

完全シンメトリック型スケールアウトユニファイドストレージソリューション

Compuverde

コアマイクロシステムズ株式会社
Core Micro Systems, Inc.

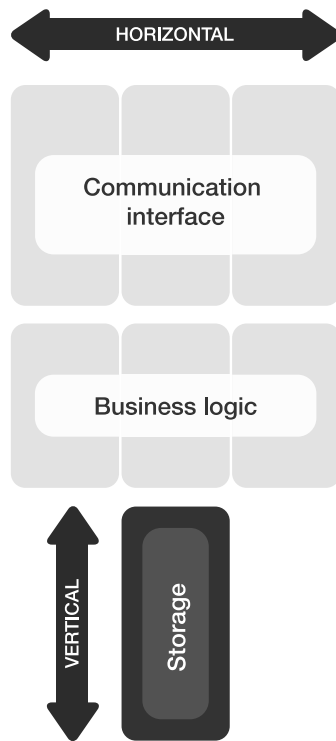
ストレージは企業構造の変化や計画に柔軟に対応/進化できることが必要

いままで



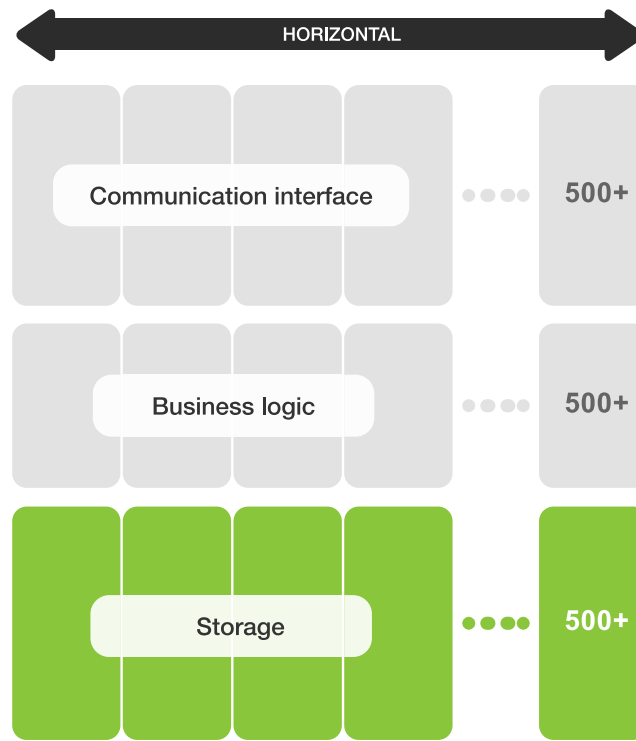
- 変更に制限がある
- 垂直的構造

現在



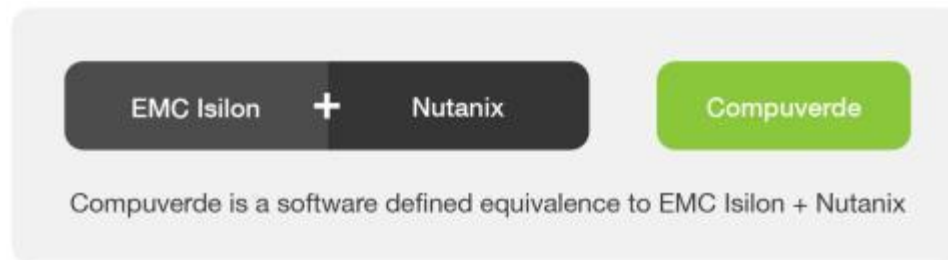
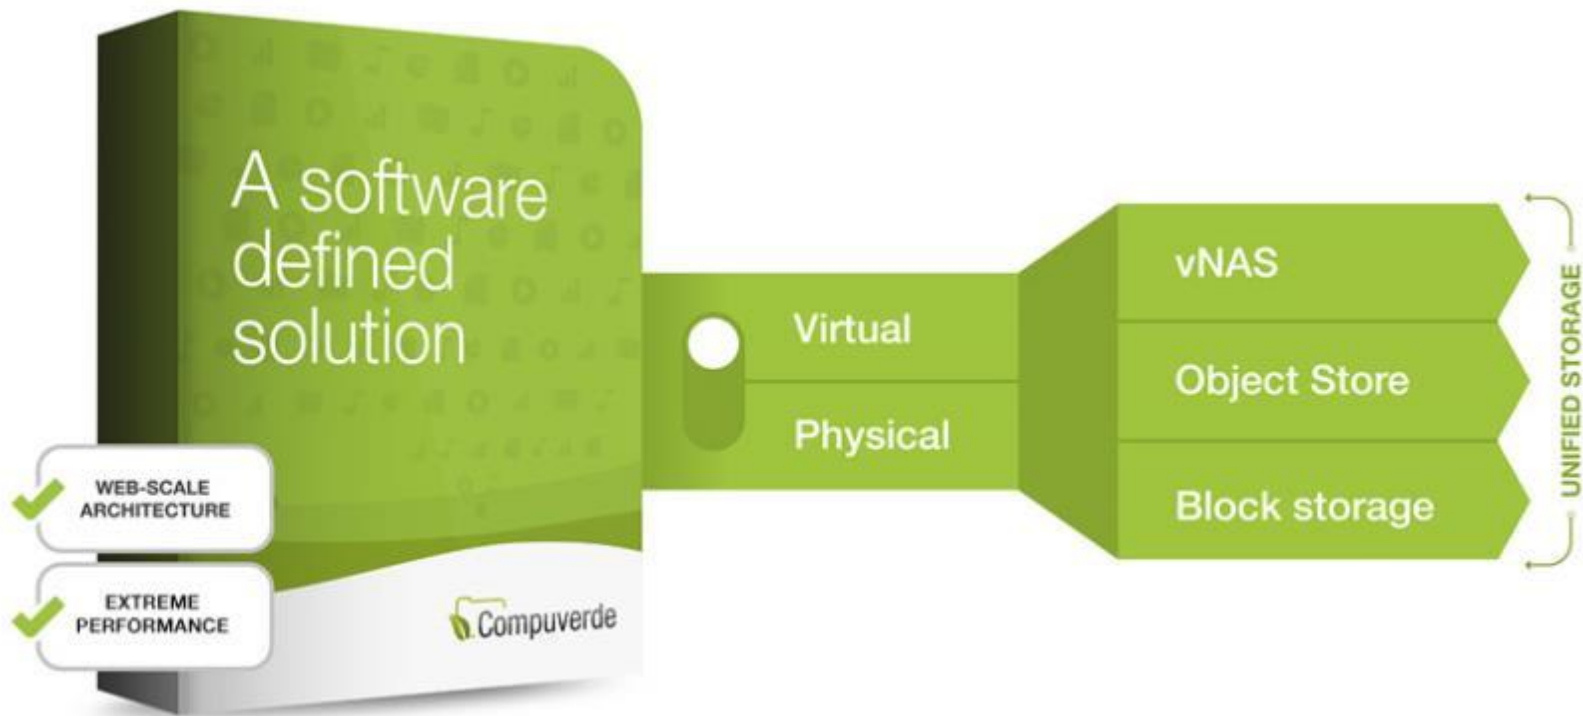
- 複数サービス & インターフェイス
- Defined Data Storage

