



Converged for Ultimate Simplicity and Performance FusionCube For Cloud



Copyright © Huawei Technologies Co., Ltd. 2018. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

Trademark Notice

 **HUAWEI**, and  are trademarks or registered trademarks of Huawei Technologies Co., Ltd. Other trademarks, product, service and company names mentioned are the property of their respective owners.

NO WARRANTY

THE CONTENTS OF THIS MANUAL ARE PROVIDED "AS IS". EXCEPT AS REQUIRED BY APPLICABLE LAWS, NO WARRANTIES OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, ARE MADE IN RELATION TO THE ACCURACY, RELIABILITY OR CONTENTS OF THIS MANUAL.

TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, IN NO CASE SHALL HUAWEI TECHNOLOGIES CO., LTD BE LIABLE FOR ANY SPECIAL, INCIDENTAL, INDIRECT, OR CONSEQUENTIAL DAMAGES, OR LOST PROFITS, BUSINESS, REVENUE, DATA, GOODWILL OR ANTICIPATED SAVINGS ARISING OUT OF OR IN CONNECTION WITH THE USE OF THIS MANUAL.

HUAWEI TECHNOLOGIES CO., LTD.

Bantian, Longgang District
Shenzhen 518129, P. R. China
Tel: +86-755-28780808

www.huawei.com

HUAWEI TECHNOLOGIES CO., LTD.



Solution Highlights

» Hyper-Converged Architecture for Diverse Service Infrastructure Platforms

Huawei FusionCube For Cloud is an IT platform based on a hyper-converged architecture. It complies with open architecture standards, converges compute and storage, and preintegrates distributed storage engines, virtualization, and cloud management software, allowing for resource on-demand provisioning and linear expansion.

FusionCube for Cloud supports different hardware configurations requirements of different services for compute, storage, and I/O resources. The FusionCube HCI applies to cloud computing scenarios and is an ideal choice for deploying the IT infrastructure of cloud data centers.



FusionCube 2000
2U 1-node

» Fast Service Rollout via Automated Installation and Deployment

The FusionCube For Cloud allows for one-stop delivery and home appliance-like installation via preintegration and preoptimization. Users can roll out services as soon as the system goes live, gaining unprecedented deployment speeds. O&M personnel can deploy services by simply clicking the mouse on the GUI. Simplified O&M allows you to focus more on service innovation and development.

» Agile and Convenient O&M over a Unified Platform

The unified management system of the FusionCube For Cloud allows access to all physical and virtual resources in FusionCube, including switches, virtual machines (VMs), and storage volumes in one or multiple chassis. In addition, the system management software supports automatic hardware discovery and configuration. To expand a cloud platform, you only need to add new nodes, connect their cables, and power them on. In this way, the FusionCube For Cloud improves resource utilization and reduces follow-up O&M costs by up to 30%. The easy O&M enforces lower requirements on IT management personnel.

» High-Speed Distributed Storage Engines Eliminate Performance Bottlenecks

The built-in distributed storage engines virtualize all local FusionCube For Cloud storage resources into a storage resource pool, a cluster consisting of multiple nodes in a distributed architecture. This cluster architecture makes the service system free from performance and bandwidth bottlenecks caused by a traditional RAID controller while providing high availability and scalability.

Data is striped and stored across all of the built-in hard drives or SSDs, improving both storage utilization and performance. Advanced storage features, such as thin provisioning and linked clone, enable future-expandable planning and on-demand linear capacity expansion, greatly reducing initial investment.

» Simplified Management

- Compute, storage, and network all-in-one, enabling one-stop delivery
- Hardware comes preinstalled, preconnected, and is transported with all components in position
- Preintegrates a cloud computing virtualization and management platform
- Automatic hardware discovery and parameter configuration
- Unified platform for managing physical and virtual resources

» Always-On Services

- Multiple data safeguarding mechanisms, delivering up to seven-nines (99.99999%) data reliability
- VMs capable of intelligent self-healing
- One of the industry's few first to support cabinet-level security; with Huawei-unique cabinet-level data distribution technology, tolerating cabinet-wide faults
- Automatically isolates faulty hard drives and rebuilds data

Typical Configuration

Model	FusionCube 2000			
Application scenarios	Virtualization, cloud, hybrid workloads, VDI			
Node specifications	2S compute node: (2 CPU, 24 DDR4 DIMMs, 2 GPUs) 4S compute node: (4 CPU, 48 DDR4 DIMMs) 2S compute storage node-AF: (2 CPU, 24 DDR4 DIMMs, 12 NVMe SSDs / 12 SAS SSDs) 2S compute storage node-HY: (2 CPU, 24 DDR4 DIMMs, 1 GPU, 2 NVMe SSDs, 12 SAS HDDs) Support Server : 2288H V5/1288H V5/2488 V5/2488H V5			
Network	10GE			
Storage	FusionCube distributed block storage			
Management software	FusionCube Center			
Virtualization software	FusionSphere 6.3 or later			
Dimensions	12U (442mm(W)*530mm(H)*840mm(D))	4U (175 mm x 447 mm x 898 mm)	2U (436mm*86.1mm*805mm)	2U (86 mm x 447 mm x 748 mm)
Power consumption	Up to 6,000 W per chassis	Up to 2,400 W per chassis	Up to 2,000 W per chassis	Up to 12,000 W per cabinet (up to 1,200 W per node)
Remarks	Customizable configuration, with memory, hard drive, switch, cabinet and other optional accessories			

Customer Benefits

» Ultimate Performance

- Distributed storage engines improve I/O performance by 3x to 5x
- Supports NVMe SSDs, further improving storage performance
- Integrated and optimized for VDI and virtualization scenarios
- Hardware-software synergy for higher application performance