EonStor CS Family Technical Overview

This white paper describes technical design and details of Infortrend scale-out NAS product - EonStor CS family with Infortrend distributed file system.

Version: 1.3

Updated: April 2020



Contents

Contents **2**
Infortrend EonStor CS System Overview **4**
EonStor CS Nodes 4
Network 5
I/O Load Balancing 6
Provisioning Structure 7
Disk Pools 7
RAID Logical Drive 7
Data Protection **8**
Erasure Code/ Replica 8
File Write/ Read 8
Protection Level 8
Healing 10
Infortrend RAID Technology 11
Online Expanding **12**
Scale-Out 12
Scale-Up 12
Expanding Capacity and Performance 13
Firmware Upgrade **14**
Immediate Firmware Upgrade 14
EonStor CS NAS Management **15**
Management Interface 15

Web-GUI 15
SSL Connection 15
Deployment Wizard 15
Authentication and Access Control 15
Data Service Software 16
Performance Monitoring 16
Quota Management 16
Remote Replication (Rsync) 16
Write Once Read Many (WORM) 16
Self-encrypting Drives (SED) 16
Tiered Data Migration 16
SMB Multichannel 16
Deployment and Expanding Example 17
Scenario 1 – Deployment 17
Scenario 2 – Expanding Nodes 17
Scenario 3 – Expanding Capacity Only 18
Conclusion 19
References 19
Legal Information 20
Trademarks 20
Contact Information 21
Website 21
Customer Support 21

Infortrend EonStor CS System Overview

Infortrend EonStor CS is a scale-out NAS system with high performance and capacity expansion capabilities. By adding more nodes to a system, you can easily get high performance and capacity in a single namespace file system.

EonStor CS offers multiple data protection designs to guarantee that enterprise customers' data is highly protected. It provides RAID for disk protection, Erasure Code and Replica for node protection, and remote replication for cluster failure protection. The Self-healing function can further maintain the same protection level when a node failure occurs.

EonStor CS features complete integration of software and hardware, with an easy-to-use web management interface and swappable modular hardware design, user is able to deploy the system in 30 minutes and easily maintain it.

EonStor CS Nodes

EonStor CS cluster consists of multiple nodes. A node is an appliance in a rack form which includes CPU, memory, networking, and drives that provide network file storage capability by clustering multiple nodes in a single cluster; more computing power, networking options, and storage capacity can be provided.

EonStor CS provides multiple hardware models, depending on performance and capacity needs, a user can choose hardware model suitable for his applications.

The table below shows each CS series' corresponding number of CPU core and available form factors.

Series	CPU Core	Model
CS 4000	12	CS 4016G CS 4024G CS 4025B CS 4060D/4060G
CS 3000	8	CS 3016G CS 3024G CS 3060D/3060G
CS 2000	4	CS 2060D/2060G

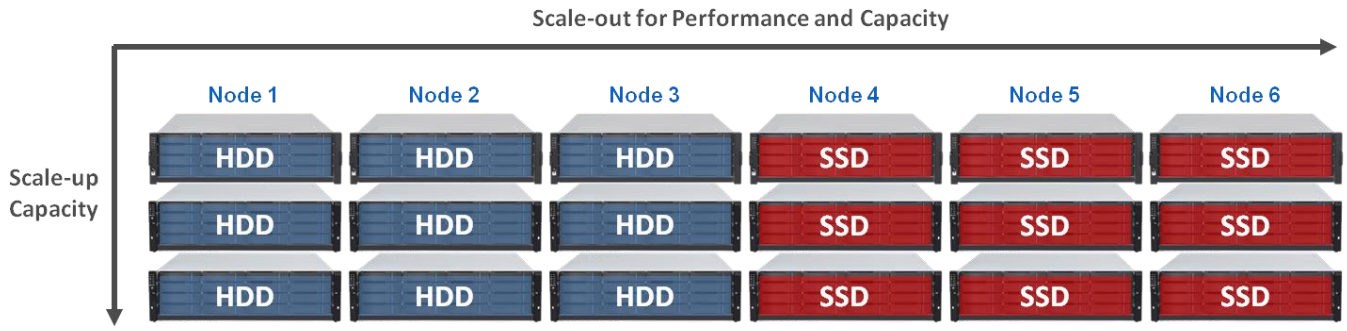
EonStor CS Family Model Name

Note:

- G indicates there is 1 node in the appliance, and D indicates there are 2 nodes in one appliance.
- CS 4025B supports SSD only.
- Only CS 4025B can be configured as an SSD node.

An EonStor CS cluster can scale out to 144 nodes for now, and minimum deployment node number is one node for Distributed mode, and three nodes for Erasure Code and Replica configurations. Currently, EonStor CS provides the HDD/SSD Node configurations. Drives in EonStor CS node should be fully populated to provide maximum performance and capacity, and both HDD and SSD are supported.

Each Node can also scale-up the node capacity through the SAS expansion channel (12 Gb/s), so that the user can expand system capacity with minimal initial configuration. The expansion enclosure will be automatically configured when connected to a head node, no further configuration is required. A node can support up to 76 drives for CS 3U 16-bay model/84 drives for CS 3U 24-bay model/100 drives for CS 2U 25-bay model/120 drives for CS 4U 60-bay model.