これから



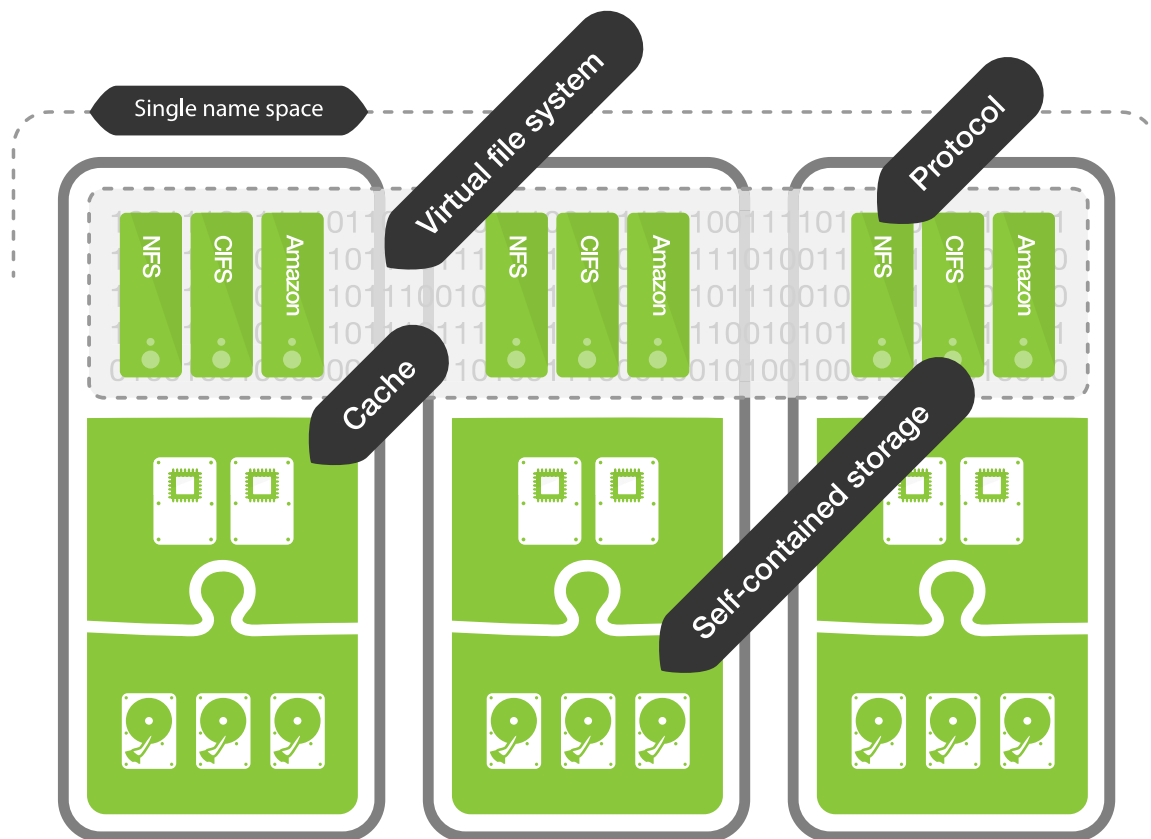
- 全ていつでも変更可能
- 拡張可能でハードウェアにとらわれない

ソフトウェア ラインナップ



ソフトウェア定義型(SDS) スケールアウトNAS/バーチャルNAS

多数プロトコル対応



☒ NFS
3/4/4.1/pNFS

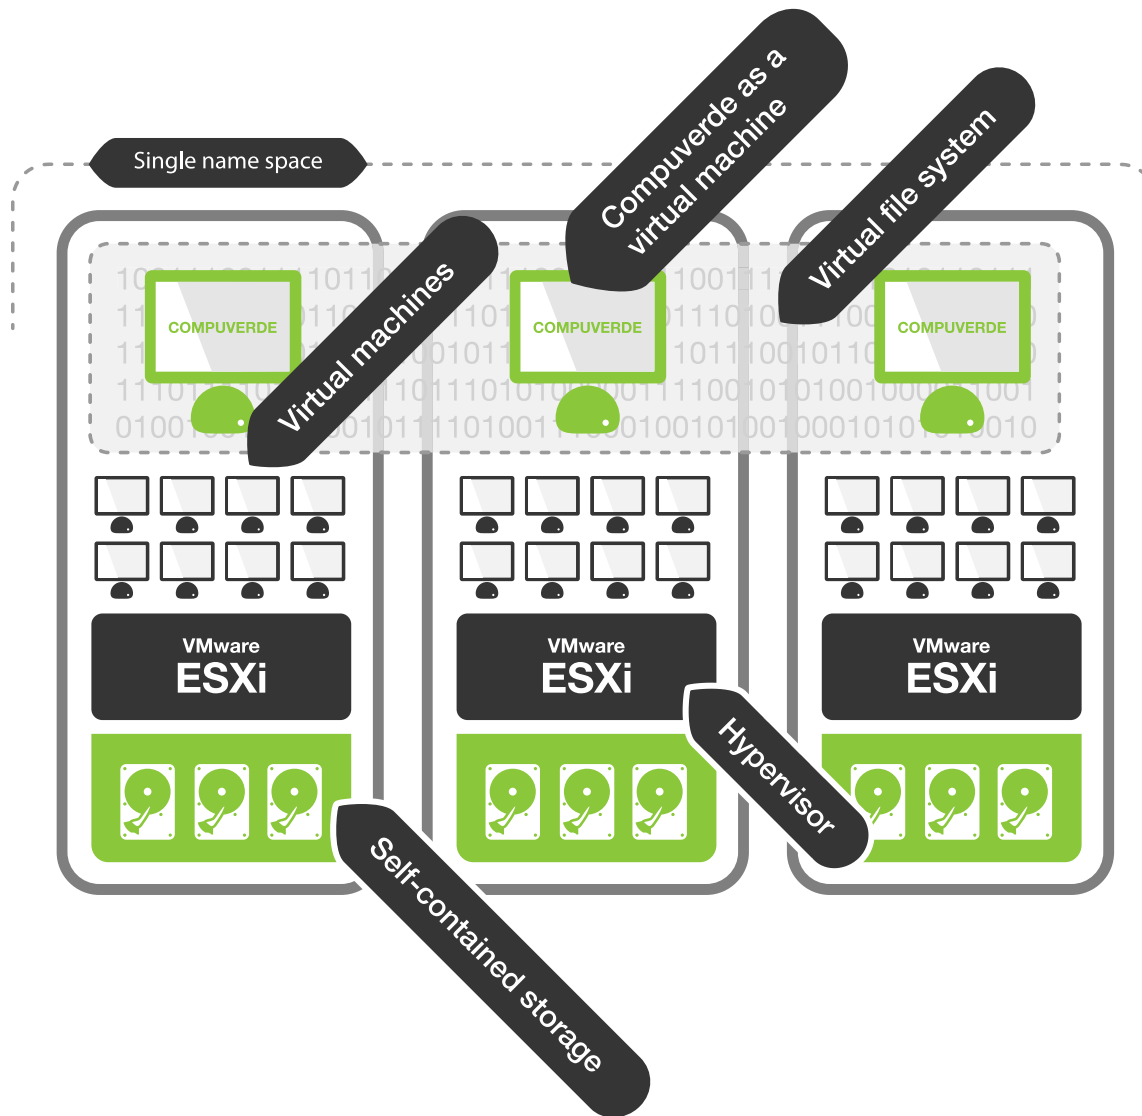
☒ CIFS
SMB 1/2/3

☒ OpenStack
Swift

☒ Amazon
S3

☒ CDMI

ソフトウェア定義型(SDS) スケールアウト仮想化プラットフォーム



各種ハイパーバイザ対応

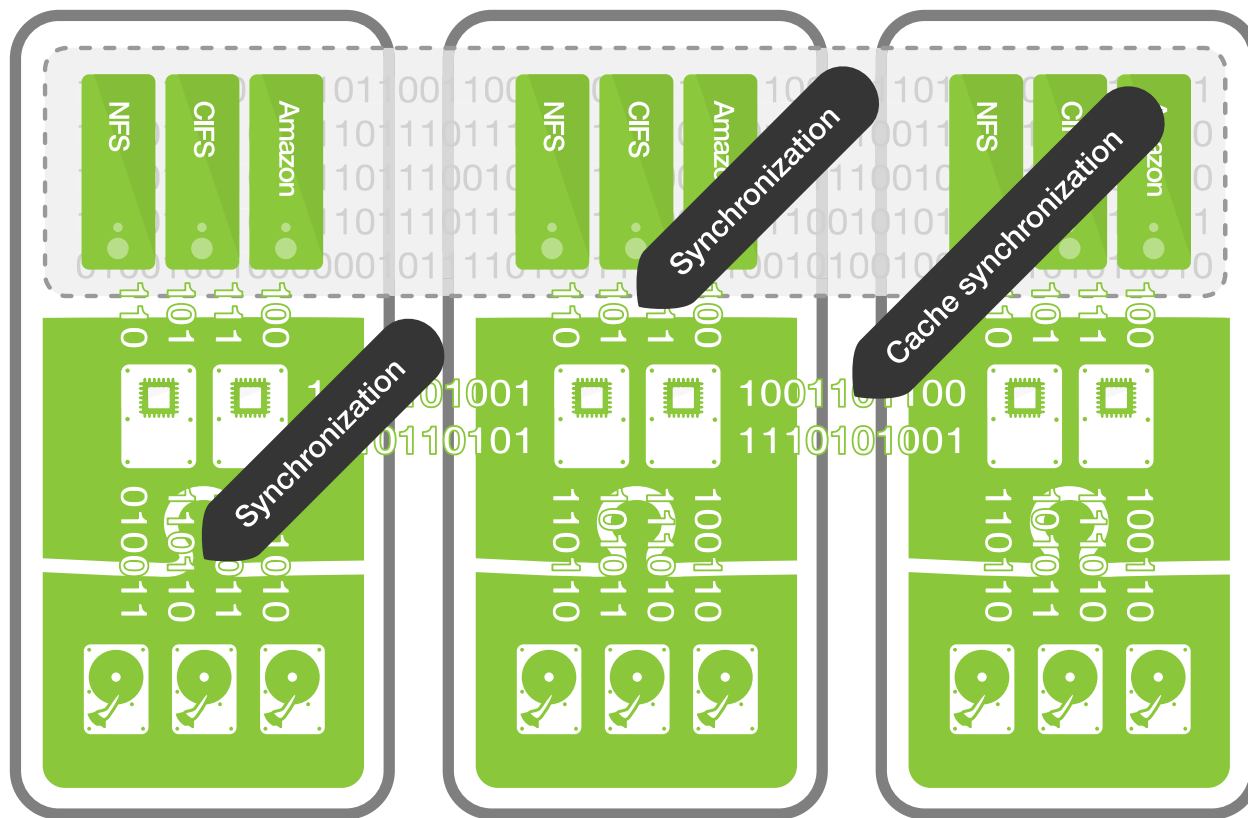
☒ VMware

☒ KVM

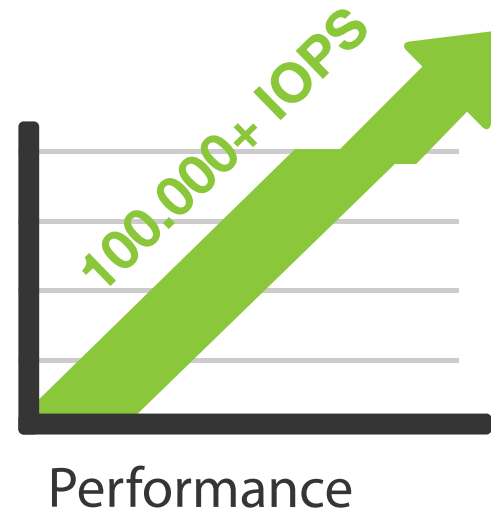
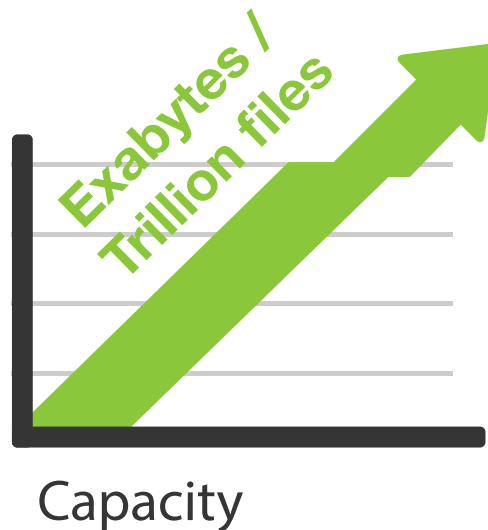
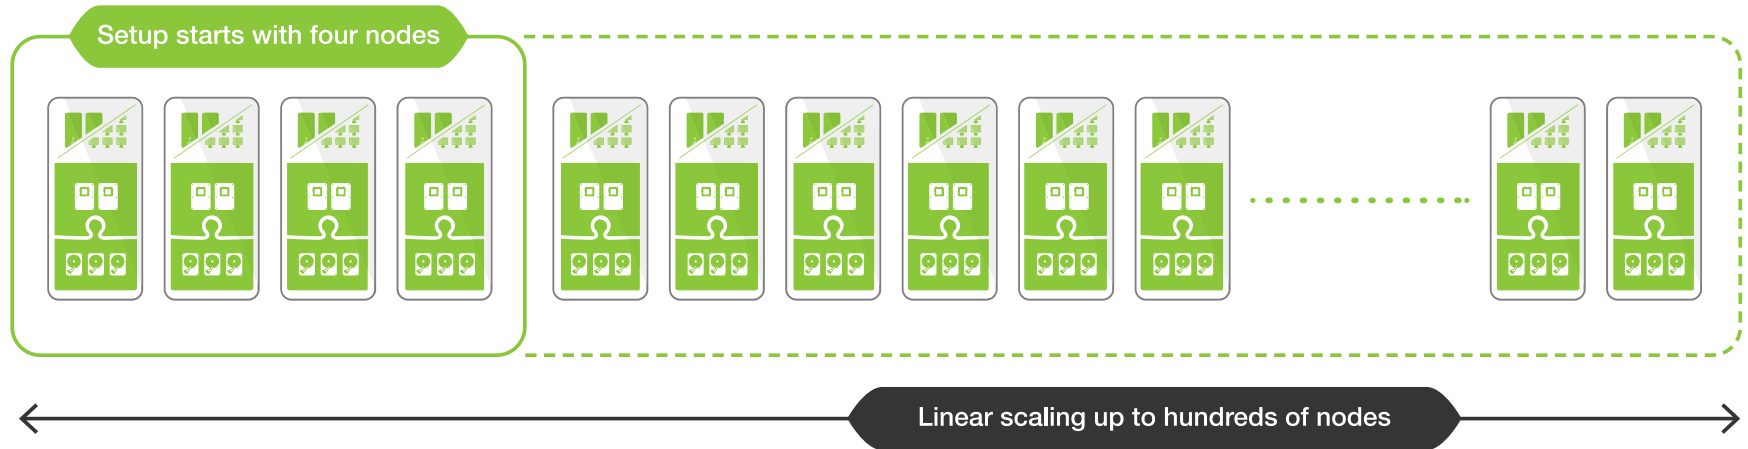
☒ Microsoft Hyper-V

☒ Xen

キャッシュ x メカニズム

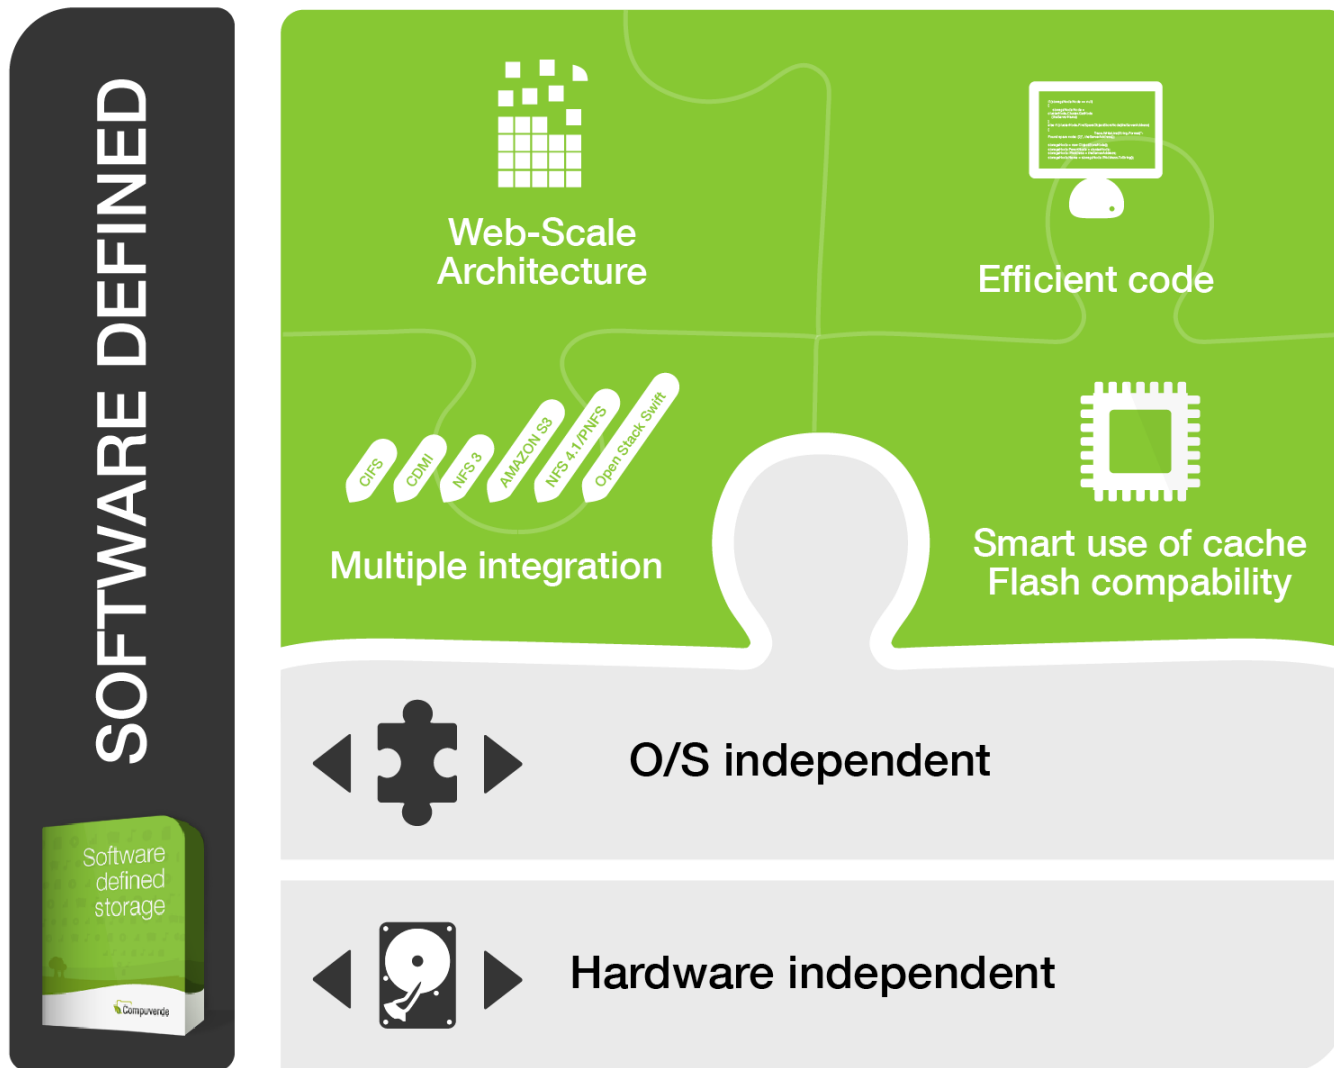


水平的にパフォーマンス&キャパシティを拡張

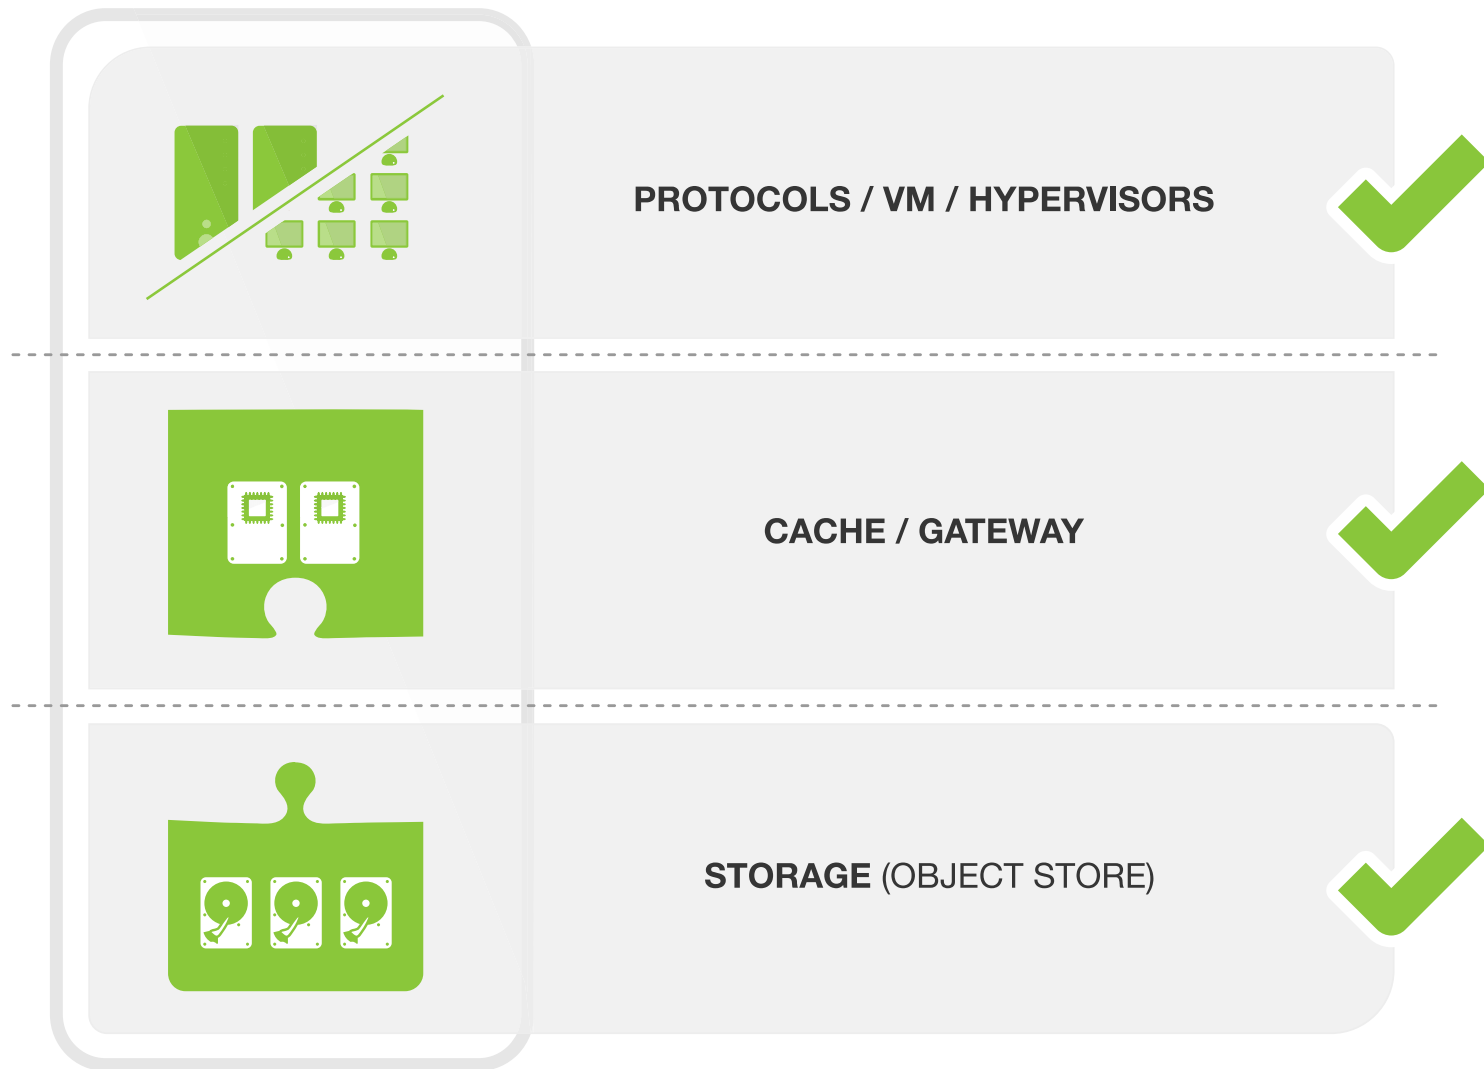




4つの機能モジュール



完全なソフトウェアスタックを提供

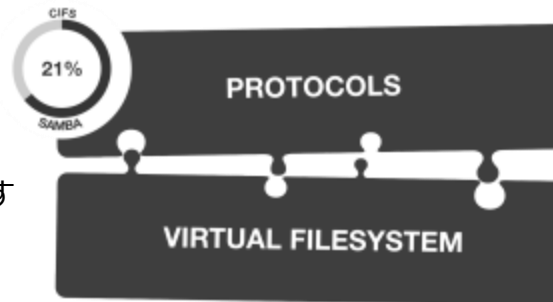


何が優れているのか？

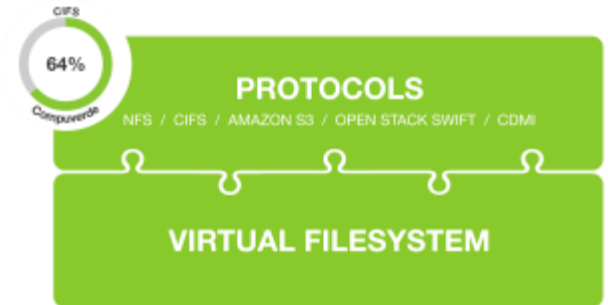
厳密な一貫性

独自の仮想ファイルシステム
と独自プロトコルの実装により
高い統一貫性が保証されています

COMPETITORS

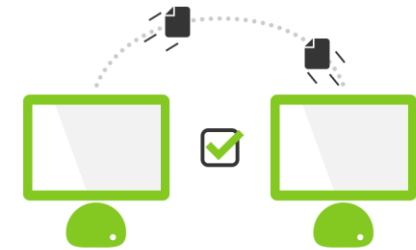
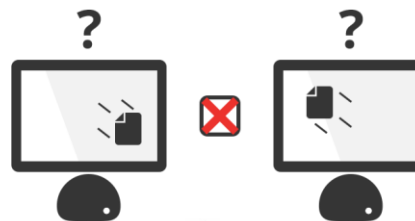


