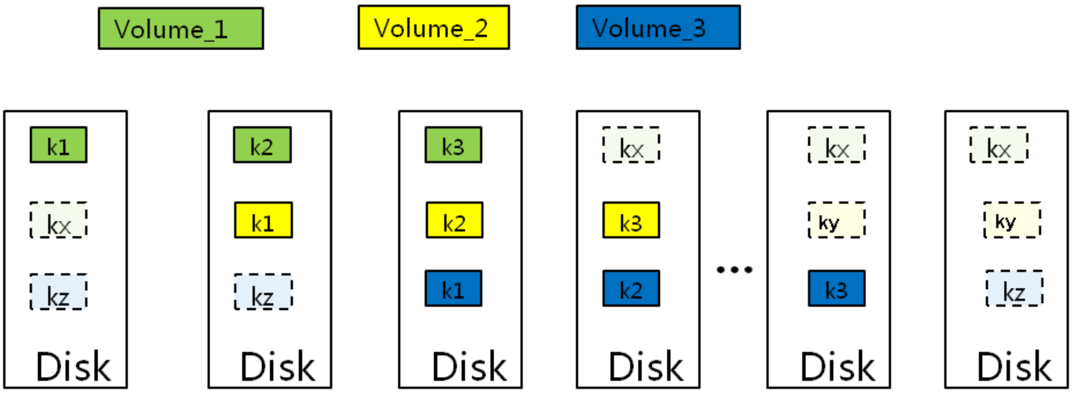

2.3 High-Performance SSD Storage

By combining with the high-speed I/O performance of flash storage, FusionCube can greatly improve storage performance and eliminate storage I/O bottlenecks. The data warehouse applications with massive data analysis impose high requirements for storage throughput. The PCIe SSD on FusionCube communicates directly with the CPU, thereby shortening the access path for data storage. Huawei FusionCube supports a distributed SSD storage system for high-performance applications and provides higher I/O performance than traditional hard disks, such as SATA/SAS disks.

Huawei SSD cards provide the following characteristics:

- High bandwidth: uses PCIe 2.0 x8 interfaces, providing 3 Gbit/s read/write bandwidth.
- High I/O: supports 4 KB data blocks and 100% random I/O, and provides 600,000 continuous random read IOPS and 220,000 continuous write IOPS.

  The distributed SSD storage system takes the following measures to enhance system reliability:

- The built-in error checking and correcting (ECC) engine and the RAID 5 engine check and correct data errors between data channels.
- The built-in data scrubbing engine periodically checks data that has been stored to prevent data errors.
- A dynamic RAID algorithm is used to share resources among channels. This ensures that the system is functioning properly even if multiple chips fail.
- Super capacitors provide power-off protection.

  FusionStorage categorizes and manages frequently accessed and less frequently accessed data and uses an advanced SSD wear algorithm to maximize the space reclamation rate, minimize the write-wear rate, and extend the life of the SSDs.

2.4 High-Speed IB Switching

FusionCube uses the industry-leading IB network that is based on the edge-cutting remote direct memory access (RDMA) technology, which enhances the database throughput, speeds up response to requests, and improves user experience.

FusionCube uses IB networking for the communication between storage nodes and database nodes (computing nodes), thereby implementing dual-link redundancy and ensuring system reliability.

FusionCube delivers the following characteristics using the high-speed IB network:

- Provides a data rate of 56 Gbit/s and allows high-speed connection setup.
- Uses standard multilayer fat-tree networking and allows smooth capacity expansion.
- Provides a communication network where congestion hardly occurs and avoids data switching bottlenecks.
- Provides minor communication delays within microseconds and transmits computing and storage information promptly.
- Provides lossless network QoS and ensures data integrity during transmission.
- Allows multi-path communication for active and standby ports and ensures communication path redundancy.