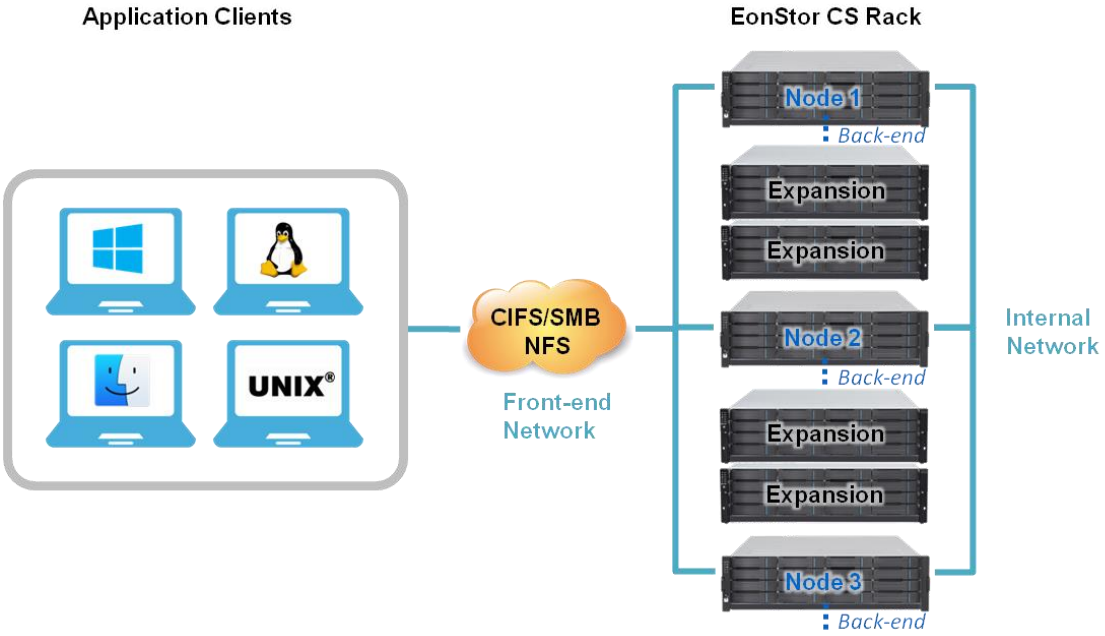


Scale-out and Scale-up of EonStor CS Cluster

Network

There are two networks for the cluster, which can be monitored through the management page.

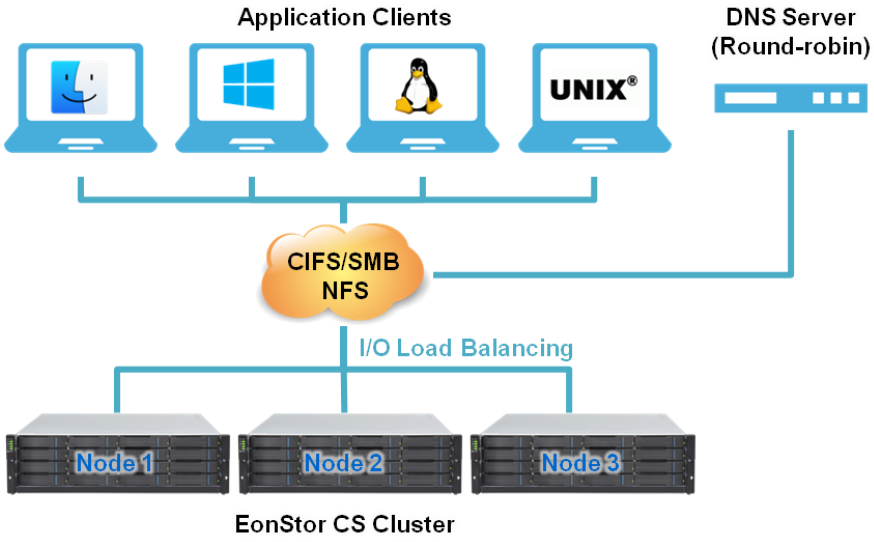
- Front-end network – This network is for user clients to access cluster nodes for data I/O. As EonStor CS supports single namespace, any client can connect to any node to access the same file. Clients can use 10GbE (SFP+)/25GbE (SFP28)/40GbE (QSFP+), and NFS/CIFS/FTP to connect to the nodes. Front-end network also supports trunk group.
- Internal network – This network is for nodes internal communication and should be isolated from a user client network. Through 10GbE (SFP+), 25GbE (SFP28), and 40GbE (QSFP+) interface, the nodes are all connected for metadata sync and data forwarding. The trunk function is enabled by default.



Network of EonStor CS Cluster

I/O Load Balancing

EonStor CS can scale out to multiple nodes to provide service for multiple clients. To make sure all nodes can provide service and maximize the overall system performance, data will be evenly distributed to all nodes. Moreover, to effectively balance I/O loads between the nodes, EonStor CS supports DNS Round–robin for clients to access to different nodes every time when they connect to CS node, this will make I/O balancing for all nodes and, therefore, the system can provide its best performance.



DNS Round-Robin for I/O Load Balancing

Provisioning Structure

Disk Pools

Infortrend EonStor CS supports disk pools for different application needs, and user is allowed to create disk pool based on different protection levels or disk types. For example, user can create Replica disk pool for critical data, and Distributed disk pool for application requiring high performance.