COMPUVERDE



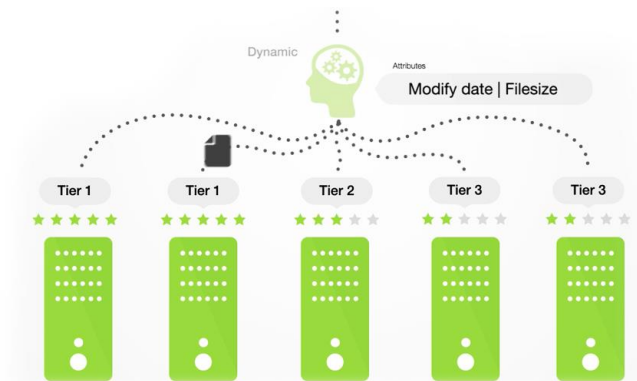
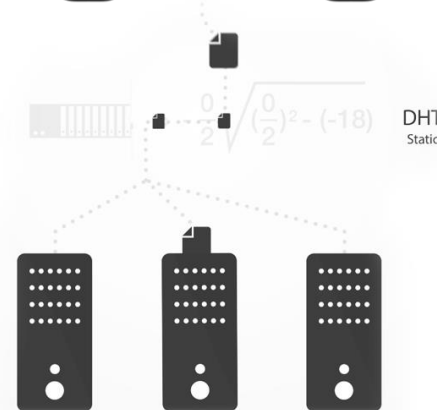
VMs間のファイル共有

vNAS機能を備えた
ハイパーコンバージドソリューション



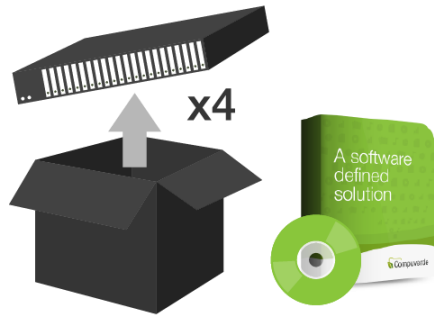
柔軟性

DHTに比べダイナミックで
インテリジェント

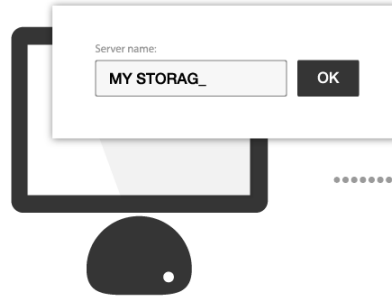


使用も管理も簡単！

(ソフトウェアレベルからのフルインテグレーションの場合)



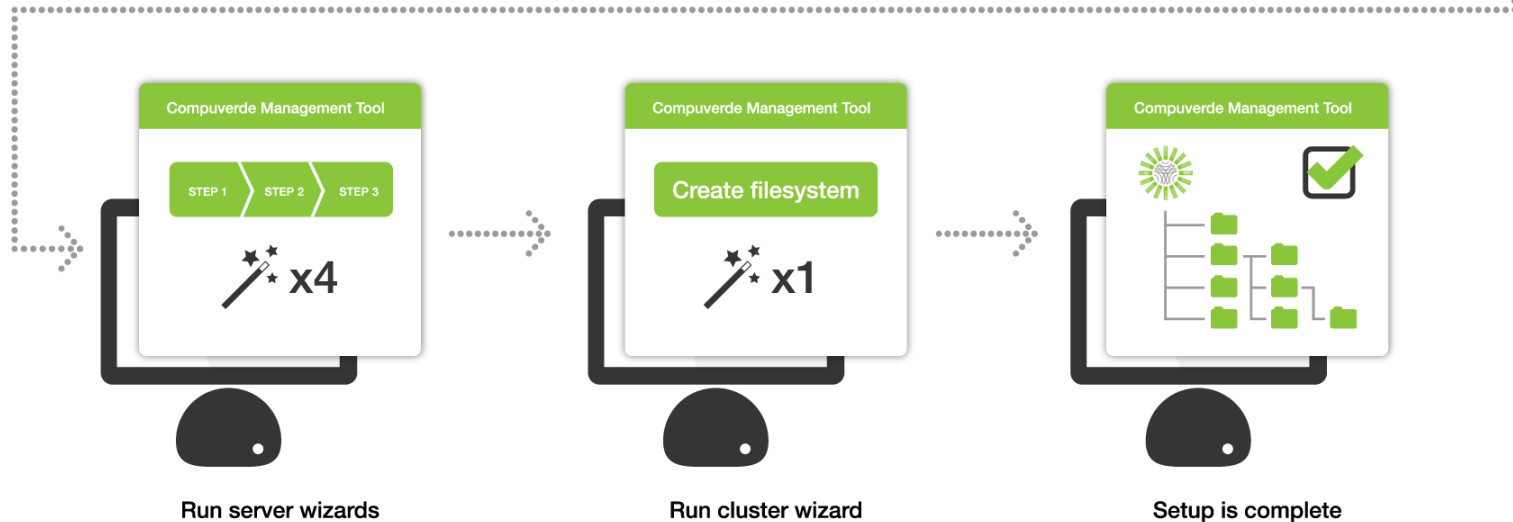
Unconfigured servers + Compuverde software



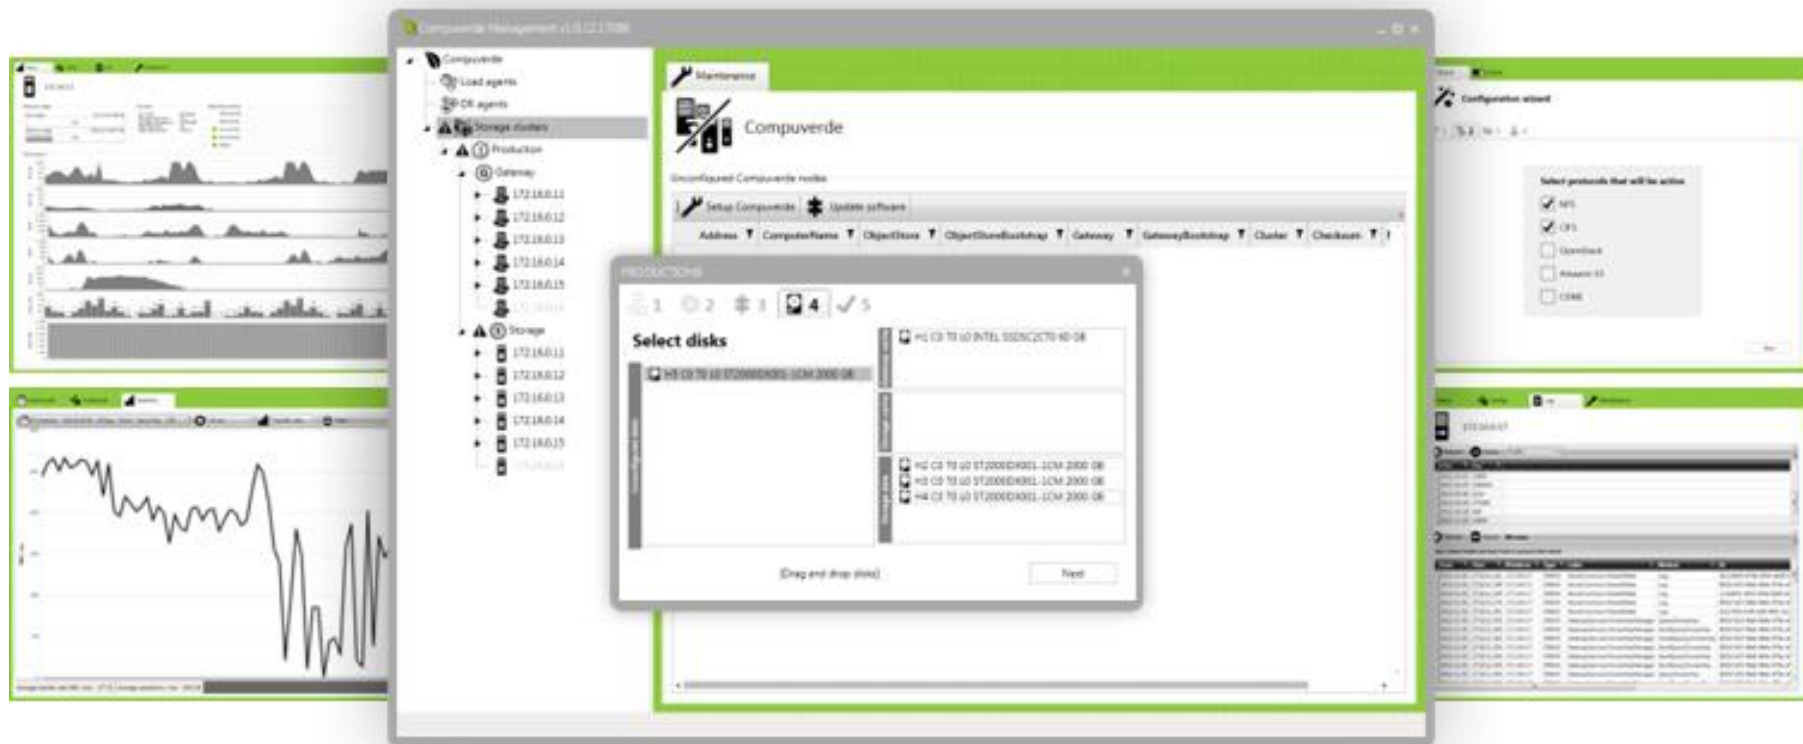
Name the servers



Connect to a network



わかりやすい管理GUI

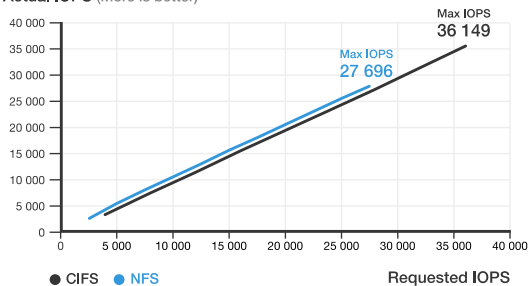


テストデータ (比較)

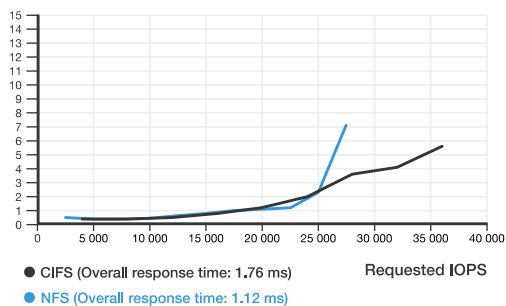
Compuverde

5.3x
faster

Actual IOPS (More is better)



Milliseconds (Less is better)



1
Node

Redundancy
(Can lose)

125
%

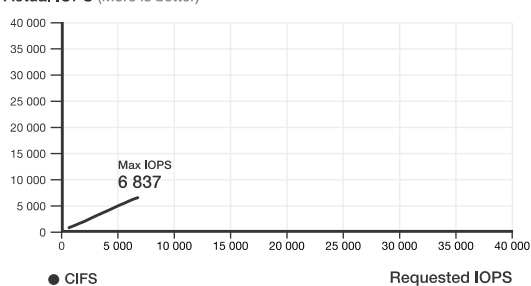
Data footprint
(Less is better)

23.0
TB

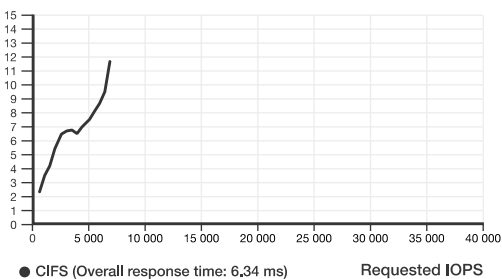
Usable space
(More is better)

Ceph (Hammer)

Actual IOPS (More is better)



Milliseconds (Less is better)



1
Node

Redundancy
(Can lose)

200
%

Data footprint
(Less is better)

14.4
TB

Usable space
(More is better)

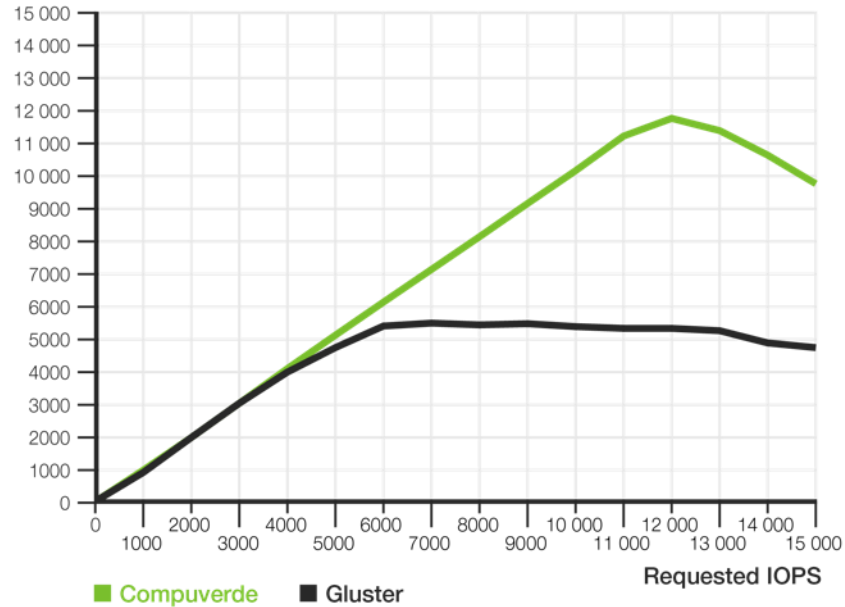
Hardware specifications (both systems)

CPU	2 x Intel Xeon E5-2620 (2.0 GHz)	x8
Memory	32 GB	
Cache	240 GB Flash	
Internal data disks	12 x 300 GB SAS 10K	
Raid controller	2 x LSI 9271 (1GB NVRAM)	
Network	10 GbE (1500 MTU)	

SPECsfs TESTS – COMPUVERDE VS GLUSTER

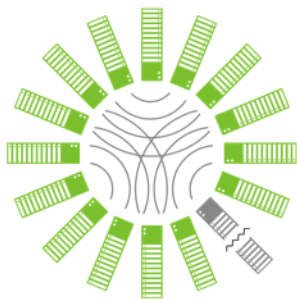
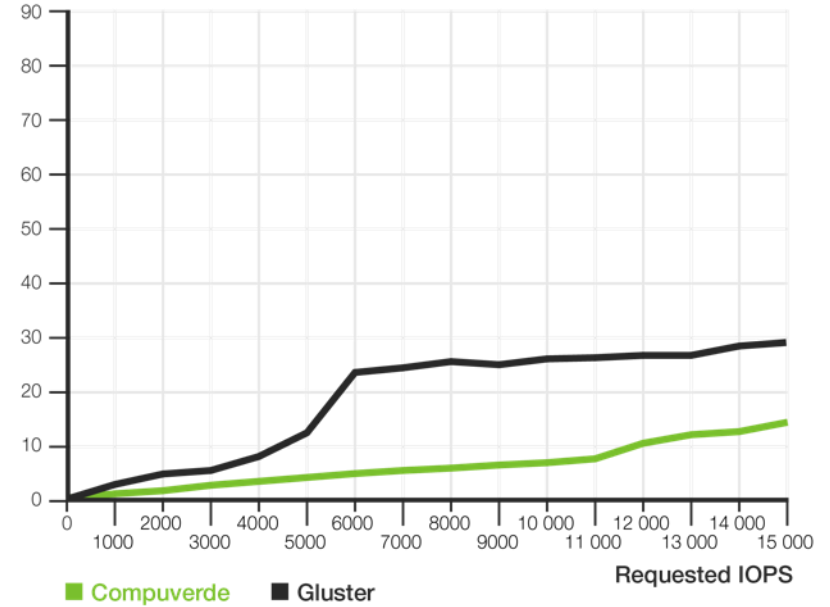
Actual IOPS

(More is better)

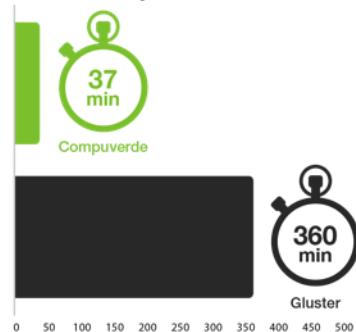


Milliseconds

(Less is better)



Time to recreate missing data (minutes)



