# SUSE® Enterprise Storage

## Vorstellung

Thomas Grätz
Sales Engineer
thomas.graetz@suse.com

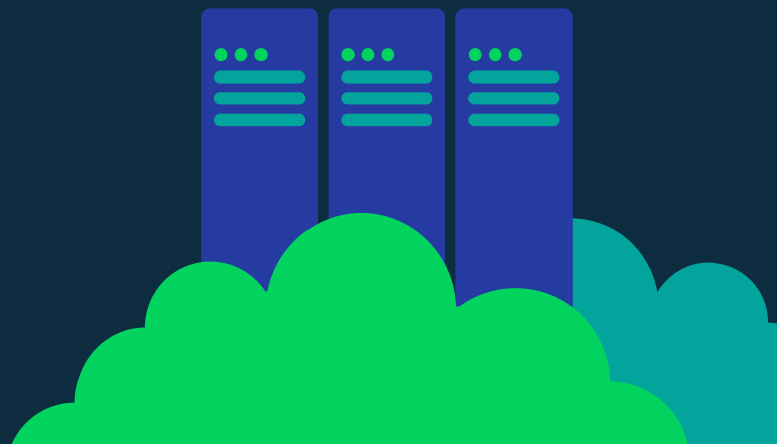

# Agenda

- Trends im Storage Markt
- SUSE Enterprise Storage
- Features und Roadmap
- Konfiguration
- Anwendungsfälle
- Zusammenfassung