All disk pools can be managed in a single namespace, which eliminates management difficulty when using NAS system. When creating a shared folder, user can specify a disk pool to store data and change to another disk pool later on with just a single click, the folder path will not be changed, and data will migrate automatically from one pool to another.



Create shared folder on selected disk pools

RAID Logical Drive

EonStor CS uses Infortrend RAID Protection technology to provide system high availability and high performance. When first-time deployment, user can specify RAID level to protect the disk, the system will automatically create Logical Drives (LDs) for all the node's disks and provide as the base unit of the disk pool. The RAID level cannot be changed after deployment: when adding a new node or JBOD, the same RAID level will be applied.

Data Protection

Infortrend EonStor CS supports multiple data protection mechanisms to protect the system from disk, node, or even cluster failure. These protection mechanisms are described below.

Erasure Code/ Replica

EonStor CS provides Erasure code and Replica to distribute data across multiple nodes, so when there is a node or LD failure, data can still be protected and accessed. Erasure code is more space-efficient, it is recommended to use Erasure Code when there are more nodes in the cluster. Replica provides higher read performance compared to Erasure code; you can also use it when there are only a few nodes in the cluster.

File Write/Read

When data is written to the node in the cluster, it will first be distributed to other nodes depending on disk pool protection level, and only when data is completely written to all nodes, the system will respond to the client that write has been completed. For better performance, a memory cache is used to store the data, and it is protected by the battery.

When the client reads data from the node in the cluster, data may be retrieved by the system from other nodes via the internal network, and then it assembles the file if necessary and returns to the client.

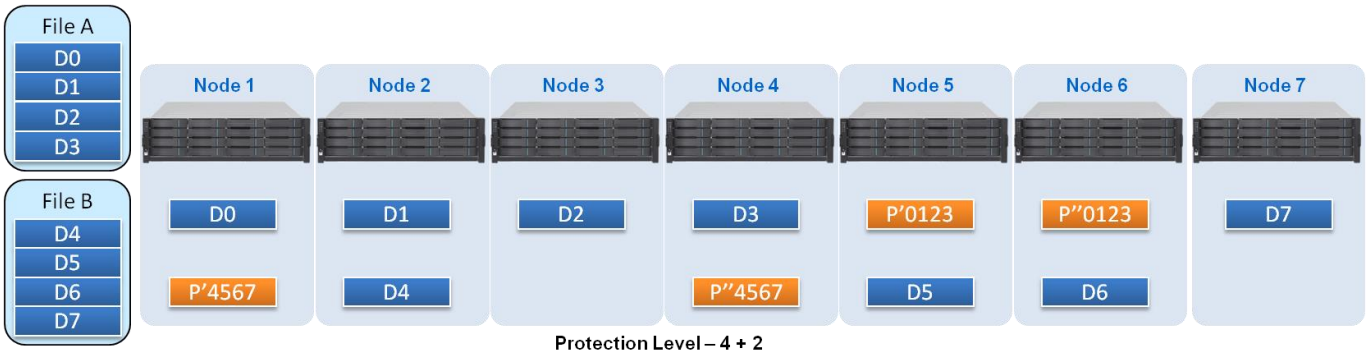
Protection Level

Infortrend EonStor CS supports the following data protection levels: Distributed, Erasure Code, and Replica. In ‘Distributed’ configuration, only one copy of a file is stored in the cluster, and, therefore, higher performance is provided. When a client tries to store a file, the system uses a hash algorithm to decide in which node to store the file, and then store the whole file within a chosen node. There is no failure tolerance for nodes, but the files are still protected by RAID.

On ‘Erasure Code’ level, user is allowed to specify N+M level protection, while N is the number of nodes to stay functioning, and M represents the number of nodes that can be lost for the cluster to continue to operate.

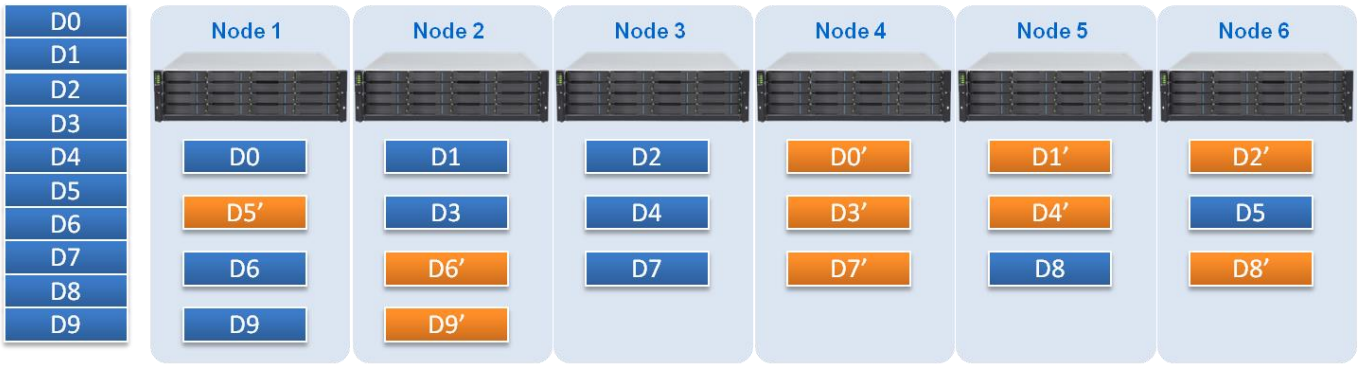
Protection Type	Protection Level (N+M)	Description
Distributed	1	Configured with 1 node. Each data will be stored in a single node.
Erasure coding	2+1	Needs at least 3 nodes to configure. Allows losing 1 node.
	4+1	Needs at least 5 nodes to configure. Allows losing 1 node.
	4+2	Needs at least 6 nodes to configure. Allows losing 2 nodes.
	8+1	Needs at least 9 nodes to configure. Allows losing 1 node.
	8+2	Needs at least 10 nodes to configure. Allows losing 2 nodes.
Replica	X2	Needs at least 3 nodes to configure, replicates 2 copies. Allows losing 1 node.
	X3	Needs at least 3 nodes to configure, replicates 3 copies. Allows losing 2 nodes.