1,000,000 files, each 1 MB in cluster

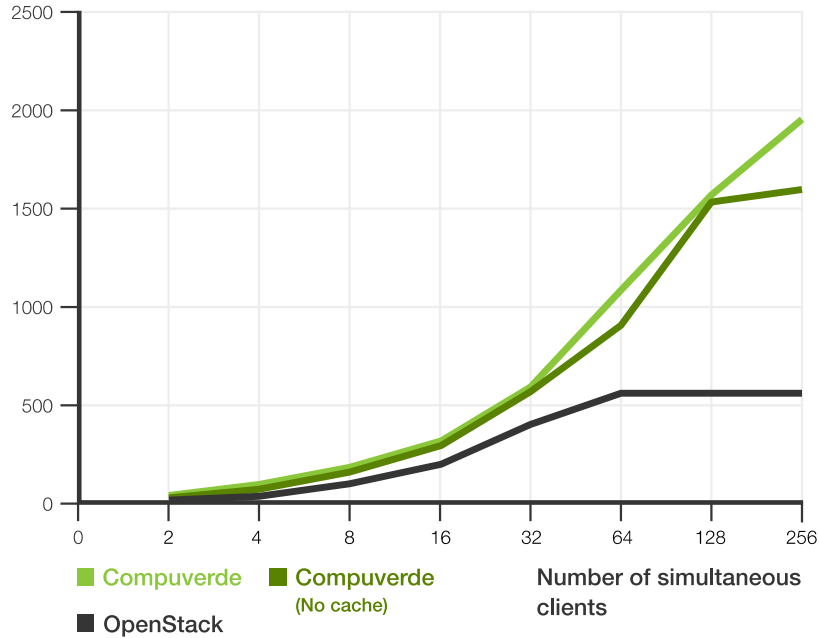
6x

Test setup

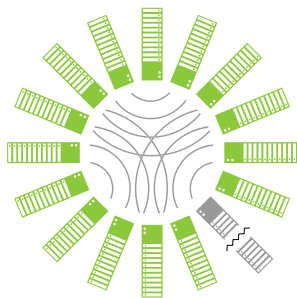
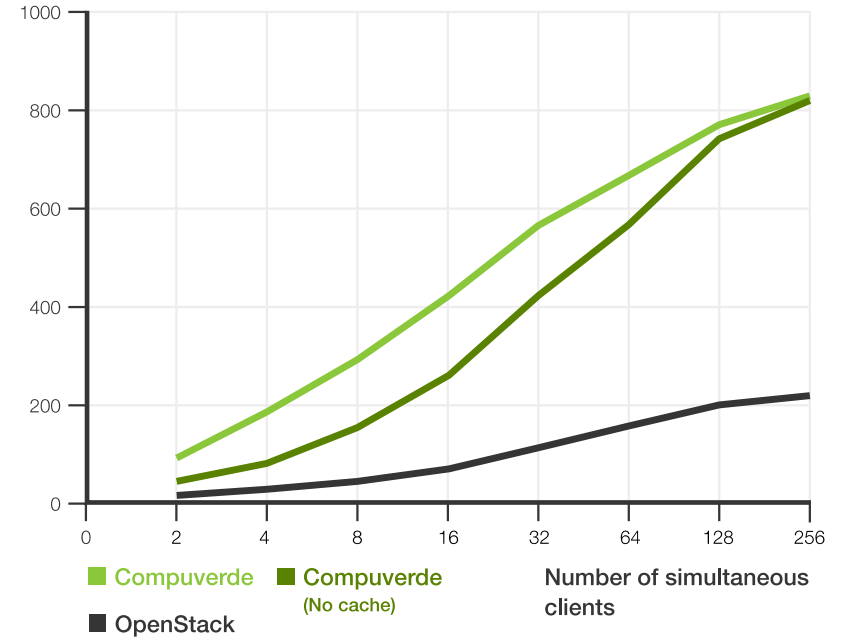
CPU: Intel Xeon E5-2620 (2.0 GHz)
 RAM: 32 GB
 Network: 2 Gbit
 Cache: 1 x Intel 330 (SSD/60 GB)
 Storage: 8 x Seagate SV35 (SATA/3 TB) JBOD

COMPARISON – COMPUVERDE VS OPENSTACK

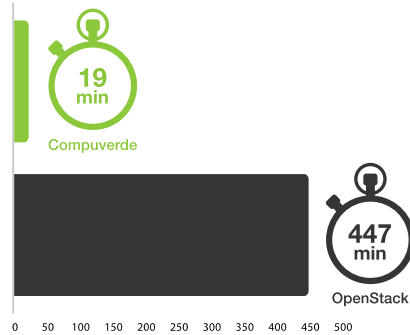
Throughput - Read
(MB/sec)



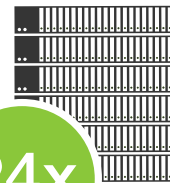
Throughput - Write
(MB/sec)



Time to recreate missing data (minutes)



5,000,000 objects, each 1 MB in cluster

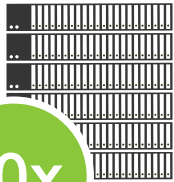
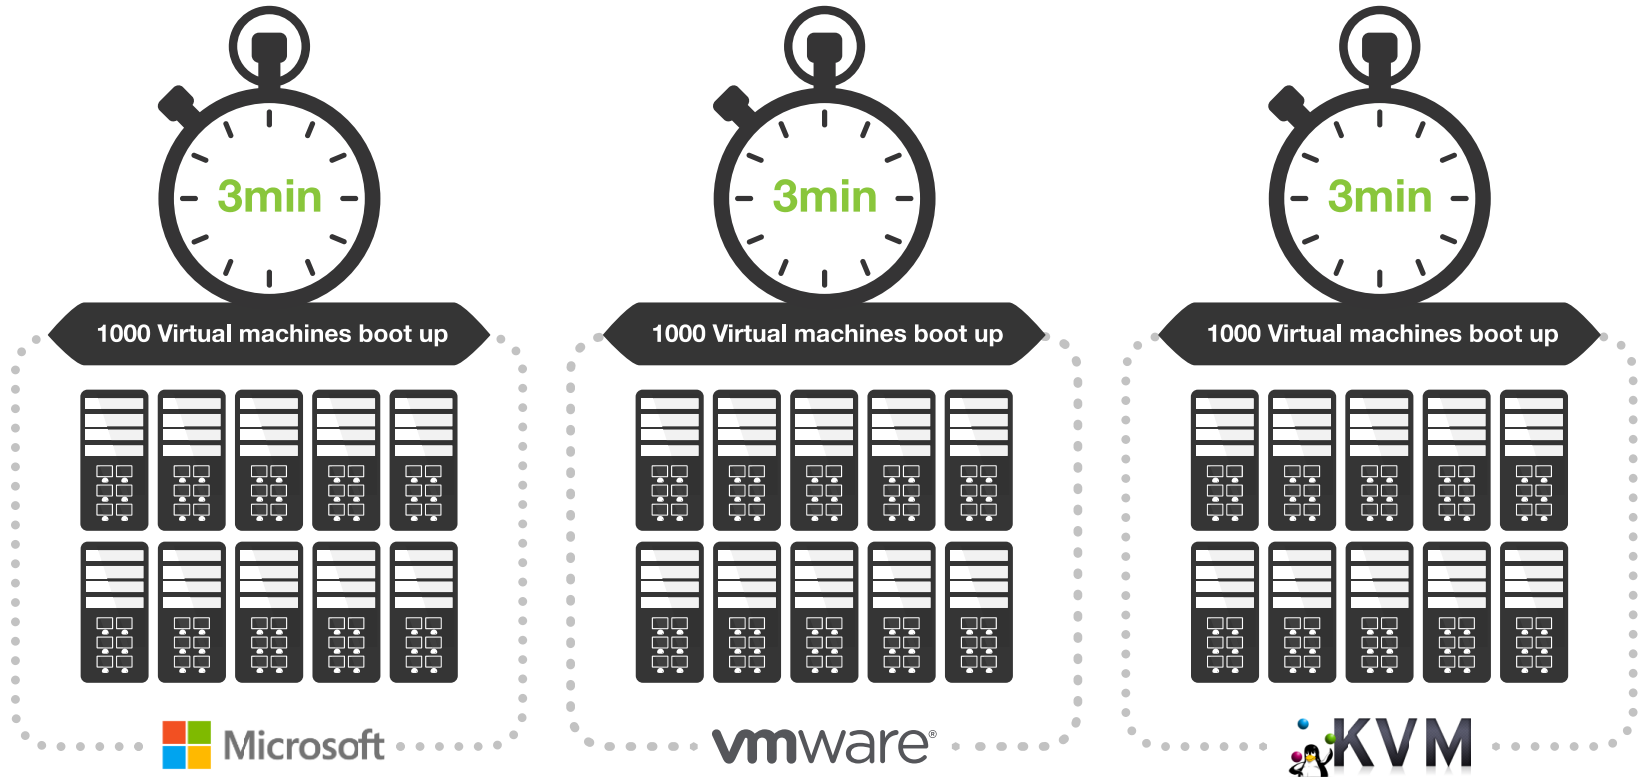


24x

Test setup

CPU: Intel Atom D525
RAM: 4 GB
Network: 1 Gbit
Cache: N/A
Storage: 16 x 2 TB SATA (Green)

PERFORMANCE – VIRTUAL ENVIRONMENT

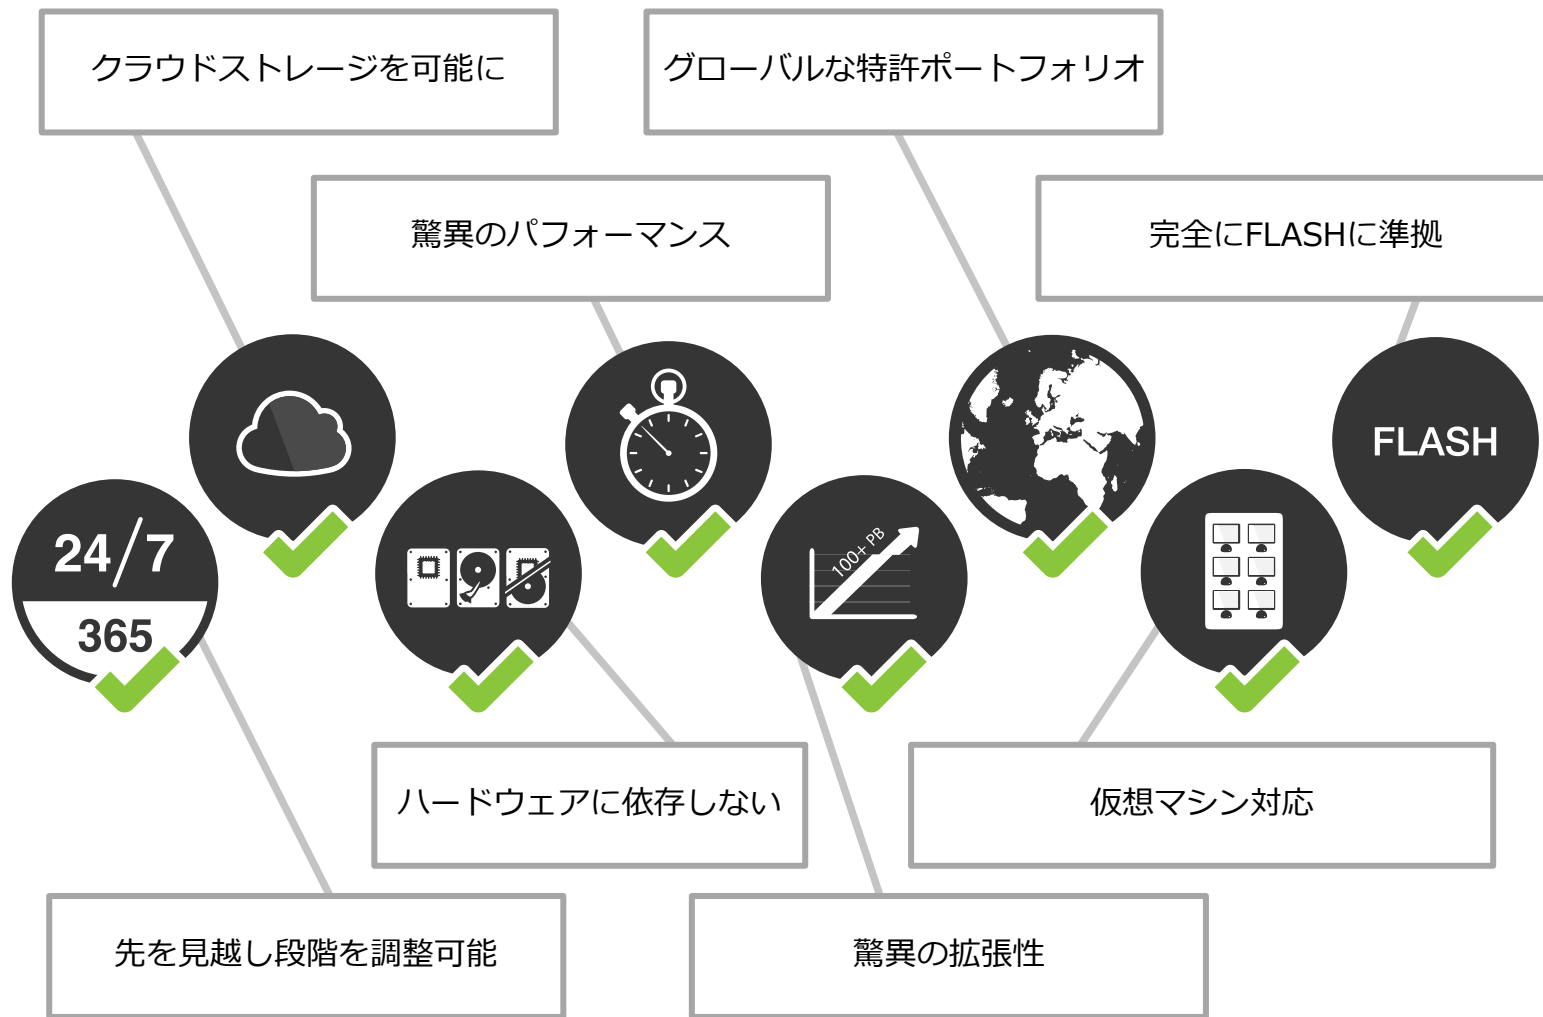


10x

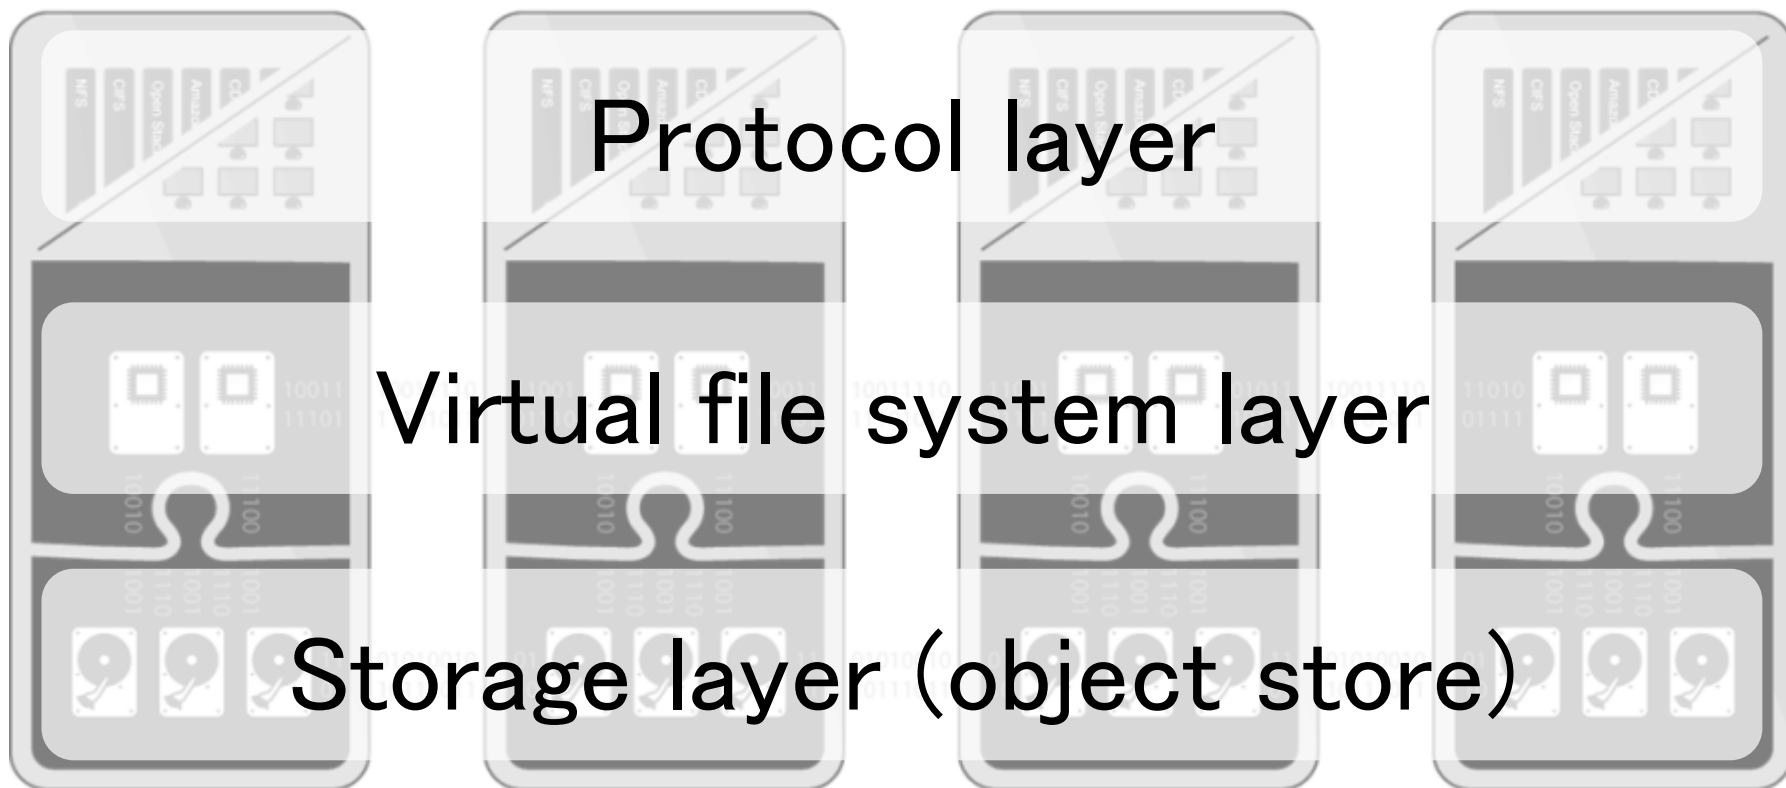
Test setup

CPU:	2x Intel Xeon E5-2620 (2.0 GHz)
RAM:	128 GB
Network:	10 Gbit
Cache:	LSI 9265-8i, 180 GB SSD
Storage:	8 x 320 GB SATA