Figure 2-6 Rate comparison of different networks

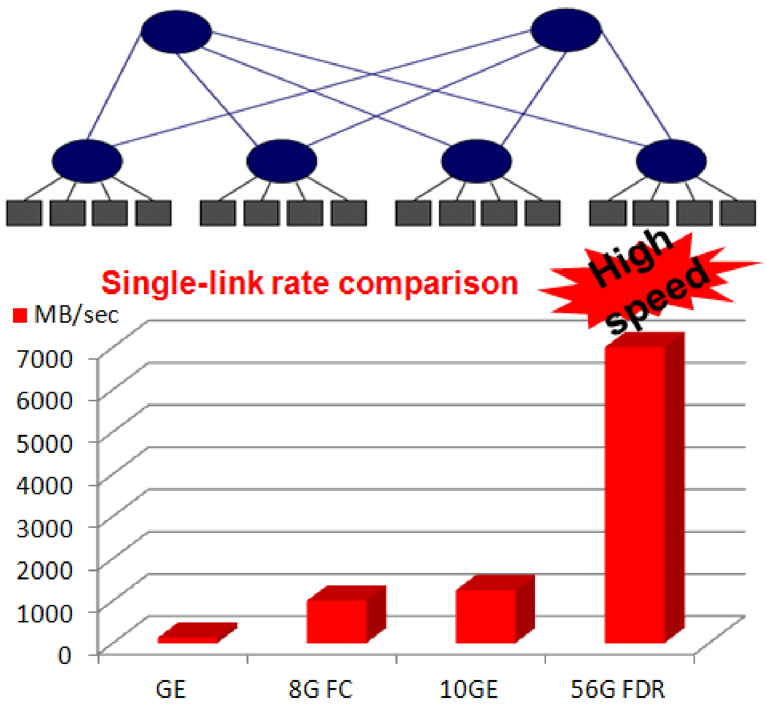

2.5 Network Reliability Design

The system networks are logically divided into the service plane, storage plane, and management plane, which are isolated to protect the basic platform.

Storage plane: The storage devices on the server communicate with each other at layer 2 over the storage network. Storage devices provide storage resources for VMs through the small computer system interface (SCSI) but do not communicate with VMs.

Service plane: It provides a channel for users to obtain services, for virtual NICs of VMs or computing blades of databases to communicate with each other, and for external applications to interact with the desktop cloud system. Access can be isolated by VLANs configured for VMs.

Management plane: It provides various functions, such as system management, service deployment, and system loading of the entire cloud computing system. In Database Infrastructure scenarios, the management plane cannot be accessed directly by external devices. The baseboard management controller (BMC) plane manages hardware on the server and the emergency channel when the network fails. It is isolated from the management plane.

Each blade server provides service ports and management ports to ensure link redundancy by NIC binding. The E9000 server has two 10GE switch modules in its subrack. They serve as the access and aggregation switches for connecting to all server blades in the subrack, cascading subracks, communicating with customer's network in the uplink.

Two 56 Gbit/s IB network ports are deployed for the storage plane to implement dual-link redundancy, thereby enhancing system reliability. The E9000 server has two IB switch modules in its subrack that connect to all server blades in the subrack and implement subrack cascading on the storage plane.

3 Computing and Storage Integrated Hardware

3.1 E9000 Subrack

Tecal E9000 (E9000 for short) is a 12-U server developed by Huawei. It can house E9000 series computing nodes, storage nodes, and capacity expansion nodes. The E9000 server has the following characteristics:

- A subrack can horizontally house 8 full-width computing nodes or 16 half-width computing nodes and supports slot division.
- A subrack supports 850 W heat dissipation for half-width slots.
- A subrack supports 1700 W heat dissipation for full-width slots.
- A half-width slot supports a maximum of two-socket CPUs and 24 DIMMs.
- A full-width slot supports a maximum of four-socket CPUs and 36 DIMMs or 48 DIMMs.
- A single subrack supports a maximum of 32 CPUs and 12 TB memory.
- The backplane supports a maximum switching capability of 15.76 Tbit/s.
- Four (two pairs of) switching slots support multiple types of switching protocols, such as Ethernet, FC, FCoE, and IB, and support straight-out I/O interfaces.

Figure 3-1 shows the appearance of the E9000 server.

Figure 3-1 Appearance of the E9000 server

3.2 E9000 Blade

FusionCube supports the following blades: 2-socket CH121 computing blade, 2-socket CH222 storage blade, 2-socket CH223 storage I/O expansion blade, 4-socket CH240 blade, and 4-socket CH242 blade. Other blades are not described in this document.