When the protection level is Erasure Code N+M, data will be sliced into N pieces and M data parity will be generated. The N pieces of data and M parity will then be randomly written to N+M nodes based on node capacity. When there is a node loss, data can be calculated through the parity. The following is an example of the stripe unit storage of Erasure coding 4+2.



Erasure Coding N+M Protection

On the other hand, when the protection level is 'Replica', data will be replicated to other nodes in the cluster to make multiple copies. The following is a sample stripe unit storage example of Replica X2.



Replica Protection

Healing

Infortrend EonStor CS supports self-healing function. When there is a node or LD failure, a system can start a healing process to maintain system protection level. Below are details of the healing process.

- Node failure healing – When there is a node or expansion unit failure, the system will fall into degraded status. If a node or expansion unit cannot come back to a normal status in some time, self-healing process will be triggered to write data from the failed node or expansion unit to other active nodes. Once the self-healing process is completed, the system will return to the normal status and data protection level will be the same as before.
- LD failure healing – Where there is an LD failure, the system will fall into degraded status, the self-healing process will be started to write data from the failed LD to other normal LDs after some time. Normally, the LD in the same node will first be used to write data, and only when there is no capacity left, the data will be written to LDs in other nodes.

Infortrend RAID Technology

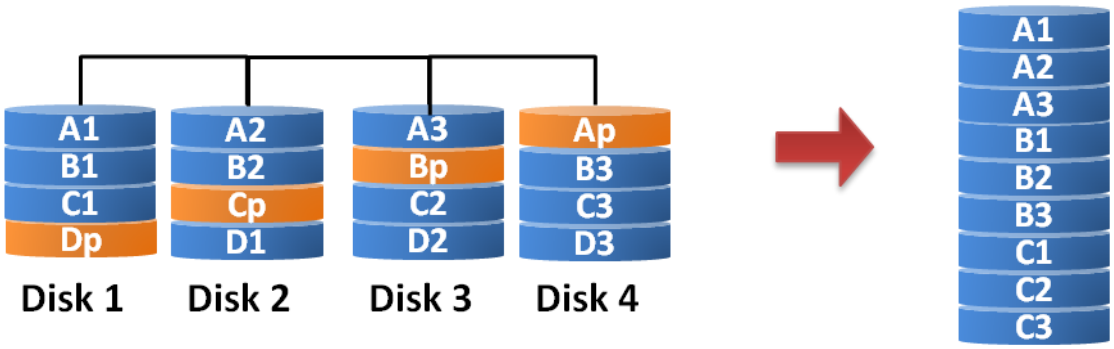
EonStor CS uses Infortrend RAID technology to protect disk in the node. RAID (Redundant Array of Independent Disks) is a storage virtualization technology that integrates multiple hard disks to provide better performance and data protection. Different RAID levels can be selected based on the user's capacity and protection level requirements. There are three configurations for the EonStor CS RAID level.

(1) Optimal Performance – This configuration provides the best system performance; it uses RAID 5 for disk protection. Several LDs will be created to maintain certain level disk protection, and the exact LD number is dependent on the enclosure drive number.

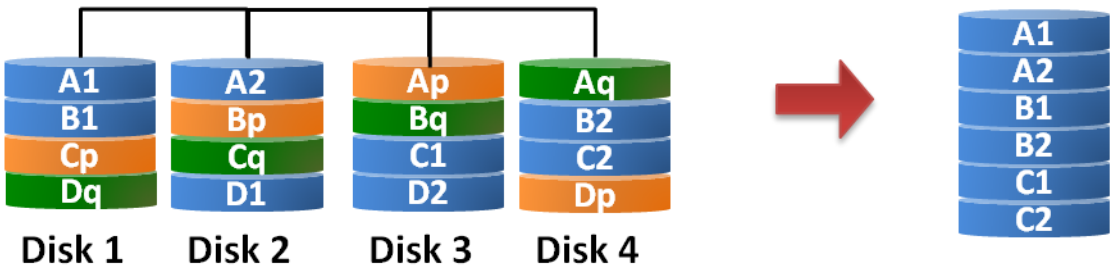
(2) Balance – The same as Optimal Performance mode, Balance mode also uses RAID 5 for disk protection, the only difference is that in Balance mode, there will be a spare drive in each enclosure, this can provide higher level protection, but on the other hand, usable capacity will be smaller.

(3) Optimal Protection – This configuration uses RAID 6 for disk protection and, therefore, allows 2 drives failure for each LD. Compared to Optimal Performance mode, there will be slight performance downgrade, but it can offer a better drive failure protection.