主な機能

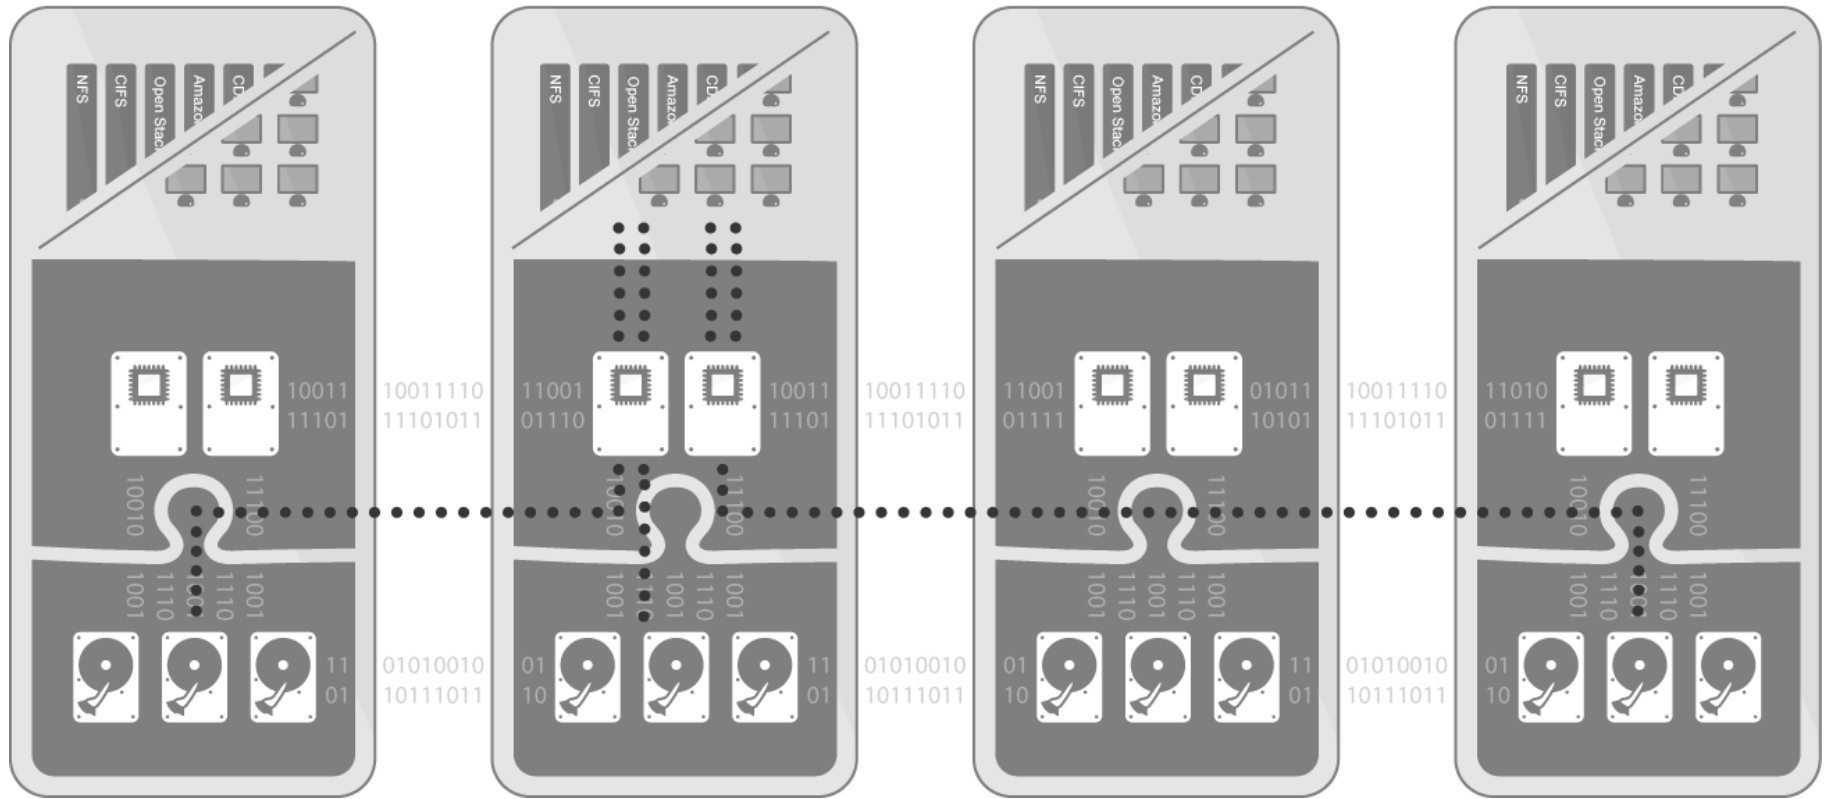


技術資料



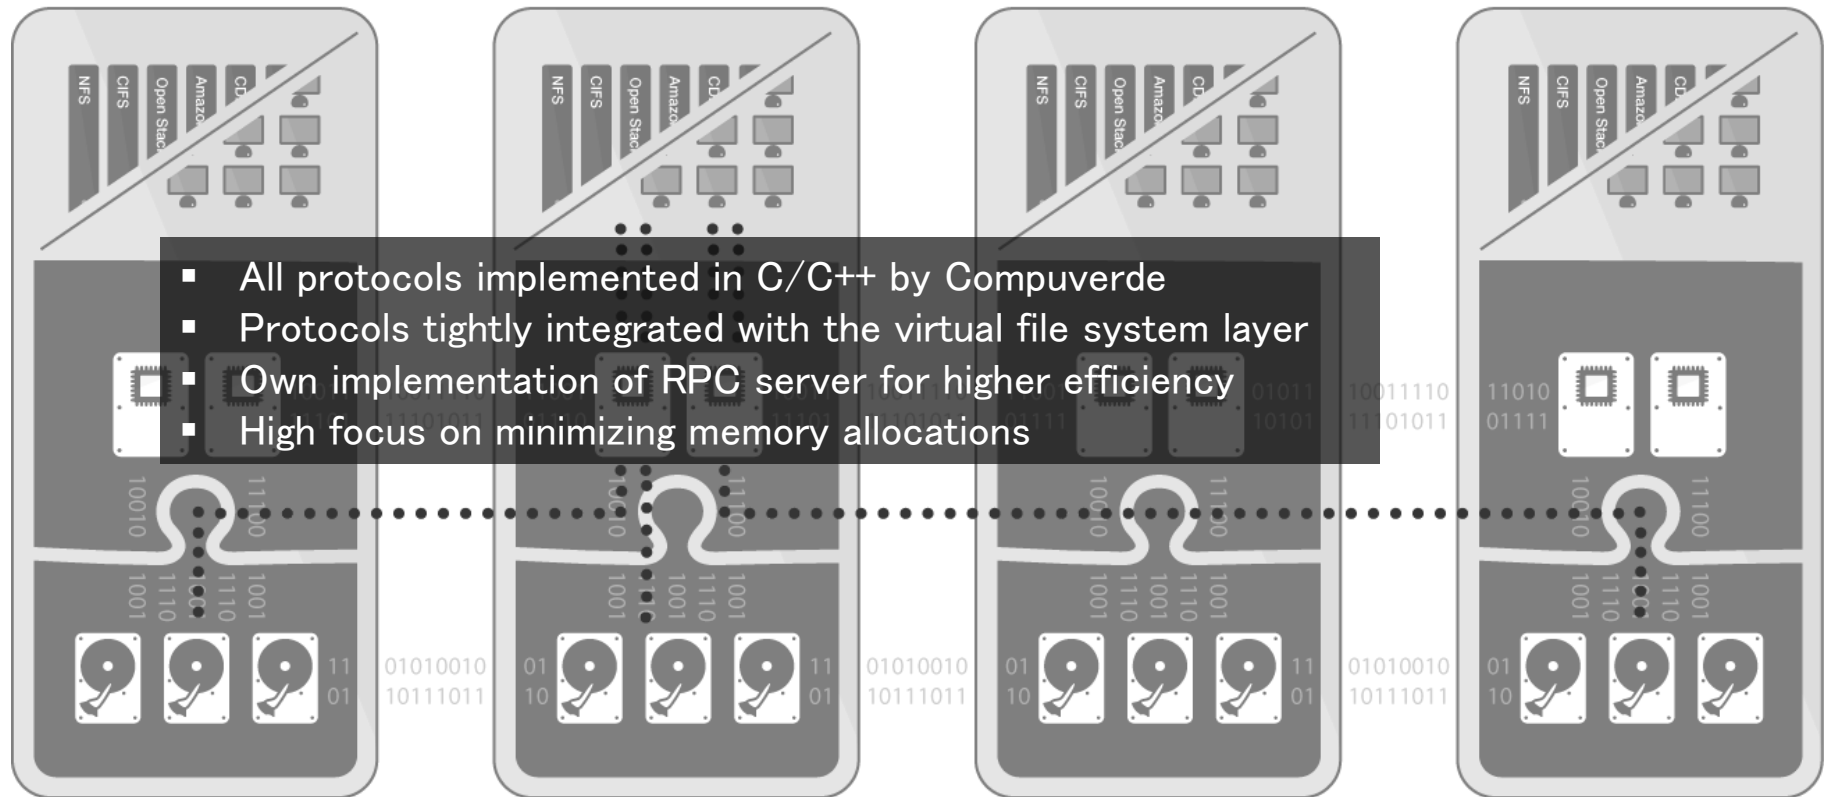
Symmetric architecture

Protocols	CIFS (SMB1/2/3)	NFS 3/4/4.1	OpenStack Swift	Amazon S3	iSCSI
Virtual FS	Support for read and write cache		Strictly consistent file system		
Storage	Extremely scalable		Full elasticity (No use of hash tables)		

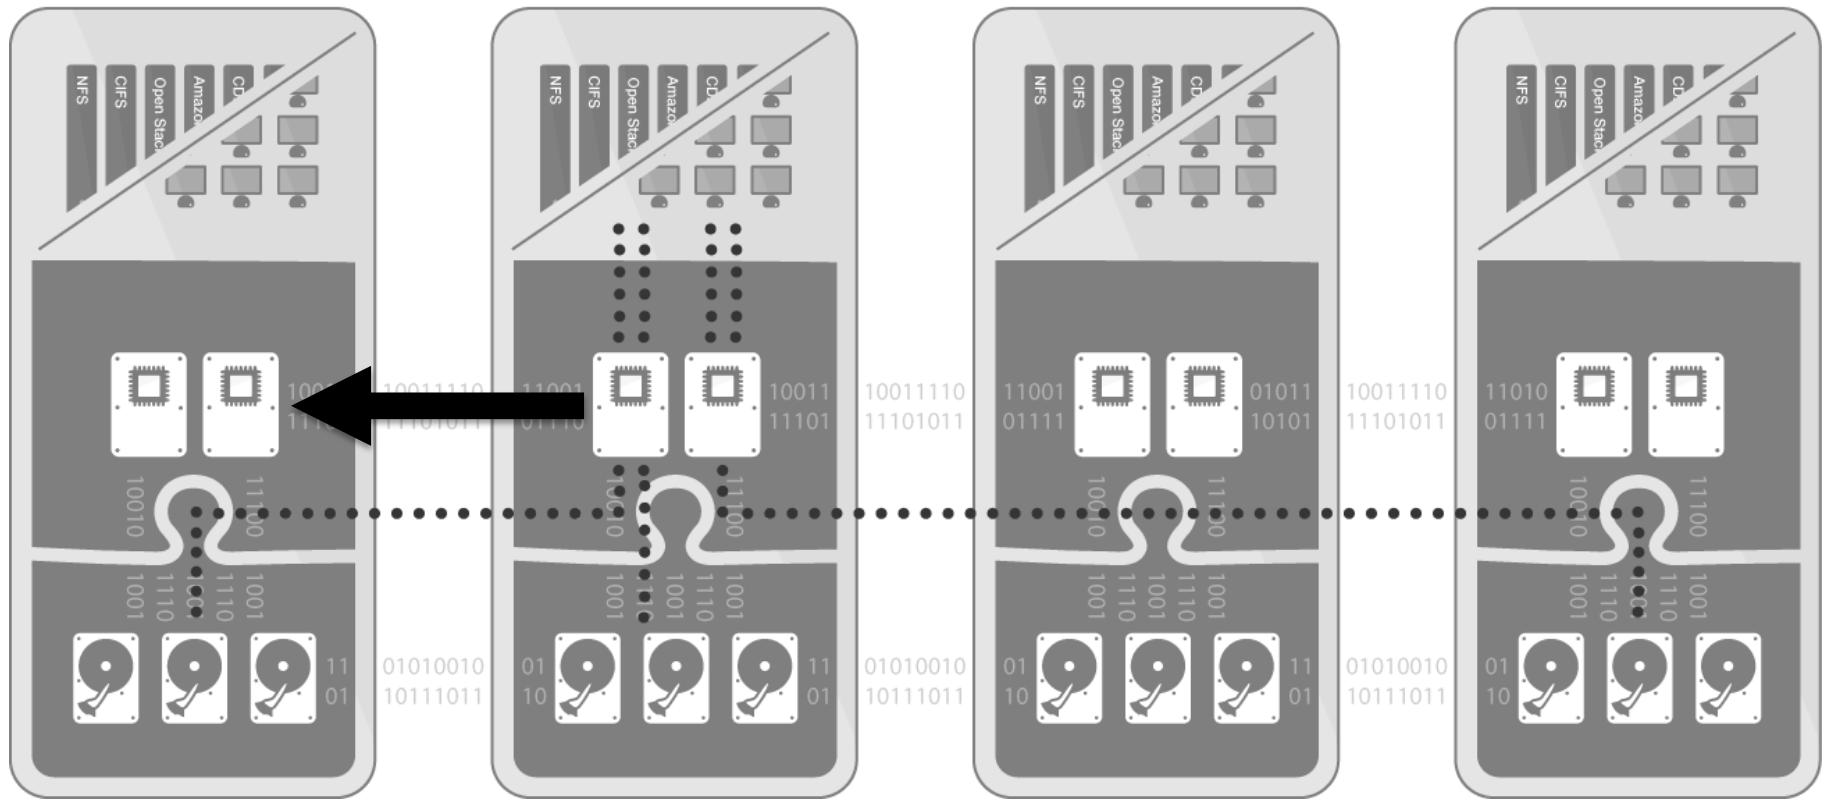


Redundancy	File copies	RAIN-5/6	Erasure coding (CRS)
Write pattern	All targets	Majority + replicate	One target + replicate
Data disparity	Rack designation	LAT/LONG/ALT	