# Trends im Storage Markt

# Die Storage Evolution…

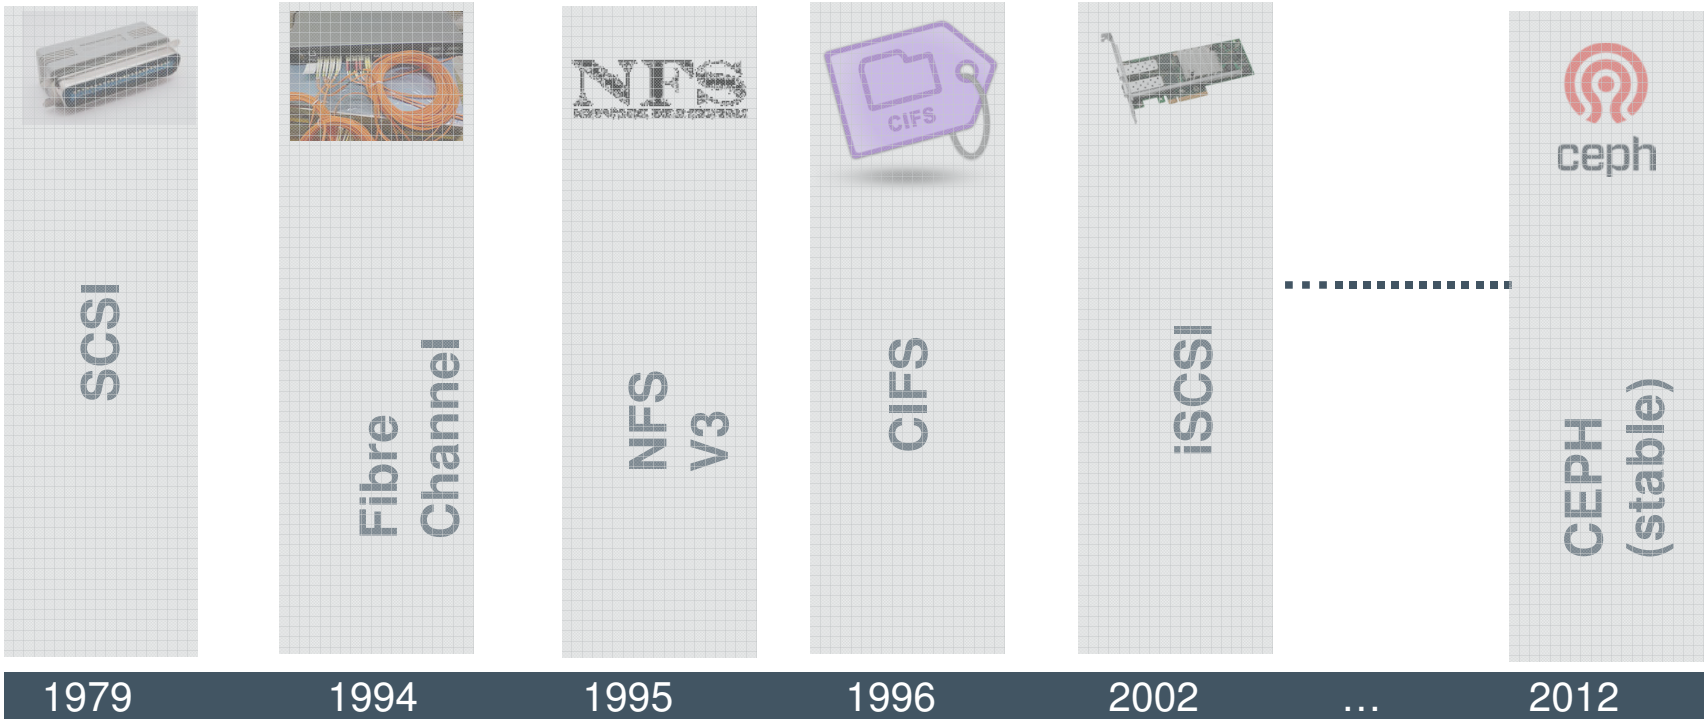


SUSE

# Die Storage Evolution…

## Dienste im Datacenter früher…





# Die Storage Evolution…

## Dienste im Datacenter heute…


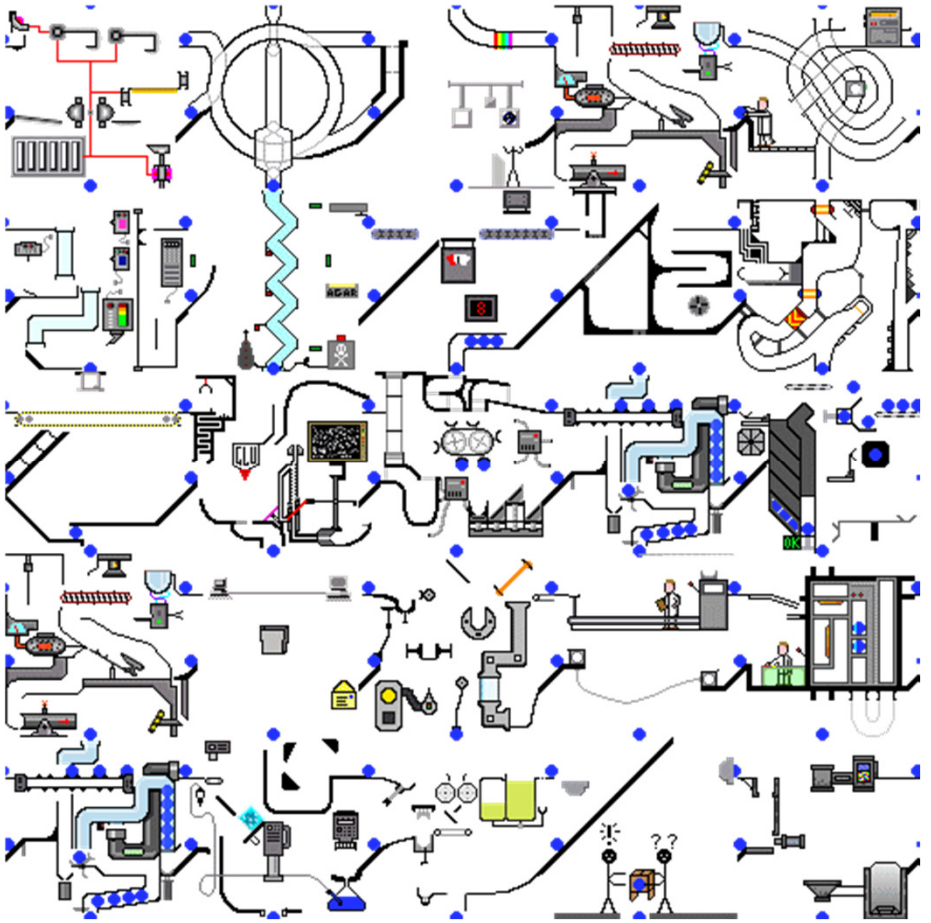

SUSE

# Und dazu das Datenwachstum…

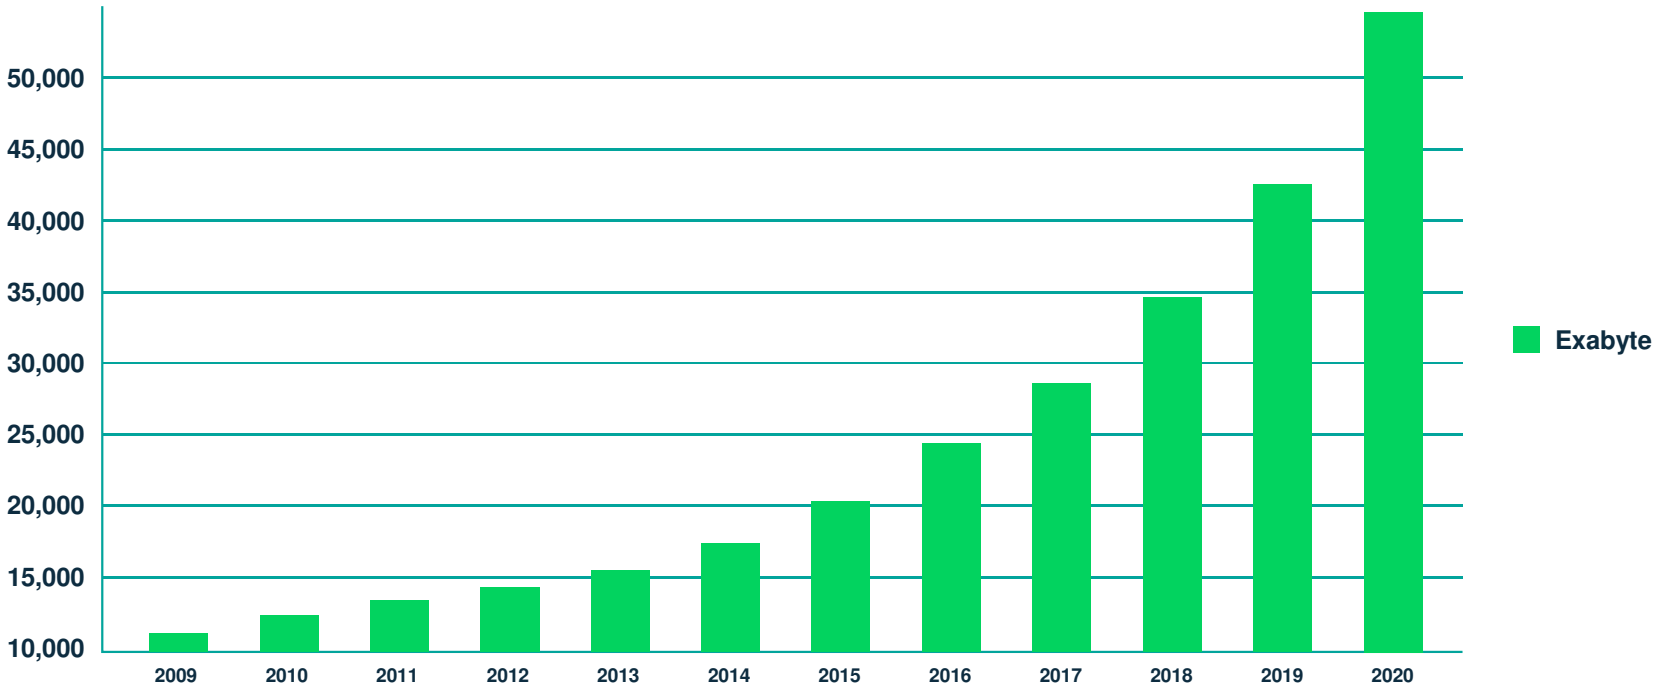


IDC Analyze the Future

Quelle: IDC / http://www.tecchannel.de/a/software-defined-storage-wird-zum-muss,2067830