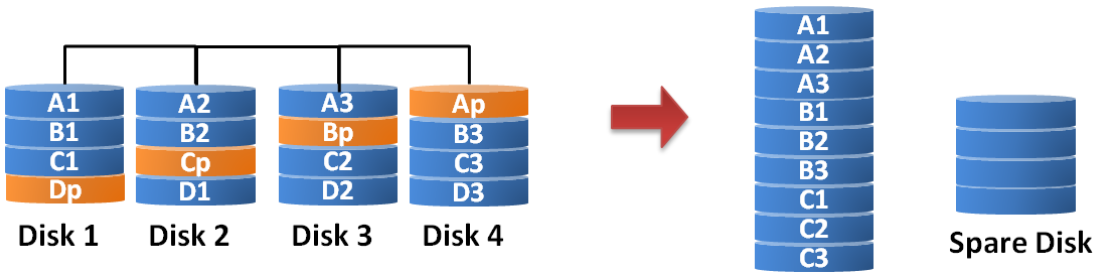
Unlike the architecture using Erasure coding/Replica across nodes, RAID technology can achieve a single node's scale-up stable extension and reduce the failure domain problem.



Optimal Performance



Optimal Protection



Balance

Online Expanding

EonStor CS supports both scale-out and scale-up online expansion. During the expansion process, the system does not need to be shut down and, therefore, maintains high availability.

Scale-Out

By adding nodes to the cluster, you can linearly increase system performance and capacity. Each node can provide extra CPU, network bandwidth, and storage capacity. Only with a few clicks, a node can be added to the cluster, and user is allowed to re-balance data in the cluster to make sure all nodes anticipate providing I/O.

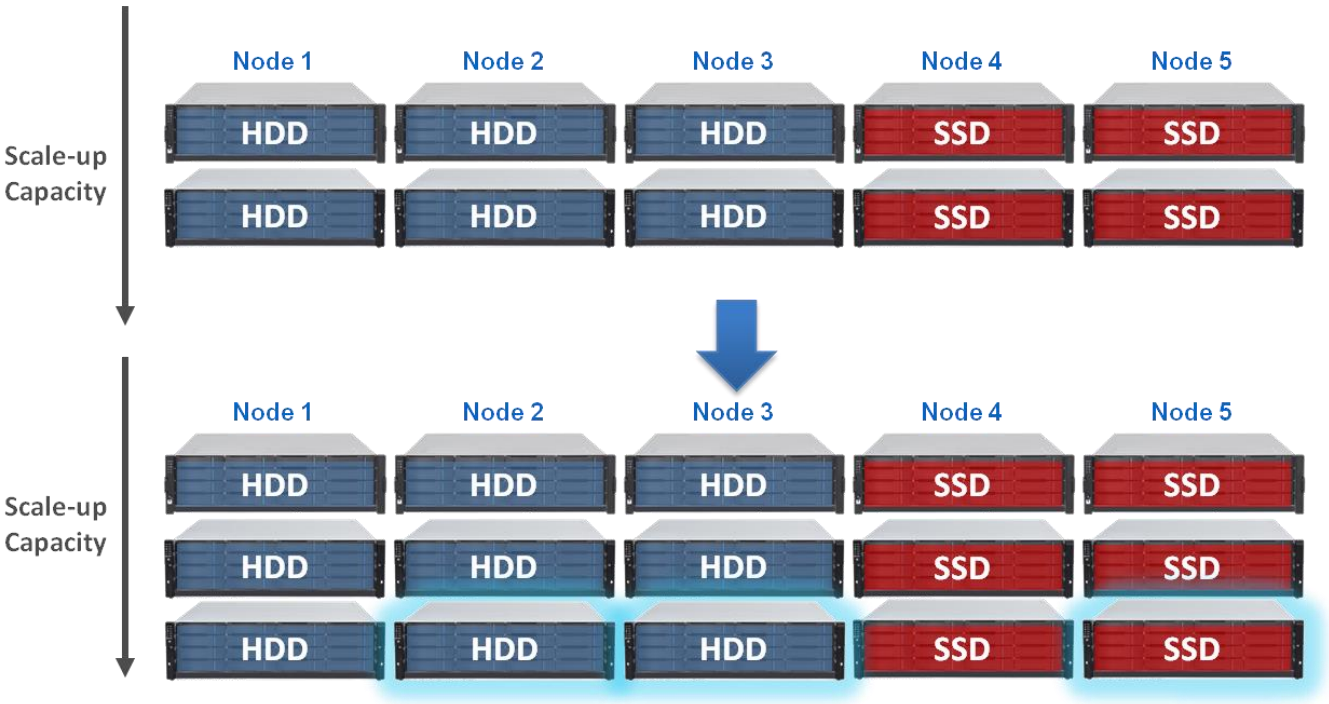
Scale-Up

If more capacity is needed, EonStor CS provides a scale-up extension by adding an expansion enclosure. It should be noted that nodes belonging to the same disk pool must extend the same expansion enclosure number to achieve the best data balance. EonStor CS family has some scale-up rules that users need to follow.

Please check the below table for the scale-up configuration rule.