Figure 3-2 shows the appearance of the CH121 blade.

Figure 3-2 Appearance of the CH121 blade

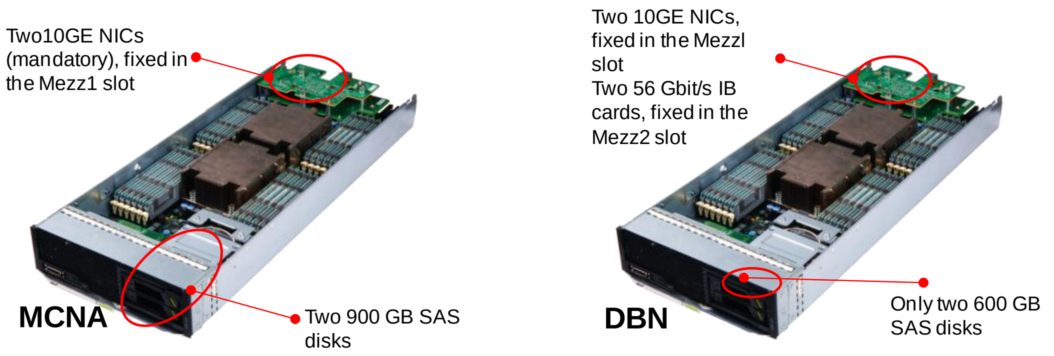

The CH121 can be used as the management computing node agent (MCNA) node or database (DB) node. It has the following characteristics:

1. Uses two Intel® Xeon® E5-2600 processors.
2. Provides 24 DDR3 DIMM slots and 2 x 2.5-inch SAS disks.
3. Provides two Mezz cards that can support the 10GE or IB network.

Figure 3-3 shows the appearance of the CH222 storage blade.

Figure 3-3 Appearance of the CH222 storage blade

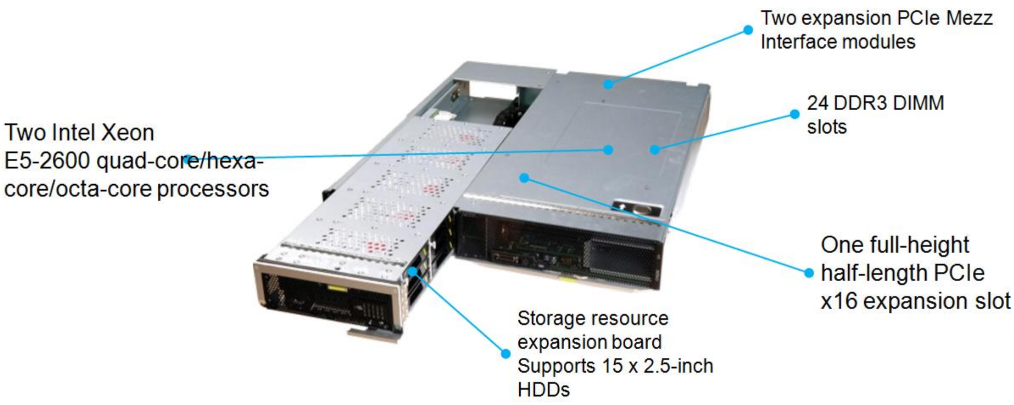

The CH222 blade is used as the DB node if large storage capacity is required. It provides the following characteristics:

1. Uses two Intel® Xeon® E5-2600 processors.
2. Provides 24 DDR3 DIMM slots.
3. Supports a maximum of fifteen 2.5-inch SSD, SAS, or SATA hard disks.
4. Provides two Mezz cards that can support the 10GE or IB network.
5. Supports hot swap to allow hard disks to be replaced without powering off the system.

Figure 3-4 shows the appearance of the CH223 storage blade.