# SUSE Storage Umfrage- Datenwachstum

Datenwachstum - Ja/Nein

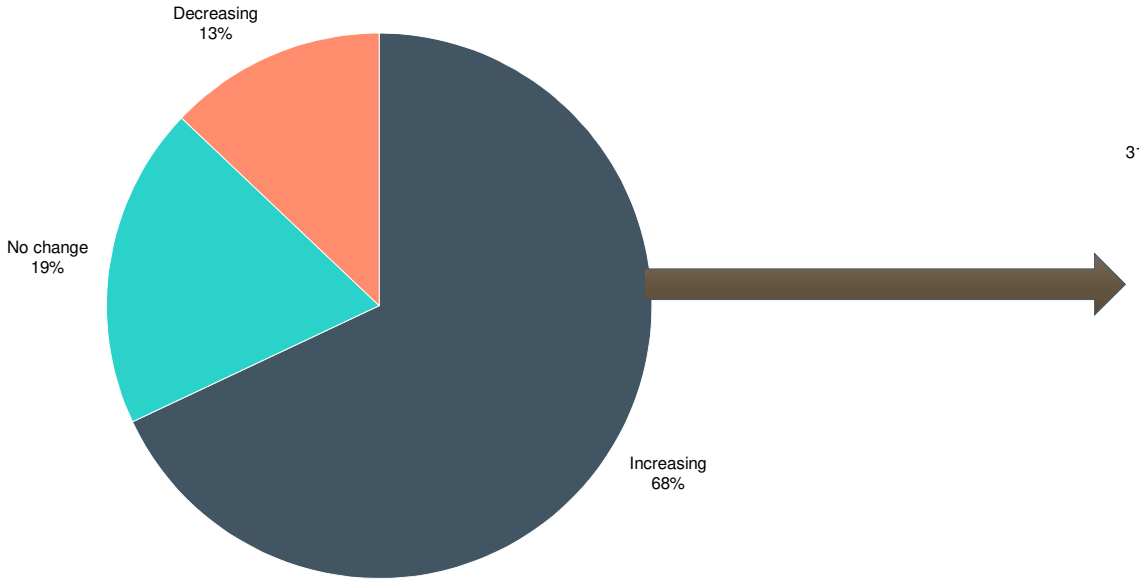


Wachstum in %

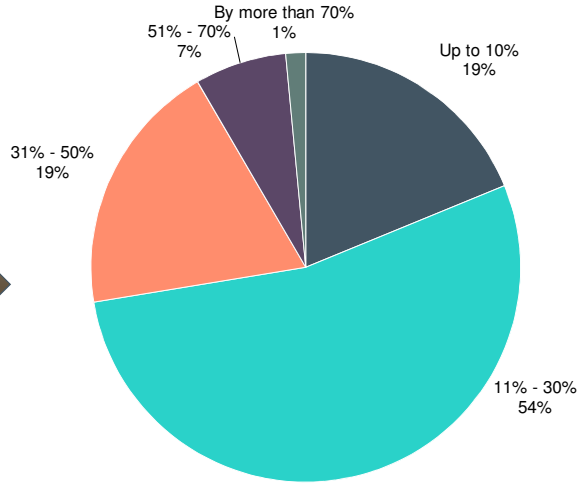


Das durchschnittliche Datenwachstum in DACH beträgt **27%**

*1202 senior IT decision makers across 11 countries completed an online survey

# SUSE Storage Umfrage 2016 - Herausforderungen


*1202 senior IT decision makers across 11 countries completed an online survey

# Aufteilung der Enterprise Daten

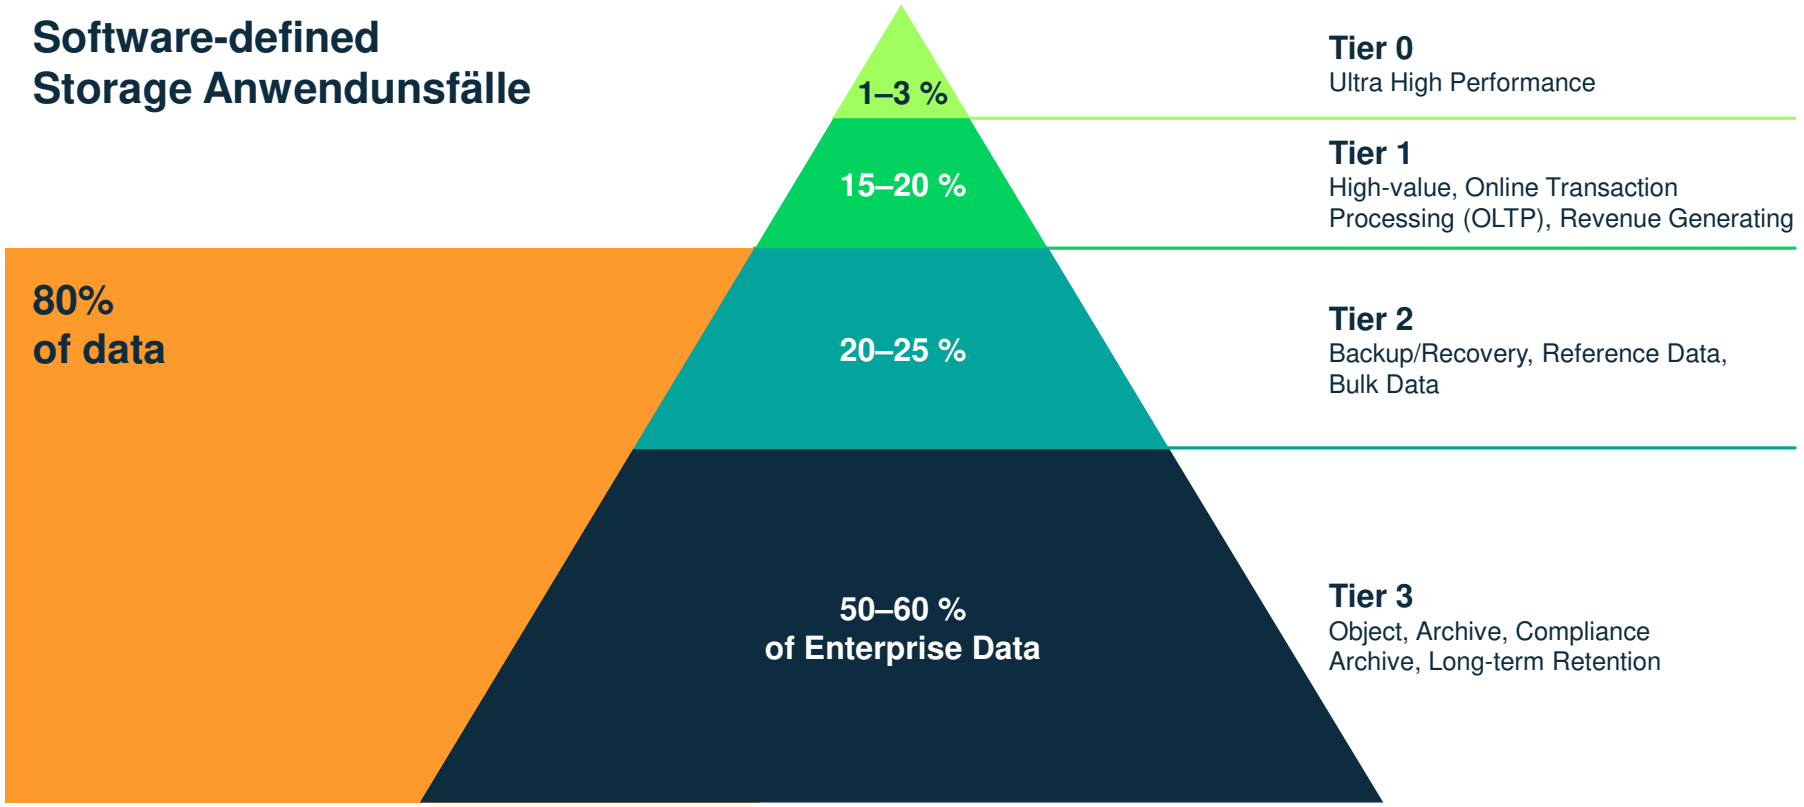

# Positionierung - SUSE Enterprise Storage

Datacenter Tiering

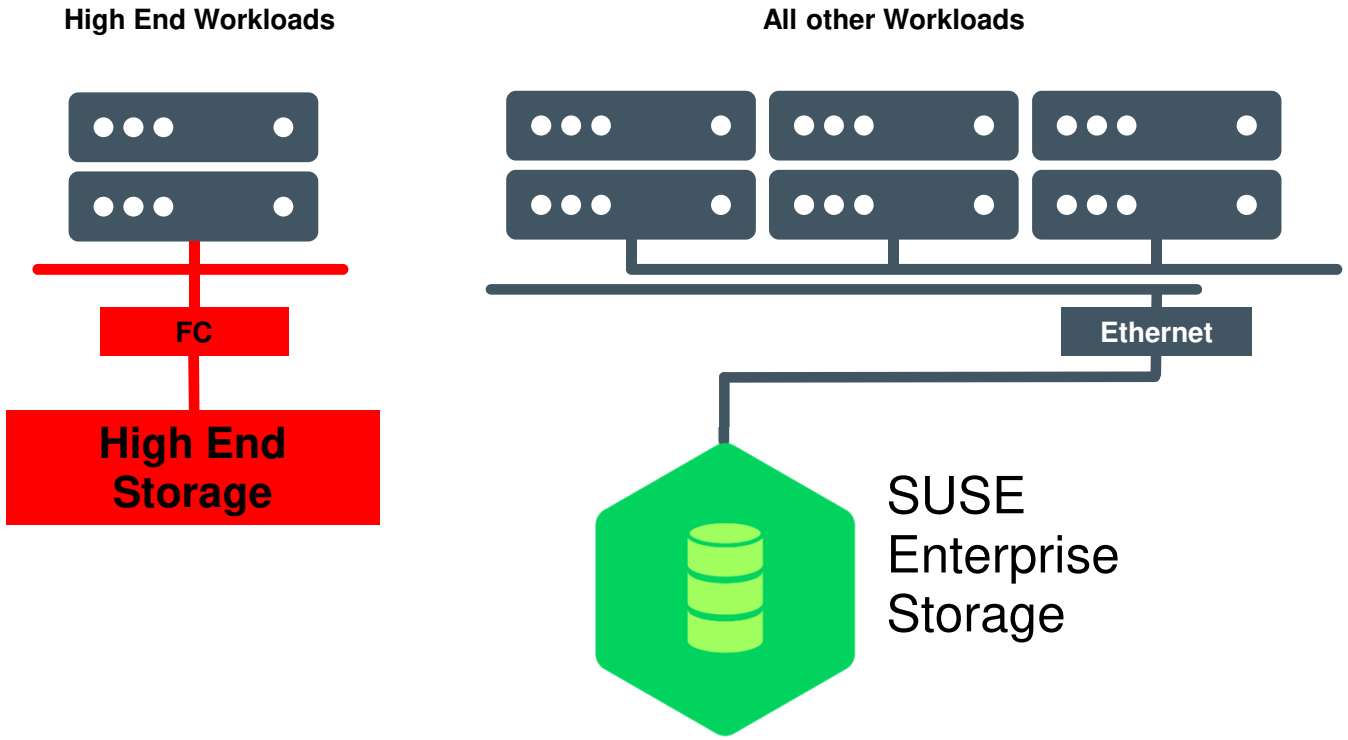