CS model	Maximum Expansion Enclosure (per node)	Total Number of Drives (per node)
CS 3U 16-bay model	JB 3016 * 3	64
	JB 3060 * 1	76
CS 4U 24-bay model	JB 3016 * 3	72
	JB 3060 * 1	84
CS 2U 25-bay model	JB 3025 * 3	100

CS 4U 60-bay model	JB 3016 * 3	78 (D model) 108 (G model)
	JB 3060 * 1	90 (D model) 120 (G model)



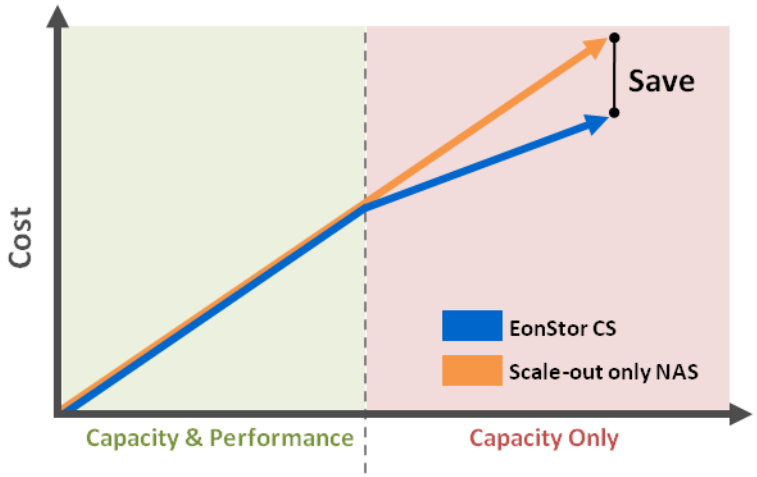
Scale-up Rule



4U 60-bay High-density Expansion Enclosure

Expanding Capacity and Performance

Unlike other scale-out NAS products that can only add a node to increase both performance and capacity, EonStor CS has the flexibility to add expansion enclosure to expand capacity only. This is useful when user needs more capacity for the cluster, and it is more cost-effective compared to other scale-out NAS products.



Cost saving by scale-up

Firmware Upgrade

EonStor CS provides immediate firmware upgrade for new features and bug fixes.

Immediate Firmware Upgrade

Immediate firmware upgrade will stop the service and upgrade all nodes at the same time, there will be a system downtime and, therefore, it is recommended to do this during off-hours.

EonStor CS NAS Management

Management Interface

Web-GUI

EonStor CS provides a central web management interface for system management, administrator can connect to EonStor CS management IP to use EonOne to do the management, and no additional management server is needed. The EonOne management interface can manage the whole cluster system and all nodes, and when there is no failure, EonOne can still operate without impact.

SSL Connection

When log in to the management interface, data is encrypted by SSL to ensure connection security.

Deployment Wizard

A Deployment Wizard is provided to guide user step-by-step to build the cluster in just a few minutes.

Authentication and Access Control

EonStor CS provides the following user management for client to connect to the cluster for accessing the file system.

- Active Directory (AD)
- LDAP (Lightweight Directory Access Protocol)
- NIS (Network Information Service)
- Local Users & Groups

Data Service Software

EonStor CS provides comprehensive data service software features.

Performance Monitoring

The EonOne management interface can display real-time system performance, and administrators can monitor performance based on interfaces, such as cluster, node, and network interface.

Quota Management

The ability to set a capacity quota for the shared folder allows the administrator to plan the usable capacity of the shared folder for different purposes.

Remote Replication (Rsync)

Offsite replication supports the Rsync protocol, which can copy data between devices that support Rsync, back up data, and backup can be periodically replicated through scheduling.

Write Once Read Many (WORM)

EonStor CS supports WORM: once data is written, it becomes ‘read only, and cannot be modified, deleted, or overwritten. WORM can be enabled by disk pool, and once enabled, it cannot be disabled again. WORM can prevent accidental modification or deletion of critical data to meet regulatory or management requirements.

Self-encrypting Drives (SED)

EonStor CS supports SED that automatically and transparently encrypts data written to the disk without degrading performance. The SED uses AES 256 to encrypt data. The encryption key is generated randomly and, therefore, is secure enough to protect user data.

Tiered Data Migration

With tiered data migration, users can migrate the data between disk pools. Users can keep the frequently accessed data in SSD disk pool. When users no longer need to access these data very often, users can migrate these data to HDD disk pool for archiving. With this feature, EonStor CS cluster can leverage the benefits of different drive types.

SMB Multichannel

SMB Multichannel enables multiple network connections between cluster and clients simultaneously. With SMB multichannel enabled, multiple paths are available between the SMB 3.0 clients and the SMB 3.0 cluster, so it facilitates aggregation of network bandwidth and network fault tolerance.

Deployment and Expanding Example

This section will explain how to deploy and expand EonStor CS cluster through several scenarios.

Scenario 1 – Deployment

The company purchases scale-out NAS and create a shared folder for marketing department with folder name "Marketing". There will be 3 EonStor CS nodes, and all equipped with NL-SAS hard drives.

(1) Follow Deployment Guide to connect network cables to EonStor CS external, internal, and management interfaces, - and power on the nodes.

(2) Use browser to connect to Deployment Wizard to add all three nodes to the cluster and complete the deployment process.