Protocols

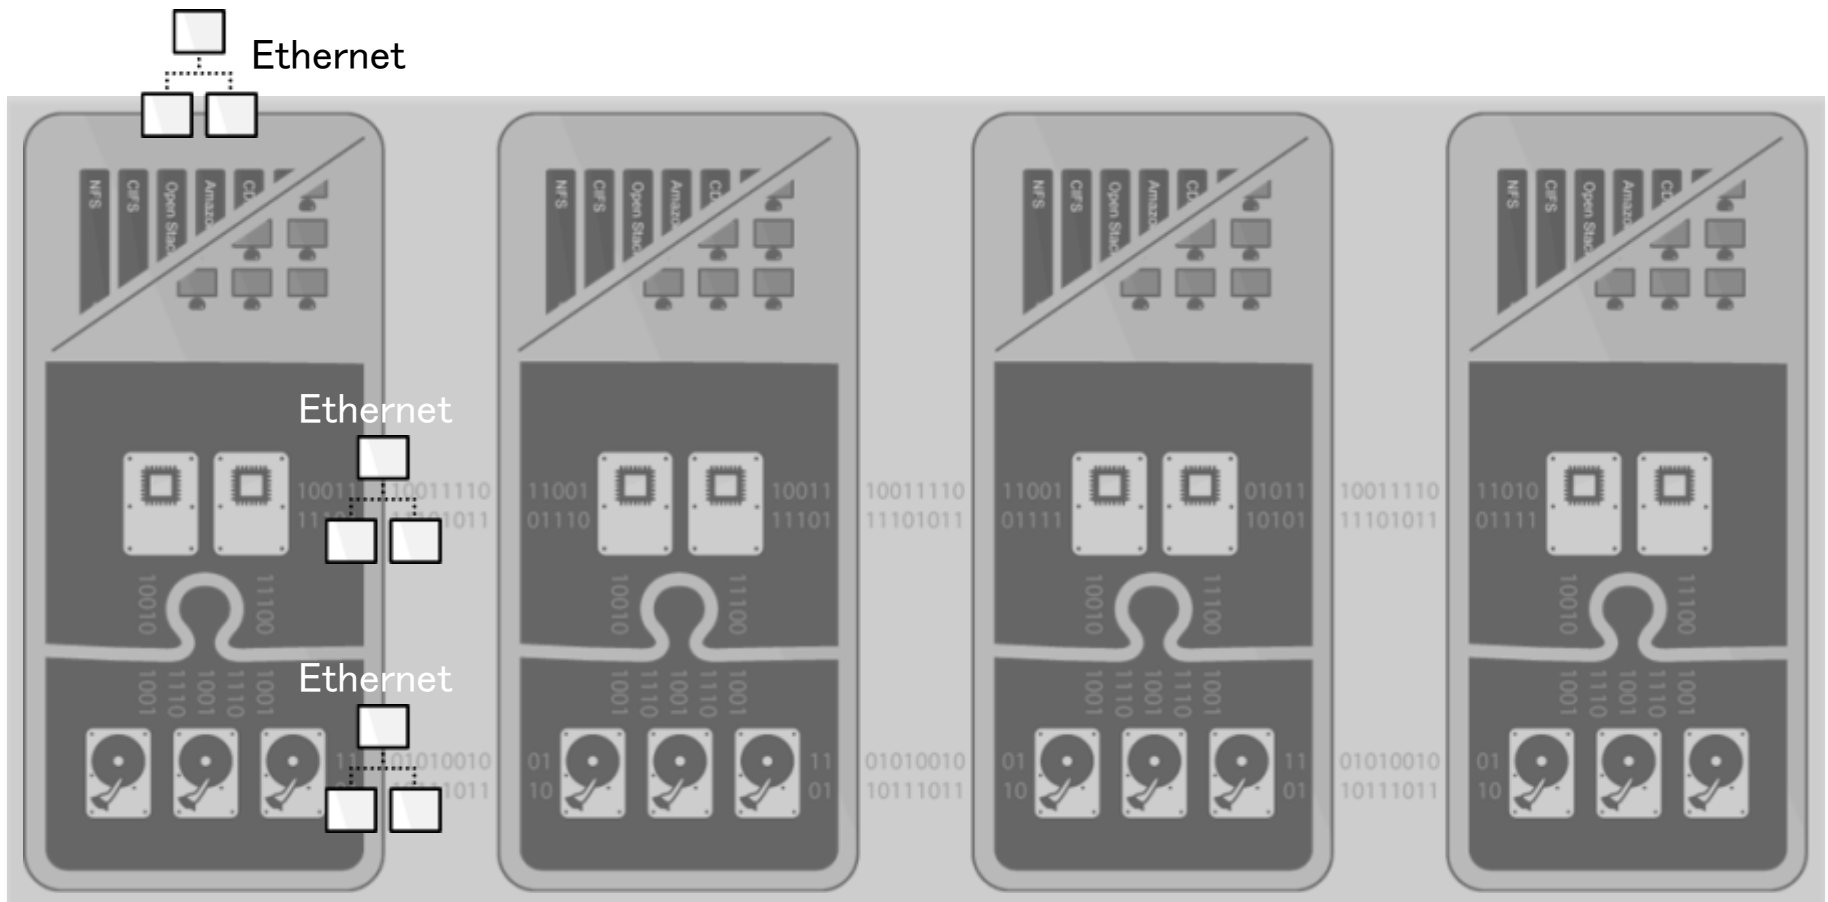


Cache



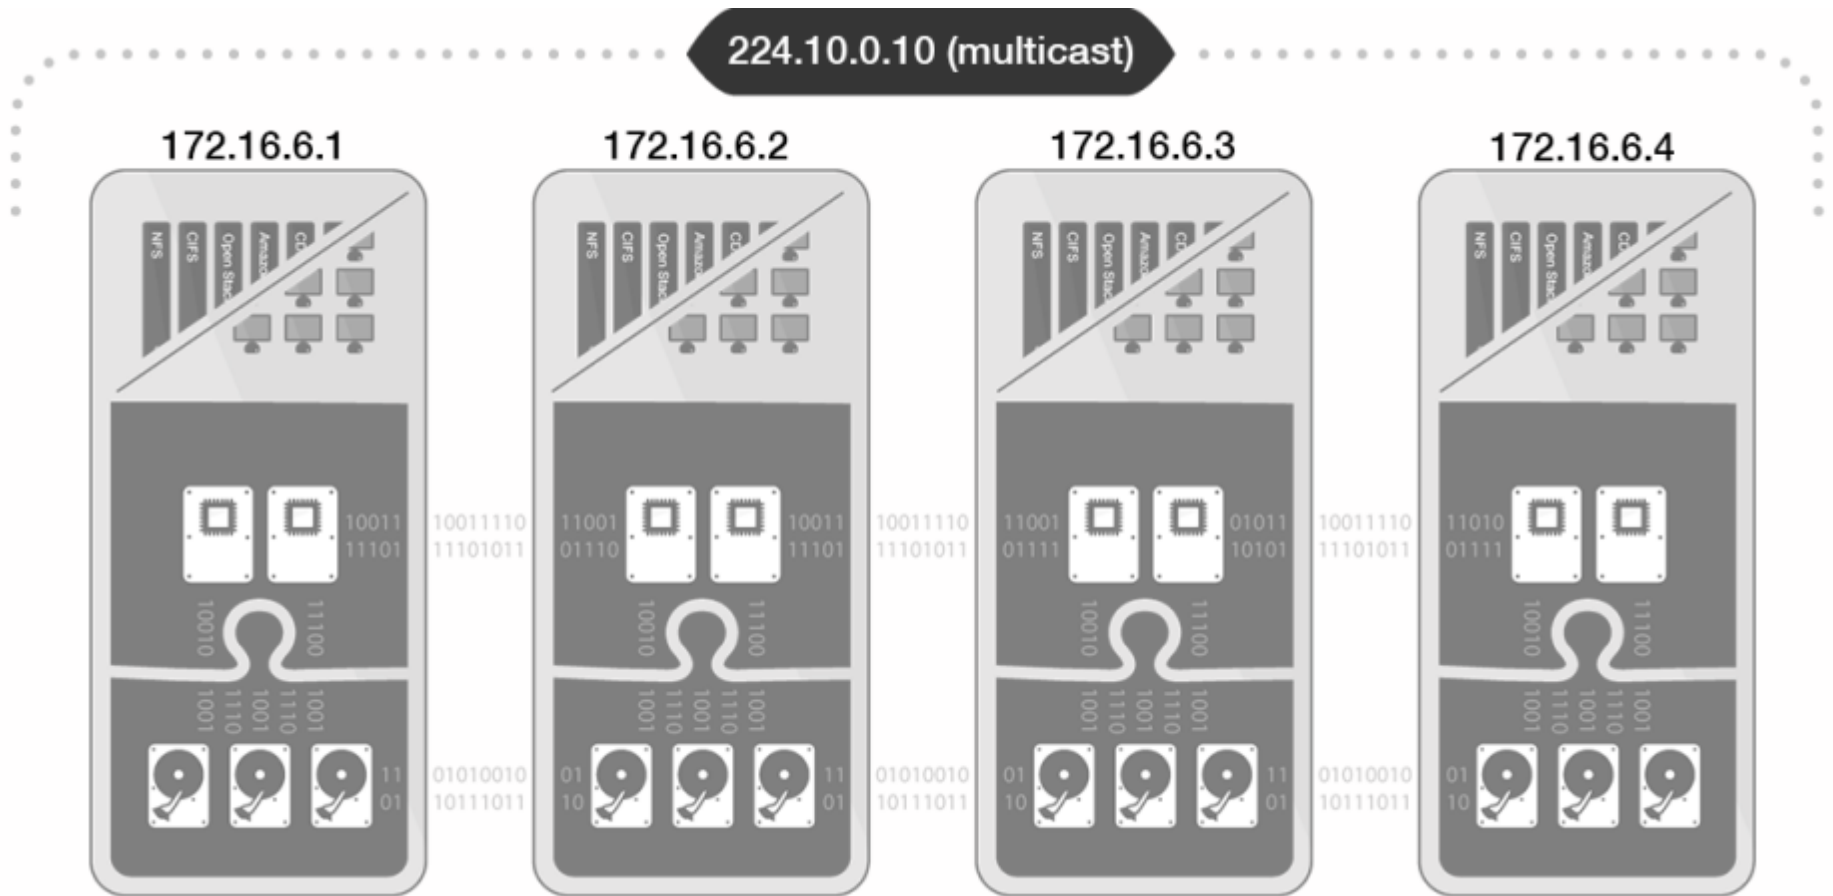
- Support for cache mirroring
- Cache partners are selected at file level
- Supports write cache for a file on single server
- Supports read cache for a file on multiple servers

Network

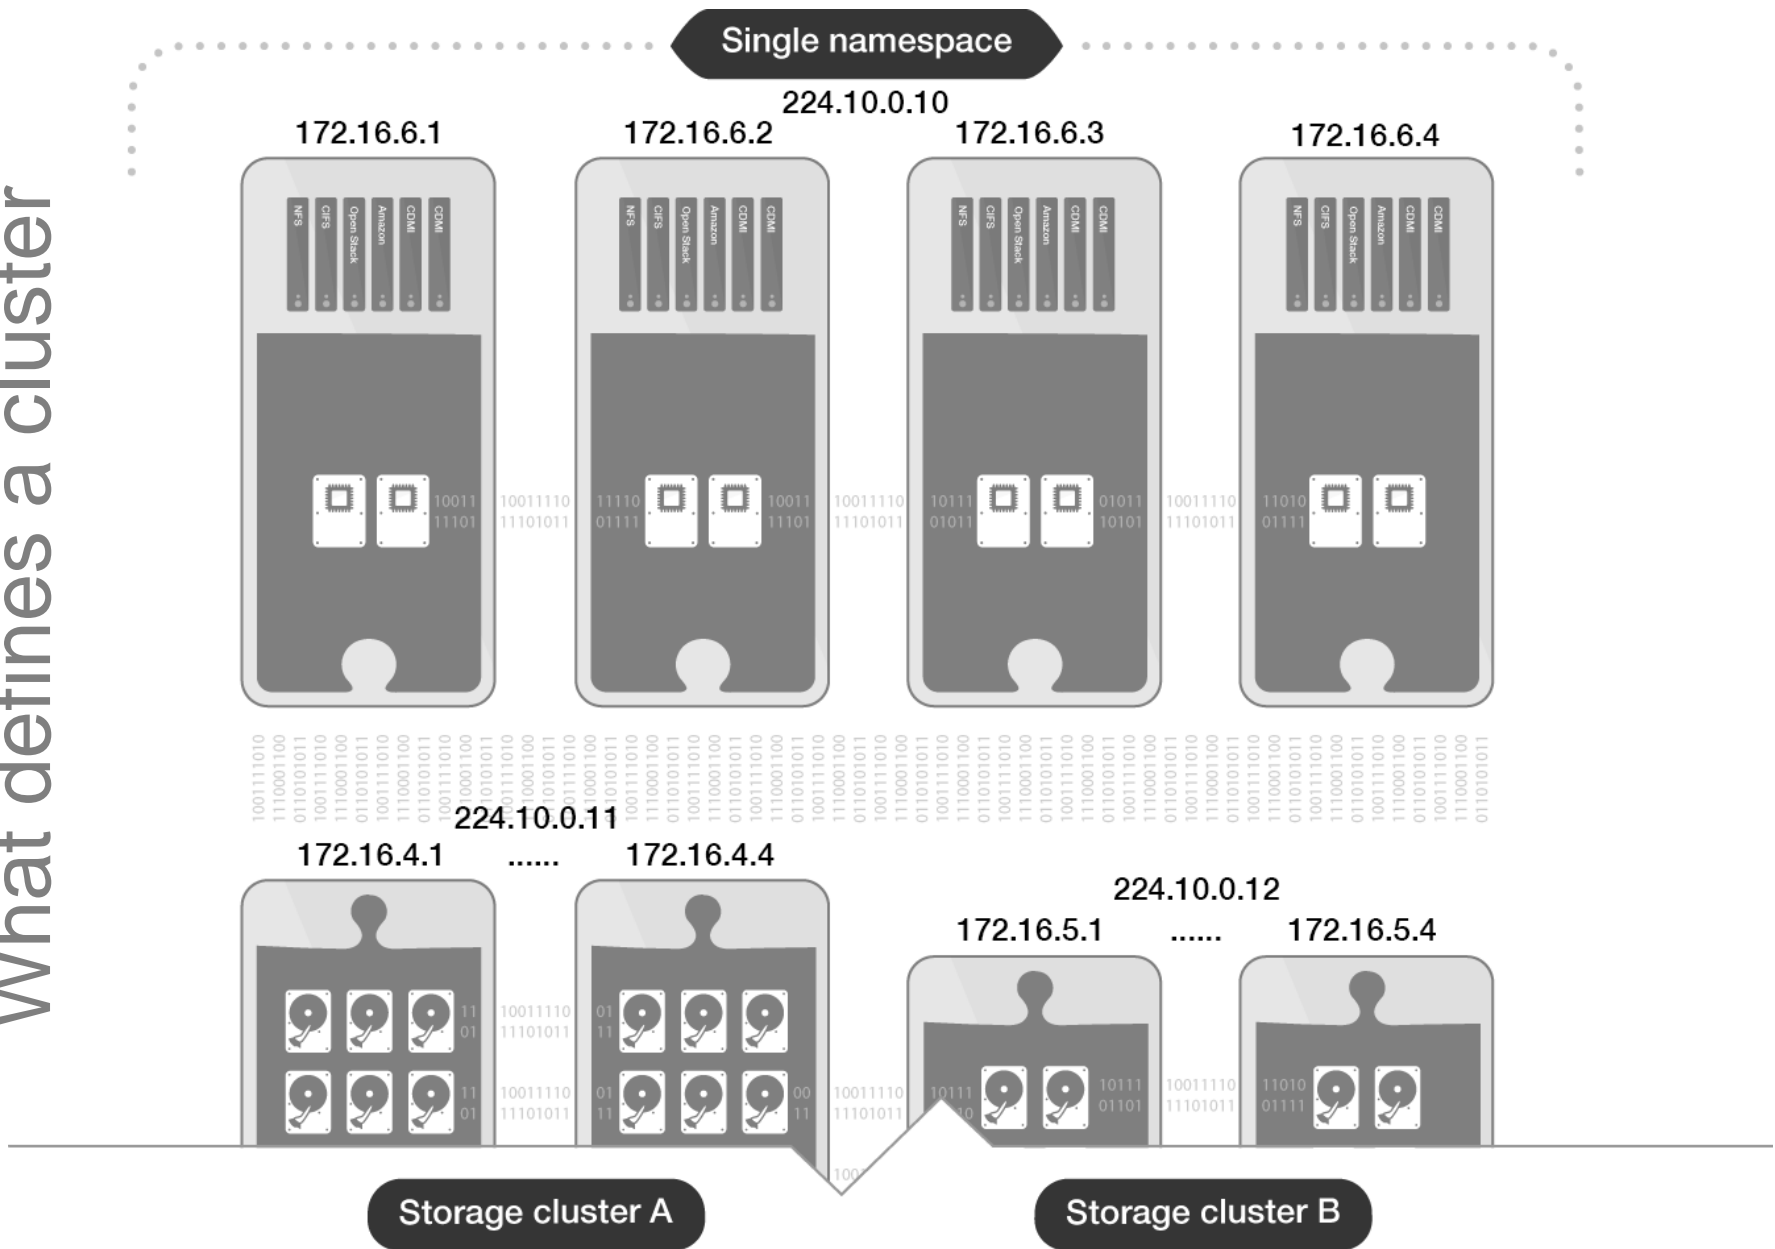


- 1–3 network ports are required depending on configuration
- TCP unicast and UDP multicast are used for communication
- All nodes have static IP-addresses
- Failover VIRTUAL IP