# Können traditionelle Systeme die Antwort sein?

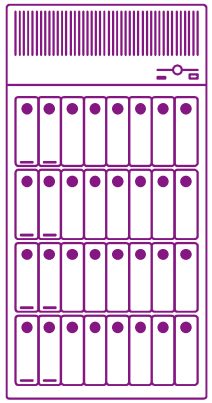

Keine nahtlose Skalierung – daher nicht zukunftssicher

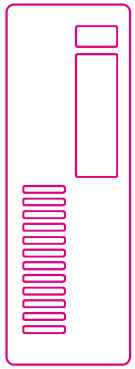

Zu teuer

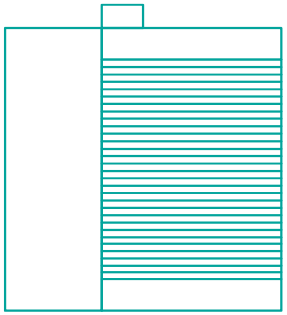

Nicht Cloud fähig

# SUSE Enterprise Storage

# SUSE Enterprise Storage

Eine **hochskalierbare, softwarebasierende** Storagelösung, die Unternehmen den Aufbau einer **kosteneffektiven** Speicherplattform, basierend auf **Standard Serverhardware** ermöglicht und zugleich **alle Enterprise Funktionen** unterstützt, die Kunden von einer derartigen Lösung erwarten.