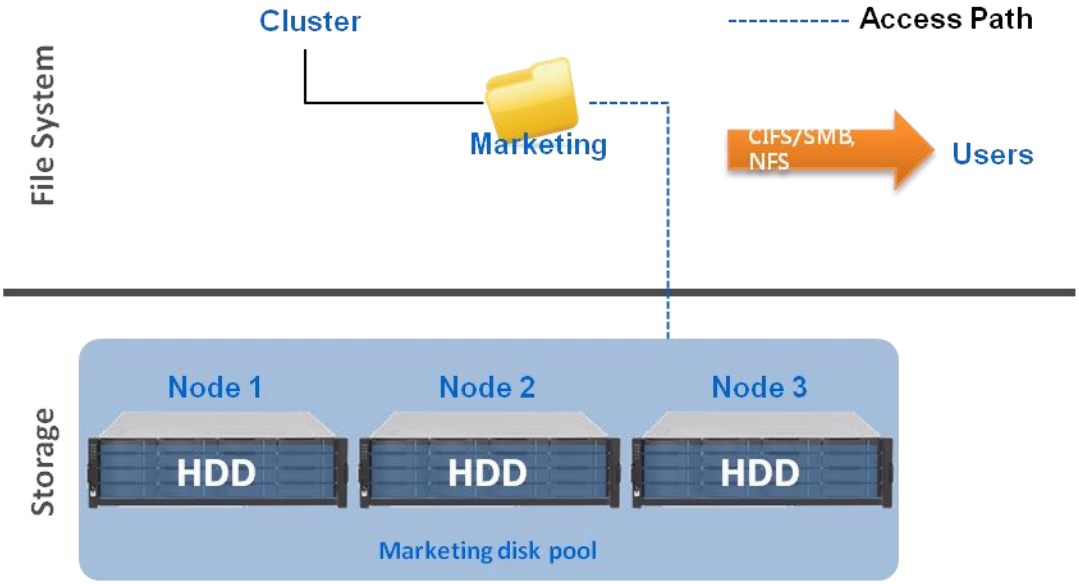
(3) Connect to EonOne management interface to create a disk pool “Marketing disk pool” and set the protection level to Erasure Code 2+1.

(4) Create a shared folder "Marketing" in the “Marketing disk pool”.

(5) Enable the network sharing protocol CIFS/SMB and NFS.

(6) Connect to Active Directory (AD) service and import all users.

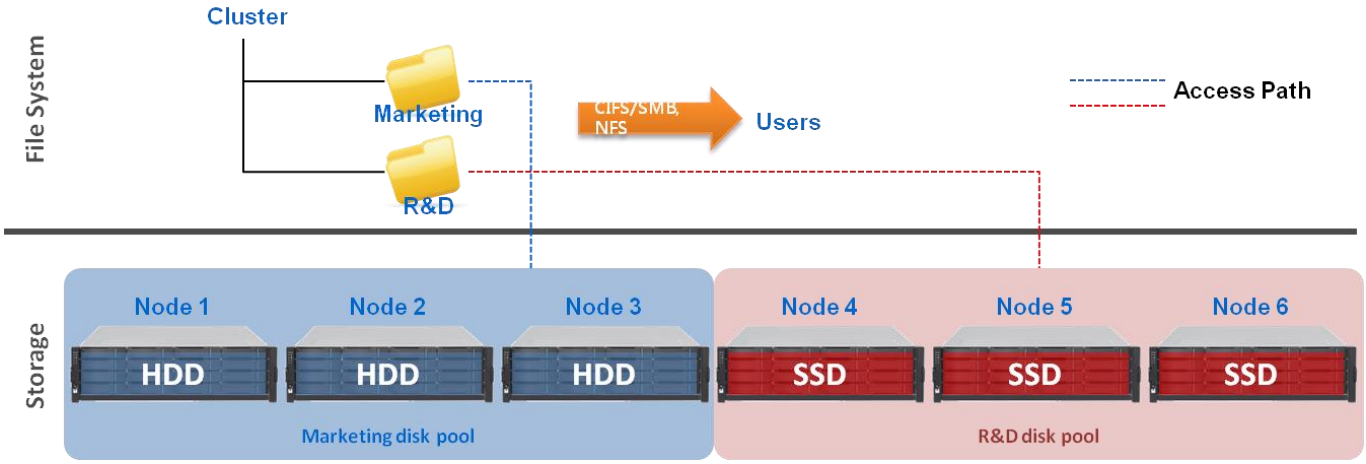
(7) Select user from marketing department and assign permissions for “Marketing” shared folder.



Scenario 2 – Expanding Nodes

The company expects to extend a scale-out NAS service to the R&D department while the 3 nodes cluster has run out of capacity and performance, and, therefore, 3 additional EonStor CS nodes will be purchased. A shared folder “R&D” for the R&D department will be created.