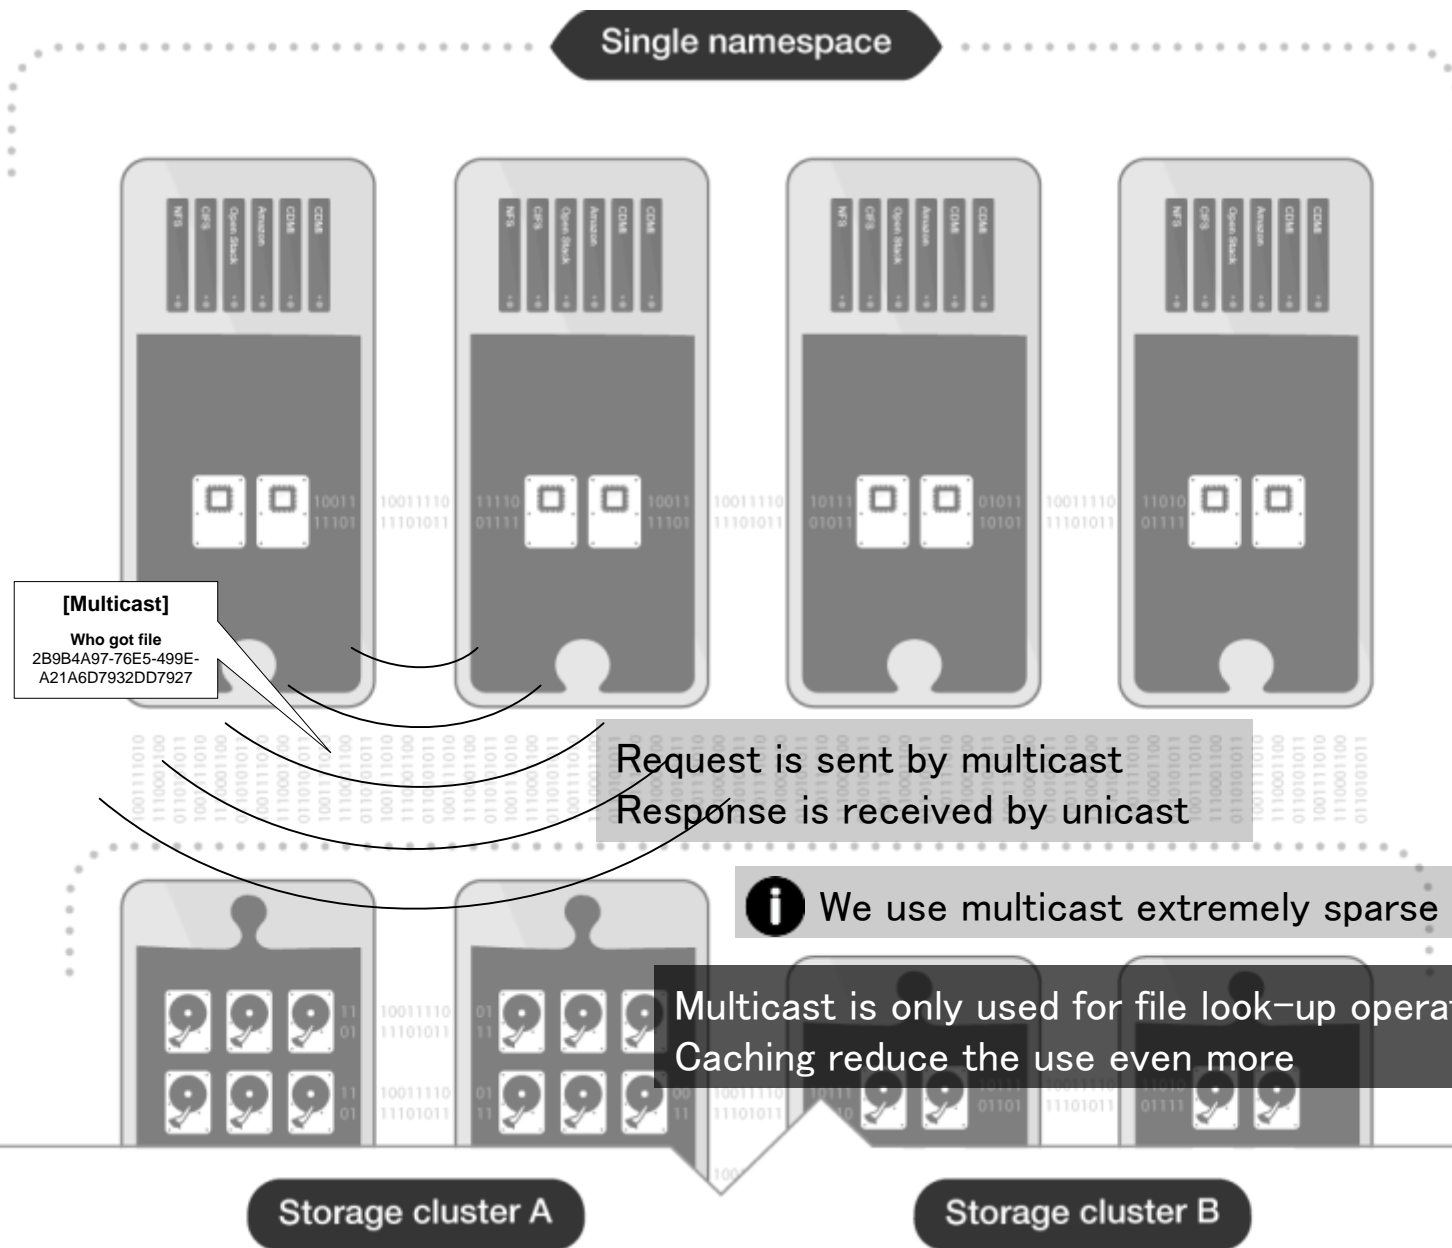
What defines a cluster



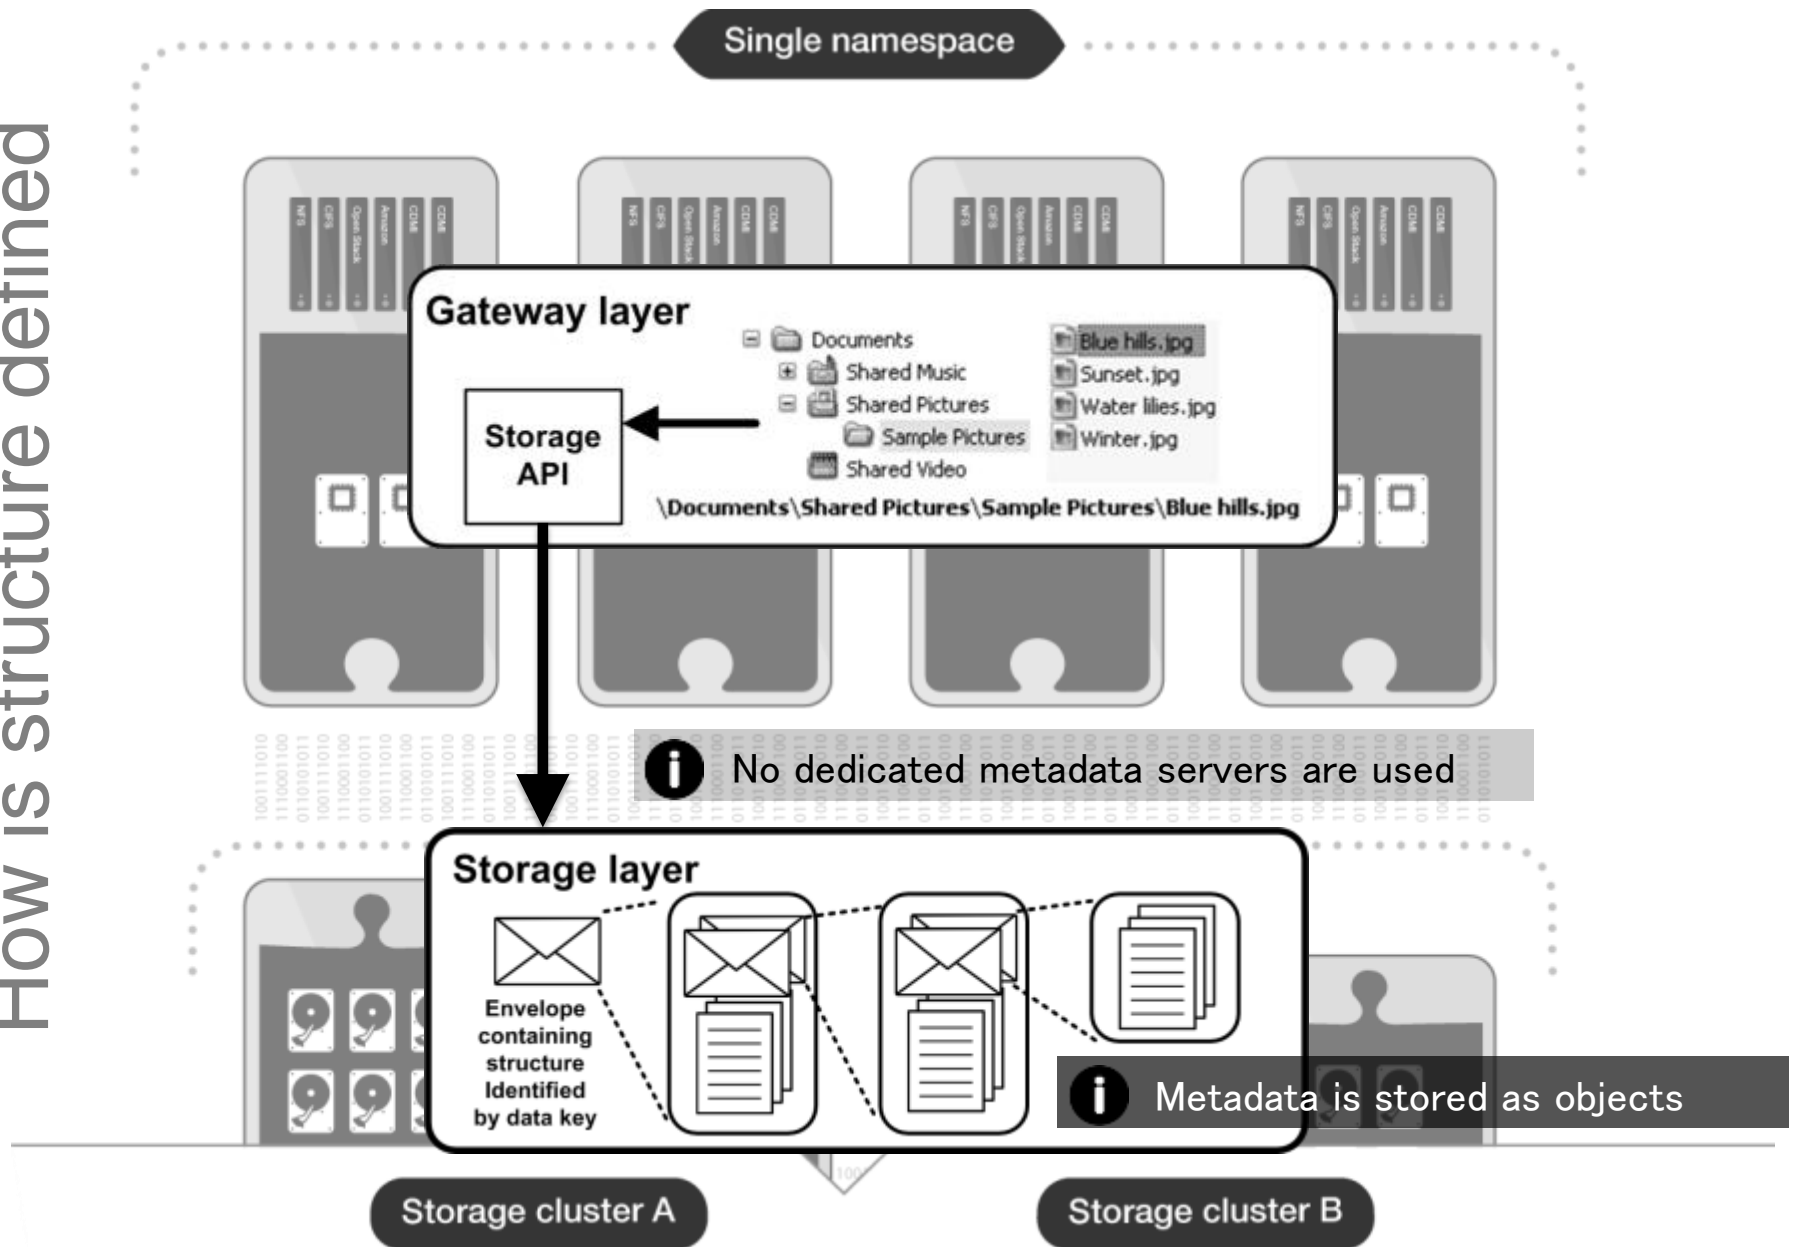
What defines a cluster



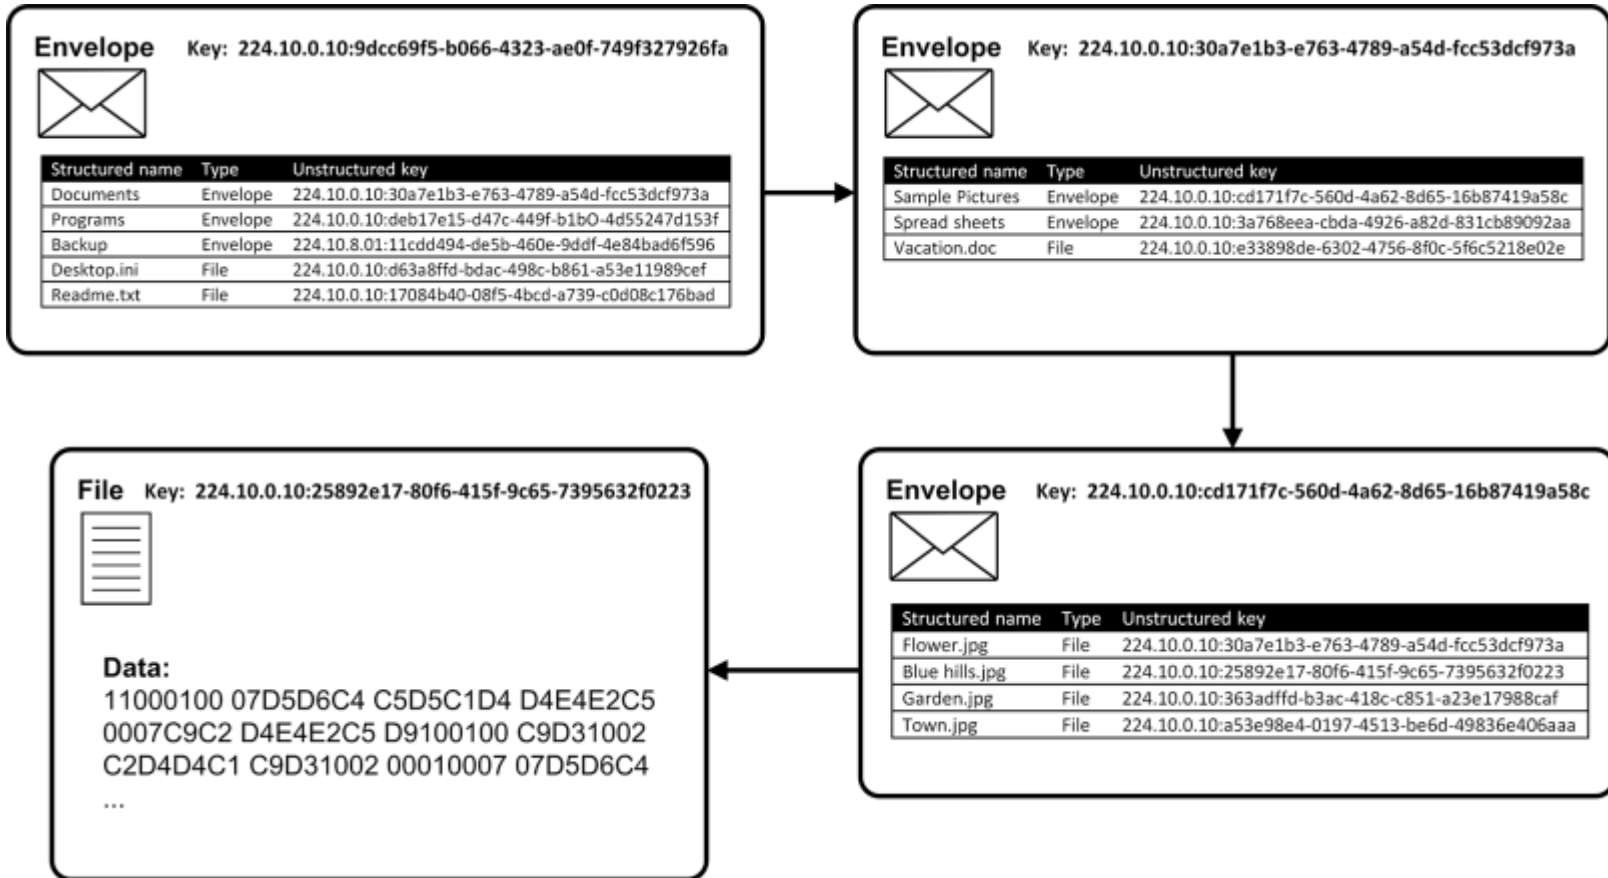
How do we find files/objects



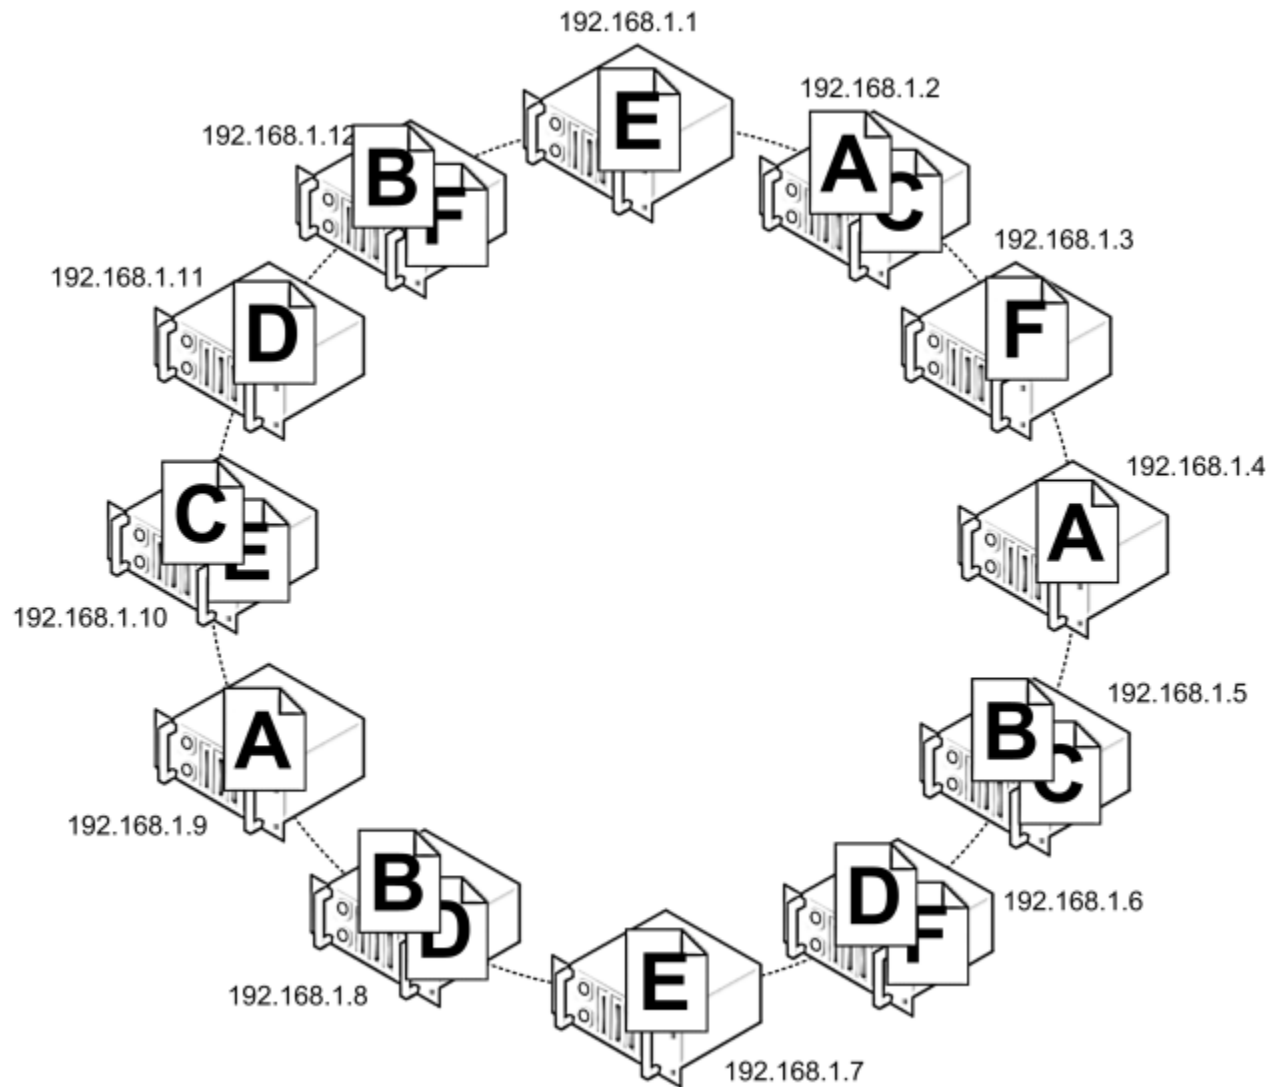
How is structure defined



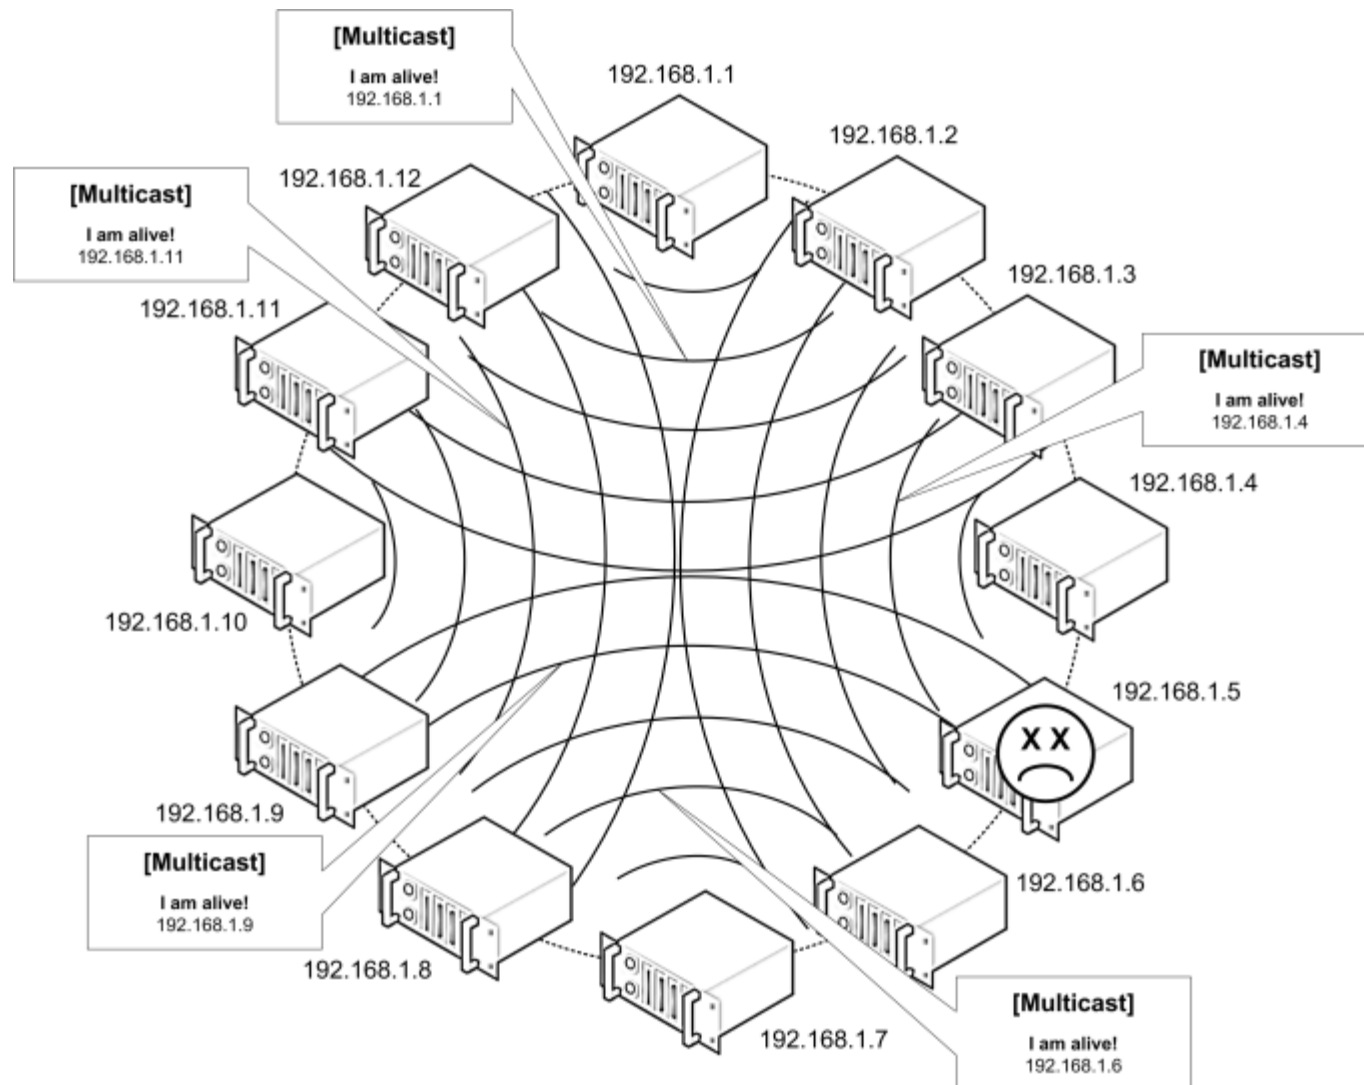
How is structure defined



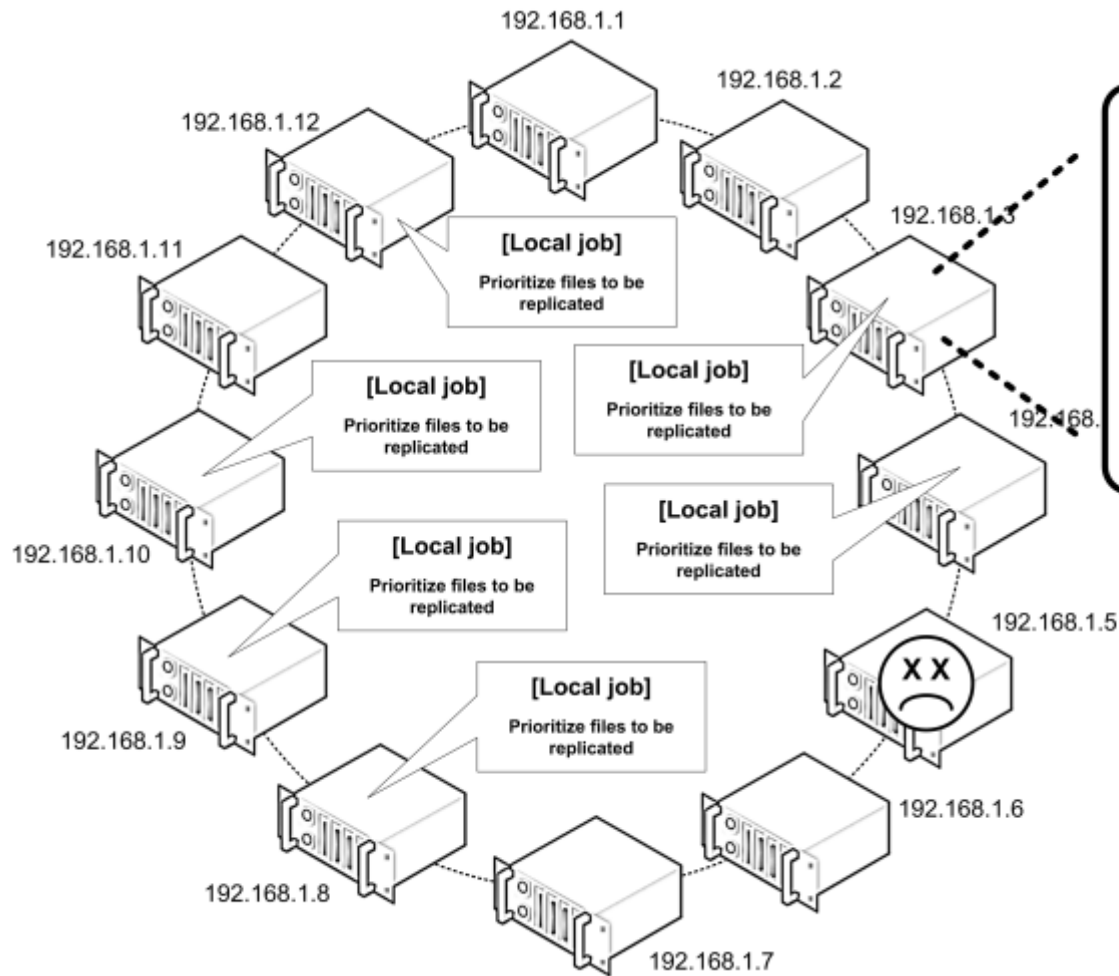
How is data distributed within a cluster



How does nodes keep track of each other