# SUSE Enterprise Storage - Architektur

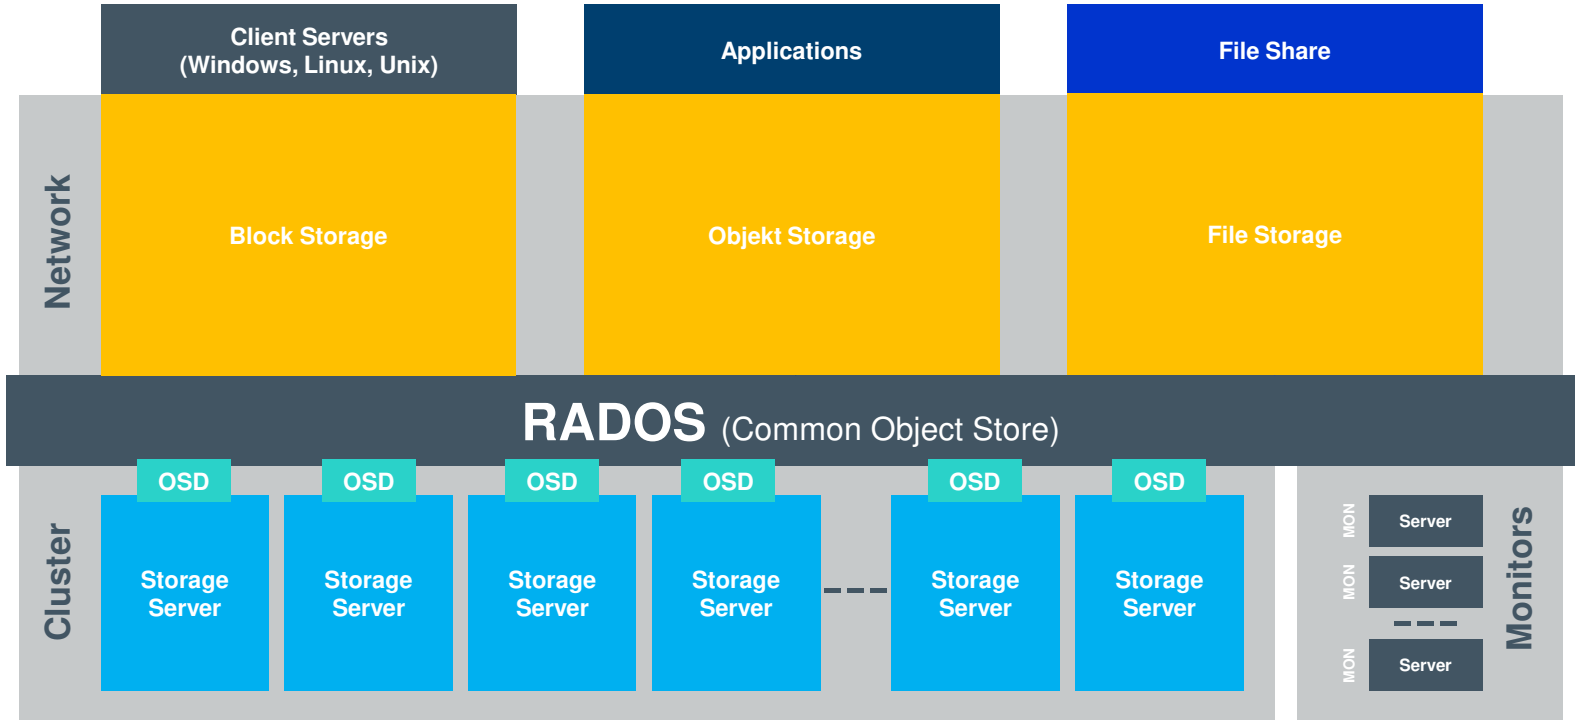

# Features und Roadmap

# SUSE Enterprise Storage 5

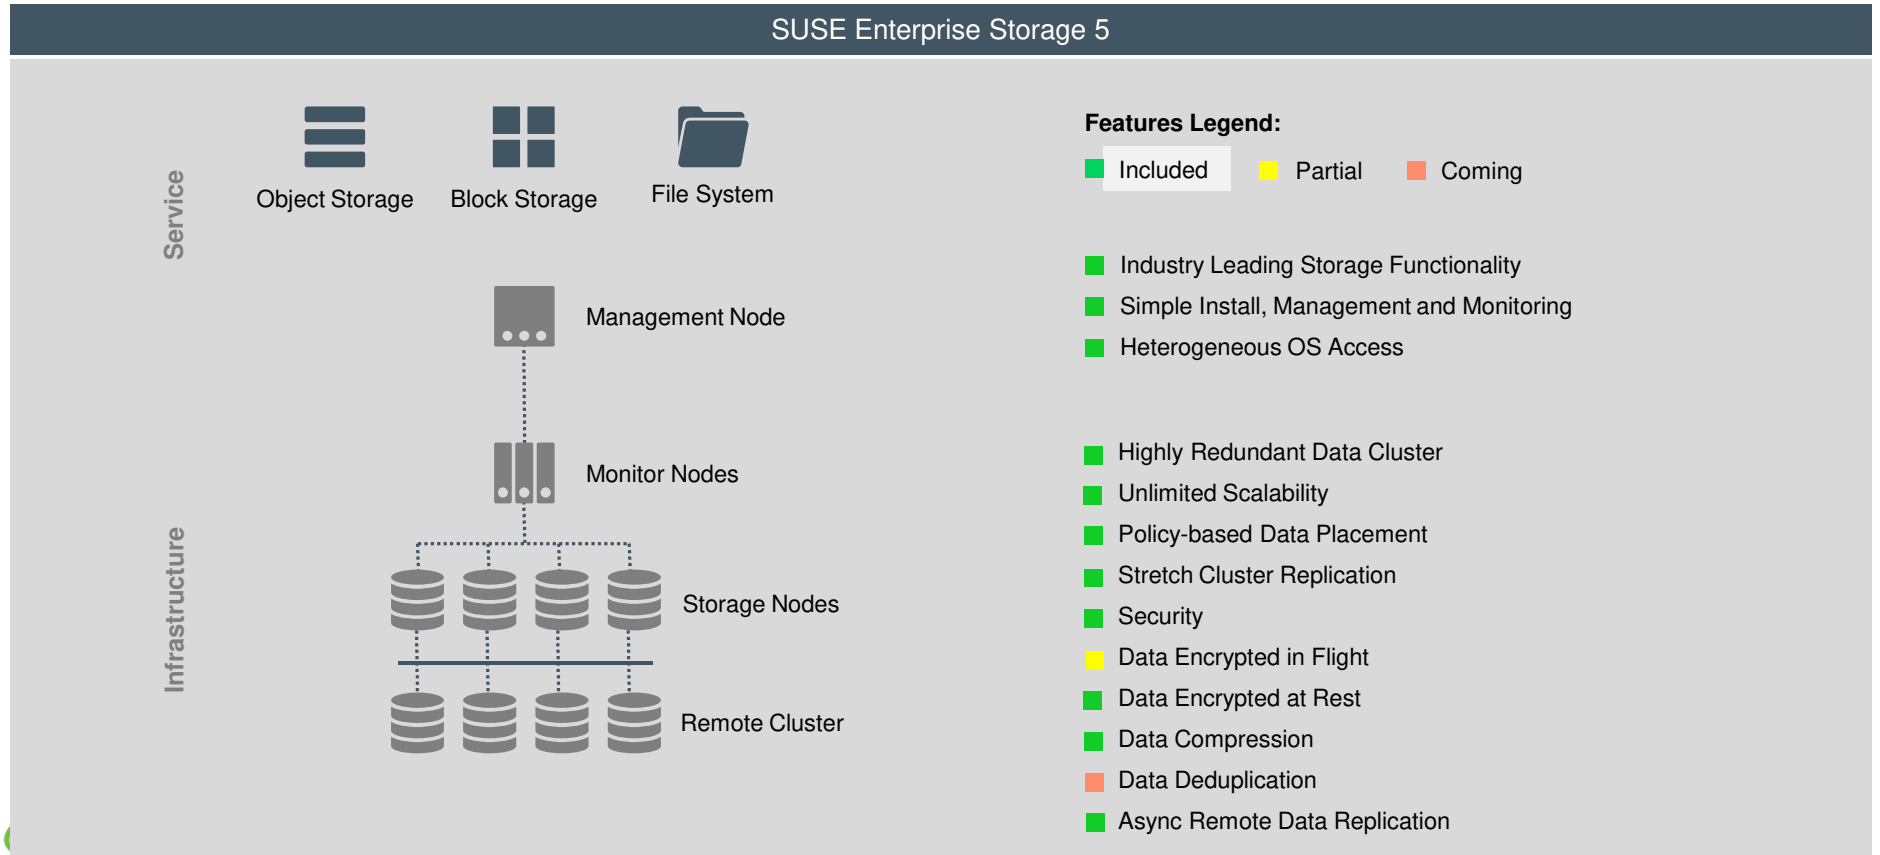

# SUSE Enterprise Storage 5 – Major Features

- Unified Block, Objekt und Dateien mit CephFS Dateisystem
- Erweiterte Hardware-Plattform mit Unterstützung für 64-Bit-ARM
- Asynchrone Replikation für Blockspeicher und multisite object replication
- Verbesserte Benutzerfreundlichkeit mit SUSE openATTIC
- Zugriff über NFS auf S3-Buckets
- „Pre Release“ Zugriff auf CIFS/SMB Shares über Samba

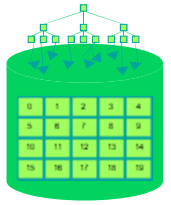