Figure 3-4 Appearance of the CH223 storage blade

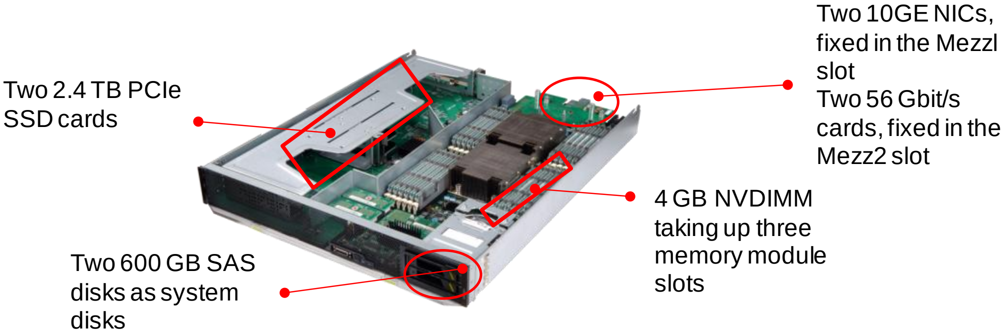

The CH223 blade is used as the storage node in high-performance database scenarios and provides the following characteristics:

1. Uses two Intel® Xeon® E5-2600 processors.
2. Provides 24 DDR3 DIMM slots.
3. Provides two PCIe x16 full-height full-length slots or four PCIe x8 full-height half-length slots, holding 2 x 2.4 SSD cards in database scenarios.
4. Provides two Mezz cards that can support the 10GE or IB network.

Figure 3-5 shows the appearance of the CH240 blade.

Figure 3-5 Appearance of the CH240 blade

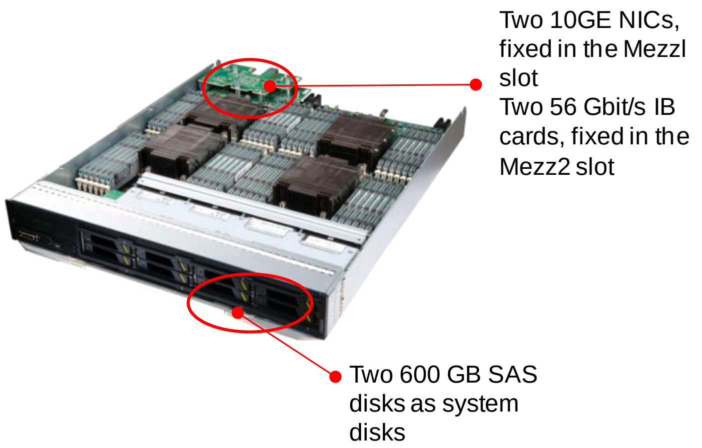

The CH240 blade is used as the computing node for the database that requires high computing capacity. It provides the following characteristics:

1. Uses four Intel® Xeon® E5-4600 processors.
2. Provides 48 DDR3 DIMM slots.
3. Supports eight 2.5-inch SSD, SAS, or SATA disks.
4. Provides two Mezz cards that can support the 10GE or IB network.

Figure 3-6 shows the appearance of the CH242 blade.

Figure 3-6 Appearance of the CH242 blade

The CH242 blade is used as the computing node for the database that requires high computing capacity or HANA database. It provides the following characteristics:

1. Uses four Intel® Xeon® E7-4800 processors.
2. Provides 32 DDR3 DIMM slots.
3. Supports eight 2.5-inch SSD, SAS or SATA hard disks.
4. Provides four Mezz card modules and one PCIe standard card. If common databases, such as Oracle database, are deployed, configure one 10GE and one IB NIC. If the HANA database is deployed, configure two MZ512 NICs and one IB NIC.

3.3 10GE Switching

FusionCube uses the CX310 switches to support 10GE networks for the service plane and the management plane. Each subrack is configured with two switches. Table 3-1 shows the switching performance indicators of the CX310.