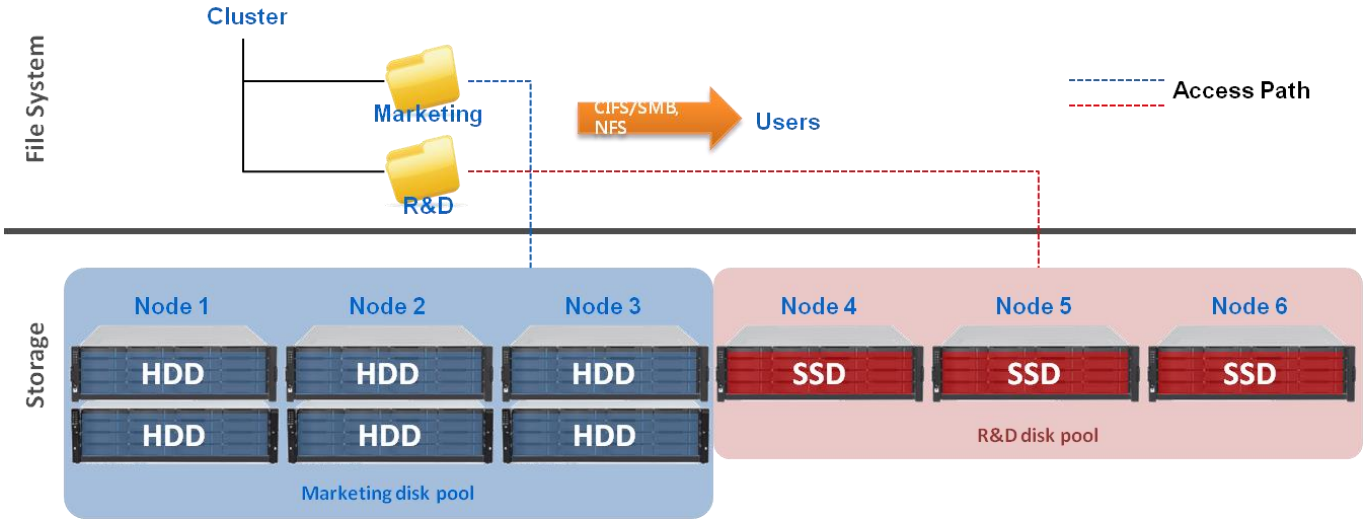
(1) Connect 3 new nodes to the existing cluster networks and power on.
(2) Login to the EonOne management interface and click ‘Add Node’ to add 3 nodes to the cluster.
(3) Create “R&D disk pool” capacity to use 3 newly added nodes and specify protection level to Replica x 3.
(4) Create a shared folder named "R&D" in “R&D disk pool”.
(5) Select users from the R&D group and assign permission for “R&D” shared folder.



Scenario 3 – Expanding Capacity Only

Demands for the Marketing department has increased, the company expects to add more storage capacity for the marketing department, therefore, 3 expansion enclosure devices, all with NL-SAS hard drives, are purchased according to the capacity requirements.

(1) The administrator connects the expansion enclosures to the back-end network with physical cables to Node 1, Node 2, and Node 3.
(2) Login to the EonOne management interface to expand "Marketing disk pool".



Conclusion

EonStor CS delivers high performance and capacity cluster file system, through EonStor CS scale-out architecture, client can access all nodes data in a single namespace; up to 144 nodes and more than 100PB capacity are supported in a single cluster.

With the support of RAID, Erasure Code/Replica, and Remote Replication, EonStor CS can provide comprehensive data protection. Through EonOne management interface, user can manage the whole cluster in a centralized web interface, with a single point of administration.

EonStor CS is ideally suited for unstructured data applications in the enterprise – including High Performance Computing (HPC), Media & Entertainment (M&E), large-scale file storage, backup and archive.

References

Category	Technical document
Document	EonStor CS Datasheet
Document	EonStor CS Best Practice
Document	EonStor CS Expansion Guide

Legal Information

All Infortrend products, including the product/s that customers have purchased from Infortrend, are subject to the latest Standard Warranty Policy available on the Infortrend website:
https://www.infortrend.com/global/Support/terms-conditions

Infortrend may from time to time modify, update or upgrade the software, firmware or any accompanying user documentation without any prior notice. Infortrend will provide access to these new software, firmware, or documentation releases from certain download sections of our website or through our service partners. Customer will be responsible for maintaining updated version of the software, firmware, or other documentation by downloading or obtaining from Infortrend, and installing designated updated code, including but not limited to firmware, microcode, basic input/out system code, utility programs, device drivers, and diagnostics delivered with Infortrend product.
Before installing any software, applications or components provided by a third party, customer should ensure that they are compatible and interoperable with Infortrend product by checking in advance with Infortrend. Customer is solely responsible for ensuring the compatibility and interoperability of the third party's products with Infortrend product.

Customer is further solely responsible for ensuring its systems, software, and data are adequately backed up as a precaution against possible failures, alteration, or loss. For any questions of hardware/ software compatibility, and the update/upgrade code, customer should contact Infortrend sales representative or technical support for assistance.

To the extent permitted by applicable laws, Infortrend will NOT be responsible for any interoperability or compatibility issues that may arise when (1) products, software, or options not certified and supported by Infortrend are used; (2) configurations not certified and supported by Infortrend are used; (3) parts intended for one system are installed in another system of different make or model.

Trademarks

Infortrend, the Infortrend logo, SANWatch, EonOne, EonStor, and EonServ are registered trademarks of Infortrend Technology, Inc. Other names prefixed with "IFT", "DS", "GS", "GSa", "GSc", "GSi", "GSe", "GSe Pro", "CS" are trademarks of Infortrend Technology, Inc. All other names, brands, products or services are trademarks or registered trademarks of their respective owners.

Contact Information

Website

For more information of Infortrend's products and services, visit:

https://www.infortrend.com/global/Home

Customer Support

Contact your system vendor or visit the following support site.

http://www.infortrend.com/global/Support/Support