File Storage

Block Storage

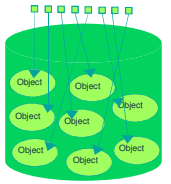


Object Storage

# SUSE Enterprise Storage 5

## SCHNELLER und EFFEKTIVER mit BlueStore Object Store

- Deutlich verbesserte Schreibperformance
- Data Compression
- Natives Block und File Erasure Coding

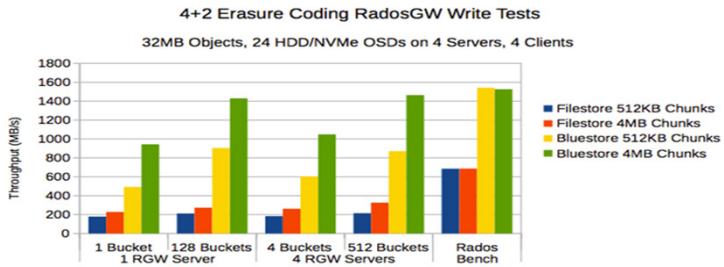


## Einfaches MANAGEMENT mit openATTIC Gen2 und DeepSea

- openATTIC Graphical User Interface für einfaches Storage Management
- Signifikante Verbesserung des openATTIC Device Monitoring
- Verbesserte Cluster Administration durch  erweiterteDeepSea Funktionen