Table 3-1 Performance indicators of the CX310

Item	Description
Appearance	

Item	Description
Number of ports on the panel	• One BMC serial port and one SYS serial port • Sixteen 10GE uplink Ethernet optical ports
Switching capability	1.28 Tbit/s
Packet forwarding rate	960 Mpps
Key data center features	• Supports TRILL. • Supports FCoE. • Supports Data Center Bridging Exchange Protocol (DCBX), packet flow context (PFC), and enhanced transport selection protocol (ETS).
VLAN	• Access, trunk, and hybrid modes • Default VLAN • QinQ VLAN and enhanced selective QinQ • MUX VLAN
QoS	• Flow classification based on the combination of layer-2, layer-3, layer-4, and 802.1p protocols • Access control list (ACL), committed access rate (CAR), Remark, and traffic scheduling • Priority queue (PQ), weighted round robin (WRR), deficit round robin (DRR), PQ+WRR, and PQ+DRR queue scheduling • Weighted random early detection (WRED) and tail drop congestion avoidance mechanisms • Traffic shaping
Security and management functions	• Hierarchical protection using commands, preventing unauthorized access • Defense against Denial of Service (DoS), Address Resolution Protocol (ARP), and Internet Control Message Protocol (ICMP) attacks • Port isolation and sticky MAC port security mechanisms • Binding between IP addresses, MAC addresses, ports, and VLANs • Various authentication modes, such as Authentication, Authorization and Accounting (AAA), Remote Authentication Dial In User Service (RADIUS), and Huawei Terminal Access Controller Access Control System (HWTACACS) • Remote network monitoring (RMON)

3.4 IB Switching

FusionCube uses the CX610 IB switches to support 56 Gbit/s data switching for the data plane. Each subrack is configured with two switches. Table 3-2 shows the switching performance indicators of the CX610.

Table 3-2 Performance indicators of the CX610

Item	Description
Appearance	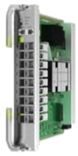
Number of ports on the panel	One BMC serial port and one SYS serial port Upstream: 18 QDR/FDR IB QSFP+ ports
Key data center features	Supports TRILL.

4 Configuration of FusionCube Database Infrastructure

Huawei FusionCube Database Infrastructure provides typical configuration of DX1 and DX4, and customized configuration of DXX. The flexible FusionCube Database Infrastructure solution meets various service load and storage requirements of different data centers. FusionCube Database Infrastructure supports various database software, such as Oracle RAC, Sybase IQ, DB2, HANA, and GBase. Oracle RAC and HANA databases have been proved and tested in actual projects.

FusionCube DX1 and DX4 components are listed as follows:

- DX1: one E9000 subrack, three two-socket computing nodes, and 6 storage nodes, with no external IB switch.
- DX4: four E9000 subracks, twelve two-socket or six four-socket computing nodes, and 24 storage nodes, with no external IB switch.

Users can choose computing nodes, storage nodes, CPUs, memory modules, and storage devices based on service requirements.

Table 4-1 lists the blades used in the FusionCube Database Infrastructure solution.

Table 4-1 Blades used by FusionCube Database Infrastructure

Blade Type	Description
Management node	CH121 or CH222 blades, used to deploy FusionManager and FusionStorage Manager
Computing node 1	CH121 blades, used to deploy database software and FusionStorage engines.
Computing node 2	CH240 blades, used to deploy database software and FusionStorage engines.
Computing node 3	CH242 blades, used to deploy database software and FusionStorage engines. Four-socket computing nodes preferentially use CH242 blades if there are no special constraints, for example, the cost and the number of cores specified in the software license.

Blade Type	Description
Storage node 1	CH223 blades, each housing two 2.4 TB SSD cards. All the storage blades constitute the storage resource pool. CH233 blades apply to scenarios where high performance is required.
Storage node 2	CH222 blades, each housing twelve 900 GB or 1.2 TB SAS disks and one SSD card for caching. All the storage blades constitute the storage resource pool. CH222 blades apply to scenarios where large capacity is required.

4.1 FusionCube DX1 Specifications

Table 4-2 lists the FusionCube DX1 specifications

Table 4-2 FusionCube DX1 specifications

Parameter	Specifications
Database	Oracle RAC 11g R2
Operating system (OS)	Oracle Linux 6.2 and RHEL 6.2
Computing node	Three CH121 blades, each with two Intel E5-2690V2 CPUs, 256 GB memory, two 900 GB SAS disks, and IB FDR and 10GE NICs
Storage node	Six CH223 blades, each with two Intel E5-2630V2 CPUs, 64 GB memory, 4 GB NVDIMM, two 300 GB SAS disks, two 2.4 TB PCIe SSD cards, and IB and 10GE NICs
	CH222 blade, with one Intel E5-2630V2 CPU, 64 GB memory, 4 GB NVDIMM, two 300 GB SAS disks, twelve 1.2 TB SAS disks, and IB and 10GE NICs
Management node	One CH121, with two Intel E5-2620 CPUs, 128 GB memory, two 900 GB SAS disks (configure only two 600 GB disks if cutover to DX4 or higher is not required), and 10GE NICs
Network interconnection	56 Gbit/s FDR IB and 10GE networks
Capacity expansion	Default capacity: two computing nodes and three storage nodes. One DB node and one storage node are added for each capacity expansion.
Storage type	Distributed storage FusionStorage
Storage capacity	28.8 TB SSD raw capacity (6 x 4.8 TB) or 86.4 TB SAS raw capacity (6 x 12 x 1.2 TB)
IOPS	600,000 (Max.)
Throughput	30 Gbit/s (Max.)