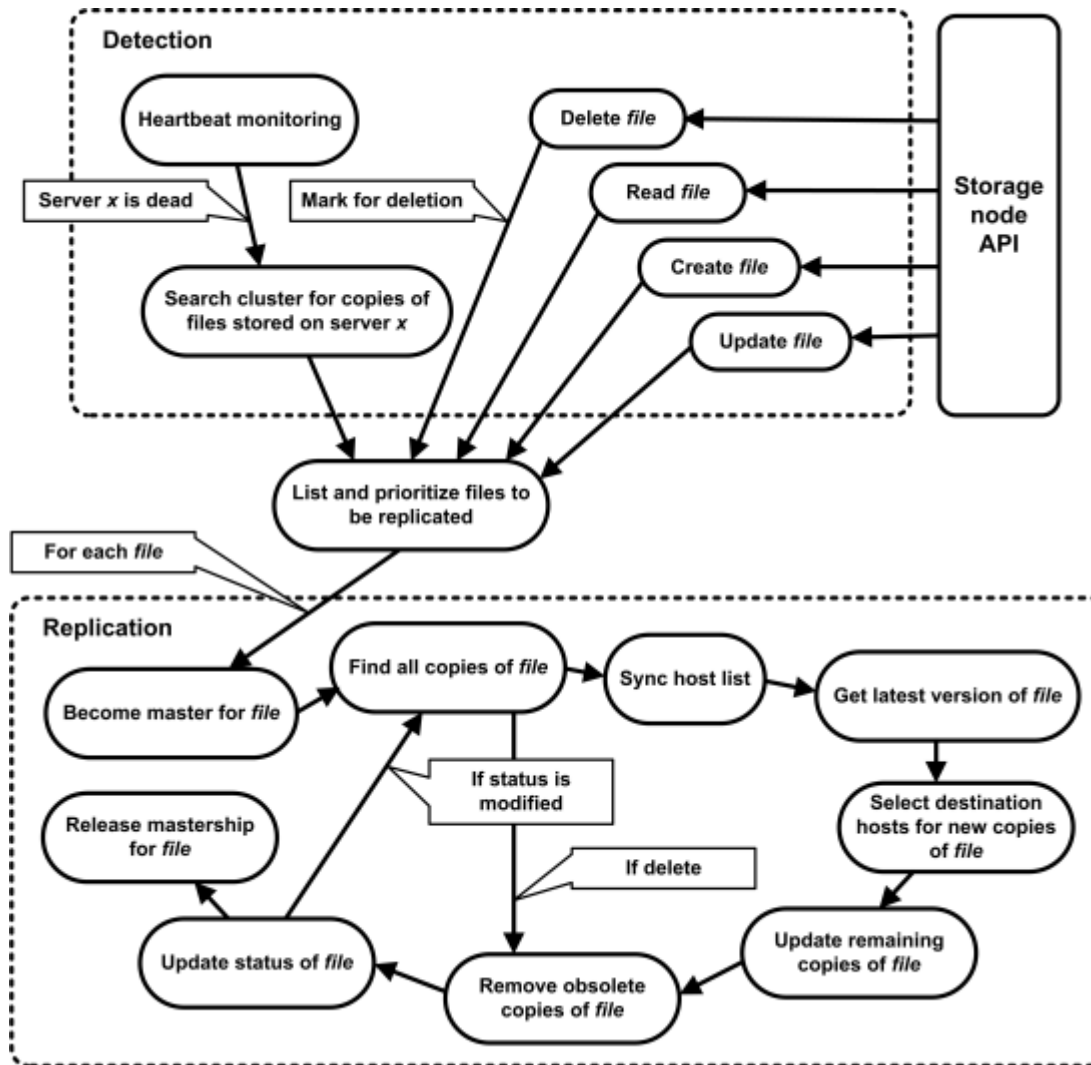
Self healing



Storage node Replication priority lists

High	Med	Low
A4099325-FD-4D... 6A206642-EB-47...	252212D8-87-4D... 717B9C3C-C2-47... 5A8E6CA8-2B-4C...	613DB68B-0A-41... 867230C1-88-4E... CBDB0010-42-45... 070326D9-03-41... E7A51F85-59-42...

Self healing





コアマイクロシステムズ株式会社

Core Micro Systems, Inc.

URL: <http://www.cmsinc.co.jp/> Mail: sales@cmsinc.co.jp

TEL: 03-5917-6451 IP Phone: 050-5558-5410 FAX 03-5917-6452

本社 〒173-0026 東京都板橋区中丸町11-2 ワコーレ要町ビル9F