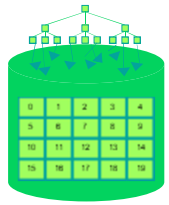

File Storage

Block Storage

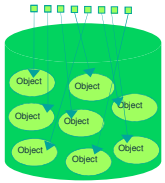


Object Storage

# SUSE Enterprise Storage 5

## Focus klar auf Performance

– SUSE Enterprise Storage 5 – Ceph BlueStore

- Bis zu 200% bessere Schreibperformance im Vergleich zum Vorgänger

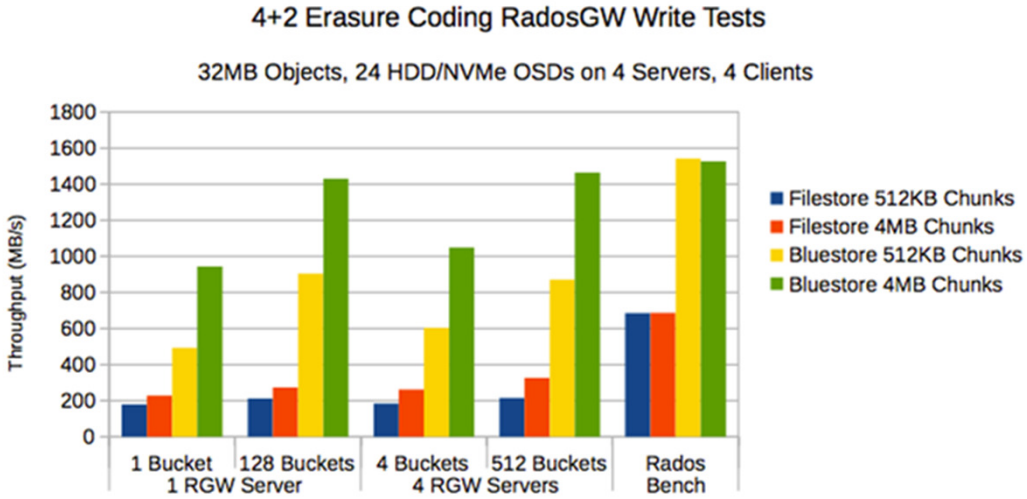


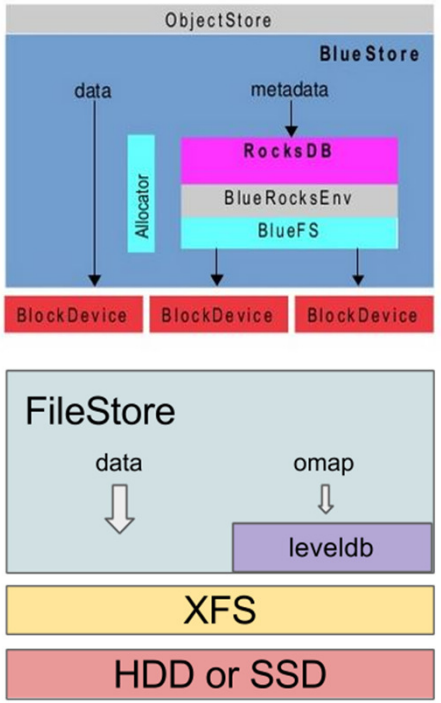

# SUSE Enterprise Storage – Neues Management

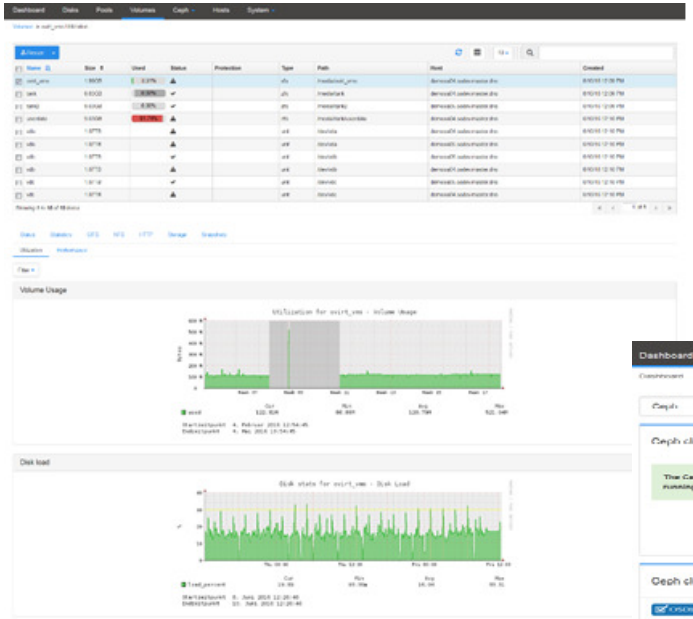

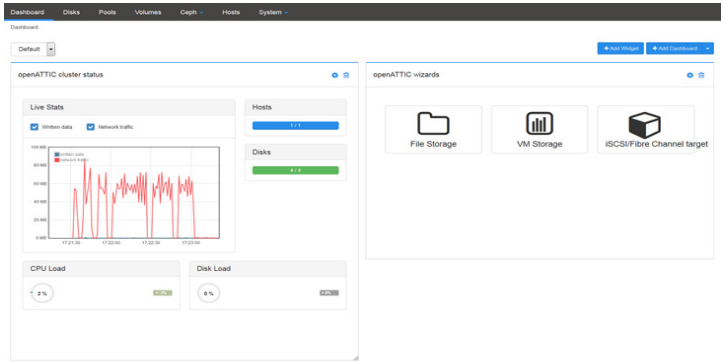


# SUSE Enterprise Storage Roadmap


## SUSE Enterprise Storage 4

**Built On**
- Ceph Jewel release
- SLES 12 SP2 (Server)

**Manageability**
- Initial openATTIC management
- Initial DeepSea Salt integration

**Interoperability**
- Arch64
- CephFS (production use cases)
- NFS Ganesha (Technology Preview)
- NFS access to S3 buckets (Technology Preview)

**Availability**
- Multisite object replication

## SUSE Enterprise Storage 5

**Built On**
- Ceph Luminous release
- SLES 12 SP3 (Server)

**Manageability**
- openATTIC management phase 2
  - Grafana monitoring dashboard
  - Prometheus event alert - email
- DeepSea Salt integration phase 2
  - Online Filestore to BlueStore

**Interoperability**
- NFS Ganesha
- NFS access to S3 buckets
- CIFS Samba (Technology Preview)
- CephFS Multi MDS support

**Availability**
- Erasure coded block and file

**Efficiency**
- BlueStore back-end
- Data compression

## SUSE Enterprise Storage 6

**Built On**
- Ceph Mimic release
- SLES 15 and CaaS Platform

**Manageability**
- openATTIC management phase 3
  - Event alert - SNMP traps
- DeepSea Salt integration phase 3
- Integration with SUSE Manager
- Automatic Metric Reporting phase 1
- IPv6

**Interoperability**
- Containerized SES
- CIFS/Samba
- Quality of Service (QoS)
- RDMA back-end (Technology Preview)

**Availability**
- Asynchronous iSCSI replication

## SUSE Enterprise Storage 7

**Built On**
- Ceph “P” release
- CaaS Platform (Server)

**Manageability**
- openATTIC management phase 4
- DeepSea Salt integration phase 4
- Integration with Kubernetes
- Automatic Metric Reporting phase 2
- Last good configuration rollback

**Interoperability**
- RDMA back-end

**Availability**
- CephFS snapshots
- Asynchronous file replication

Information is forward looking and subject to change at any time.