Parameter	Specifications
Storage latency	< 1 ms
Subrack	One 12 U subrack

4.2 FusionCube DX4 Specifications

Table 4-3 lists the FusionCube DX4 specifications.

Table 4-3 FusionCube DX4 specifications

Parameter	Specifications
Database	Oracle RAC 11g R2
OS	Oracle Linux 6.2 and RHEL 6.2
Computing node	Twelve CH121 blades, each with two Intel E5-2690V2 CPUs, 256 GB memory, two 900 GB SAS disks, and IB FDR and 10GE NICs
	Six CH242 blades, each with four Intel E5-4650 or Intel E7-4870 CPUs, 768 GB or 576 GB memory, two 900 GB SAS disks and IB FDR and 10GE NICs
Storage node	Twenty-four CH223 blades, each with two Intel E5-2630V2 CPUs, 64 GB memory, 4 GB NVDIMM, two 300 GB SAS disks, two 2.4 TB PCIe SSD cards, and IB and 10GE NICs
	CH222 blade, with one Intel E5-2609 CPU, 64 GB memory, 4 GB NVDIMM, two 300 GB SAS disks, twelve 1.2 TB SAS disks, and IB and 10GE NICs
Management node	One CH121 blade, with two Intel E5-2620 CPUs, 128 GB memory, two 900 GB SAS disks, and 10GE NICs
Network interconnection	56 Gbit/s FDR IB switching network and 10GE network, requiring no external IB switch
Capacity expansion	Default capacity: two two-socket or three four-socket computing nodes and eight storage nodes. One DB node and one storage subrack are added for each capacity expansion.
Storage type	Distributed storage FusionStorage
Storage capacity	115.2 TB SSD raw capacity (24 x 4.8 TB) or 345.6 TB SAS raw capacity (24 x 12 x 600 GB)
IOPS	2,400,000 (Max.)
Throughput	30 Gbit/s (Max.)
Storage latency	< 1 ms
Subrack	Two to four 12 U subracks

4.3 Specifications of FusionCube for SAP HANA

Table 4-4 lists the specifications of FusionCube for SAP HANA.

Table 4-4 Specifications of FusionCube for SAP HANA

Parameter	Specifications
Database	SAP HANA
OS	SLES for SAP 11 SP3
Memory	512 GB/1 TB
Number of HANA nodes	2–16
Computing node	Two to sixteen CH242 blades, each with four IvyBridge Intel E7-4880/4890 CPUs, 512 GB or 1 TB memory, two 600 GB SAS disks, and IB FDR and 10GE NICs
Storage node	Three to fifteen CH2222 blades, each with two Intel E5-2630V2 CPUs, 64 GB memory, 4 GB NVDIMM, twelve 900 GB 10k SAS disks (data disks), two 600 GB 10k SAS disks (system disks), and IB and 10GE NICs
	Fourteen CH222 blades and eight CH223 blades. Each CH223 blade is configured with two Intel E5-2630V2 CPUs, 64 GB memory, 4 GB NVDIMM, two 600 GB 10 k SAS disks, two 2.4 TB PCI-SSD cards and IB and 10GE NICs
Management node	One or two CH222 blades, each with two Intel E5-2620V2 CPUs, 128 GB memory, two 900 GB SAS disks, and 10GE NICs
Network interconnection	56 Gbit/s FDR IB and 10GE networks
Storage type	Distributed storage FusionStorage
Cabinet	One or two 42 U cabinets
Management	FusionManager 3.X