# Konfiguration

# Cluster Components: MON, OSD, MDS ???

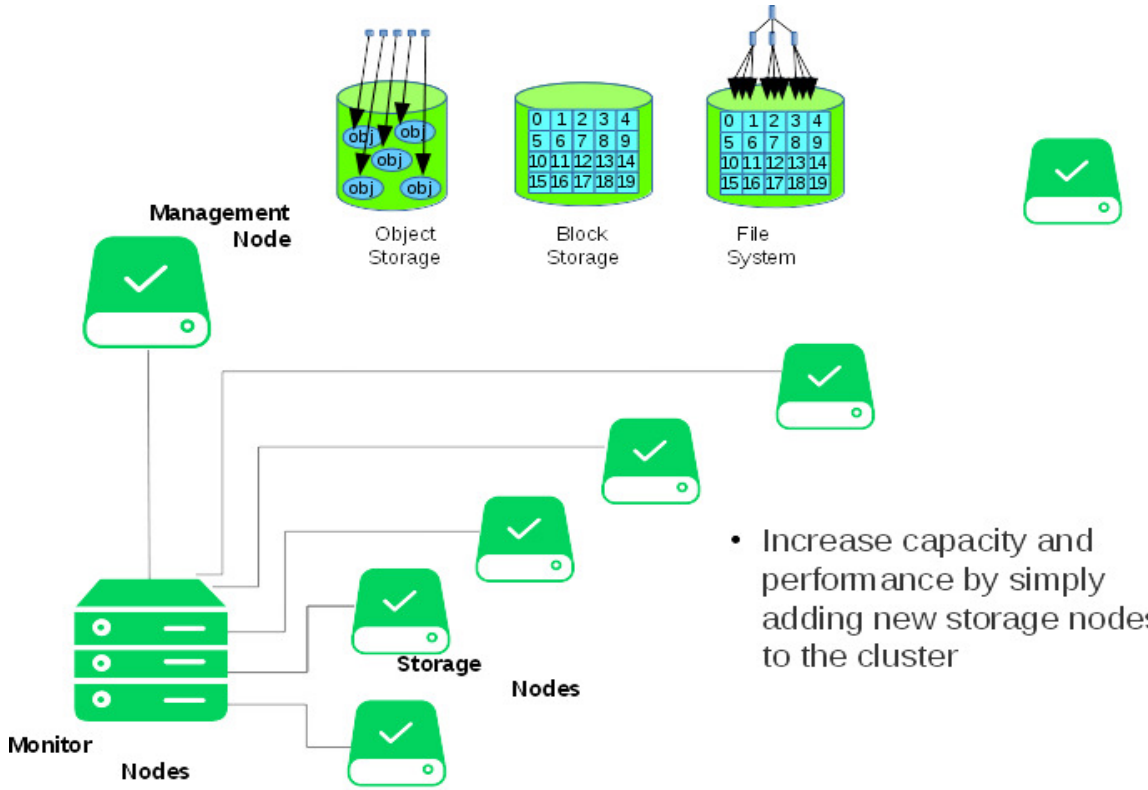


- Increase capacity and performance by simply adding new storage nodes to the cluster

SUSE

# S3 oder iSCSI

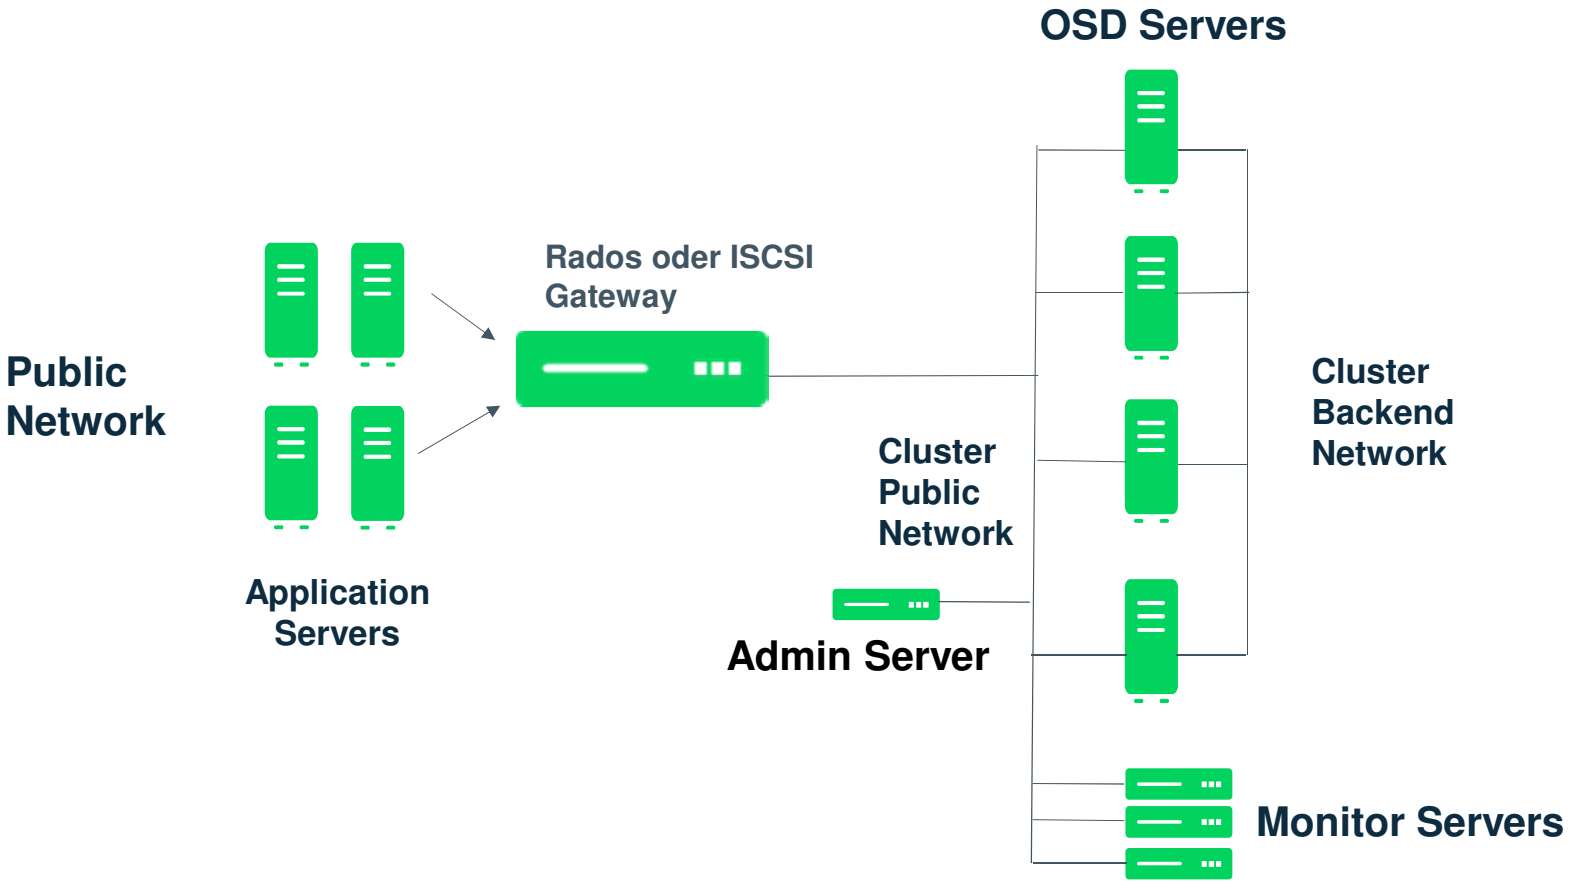

# File Transfer via CephFS

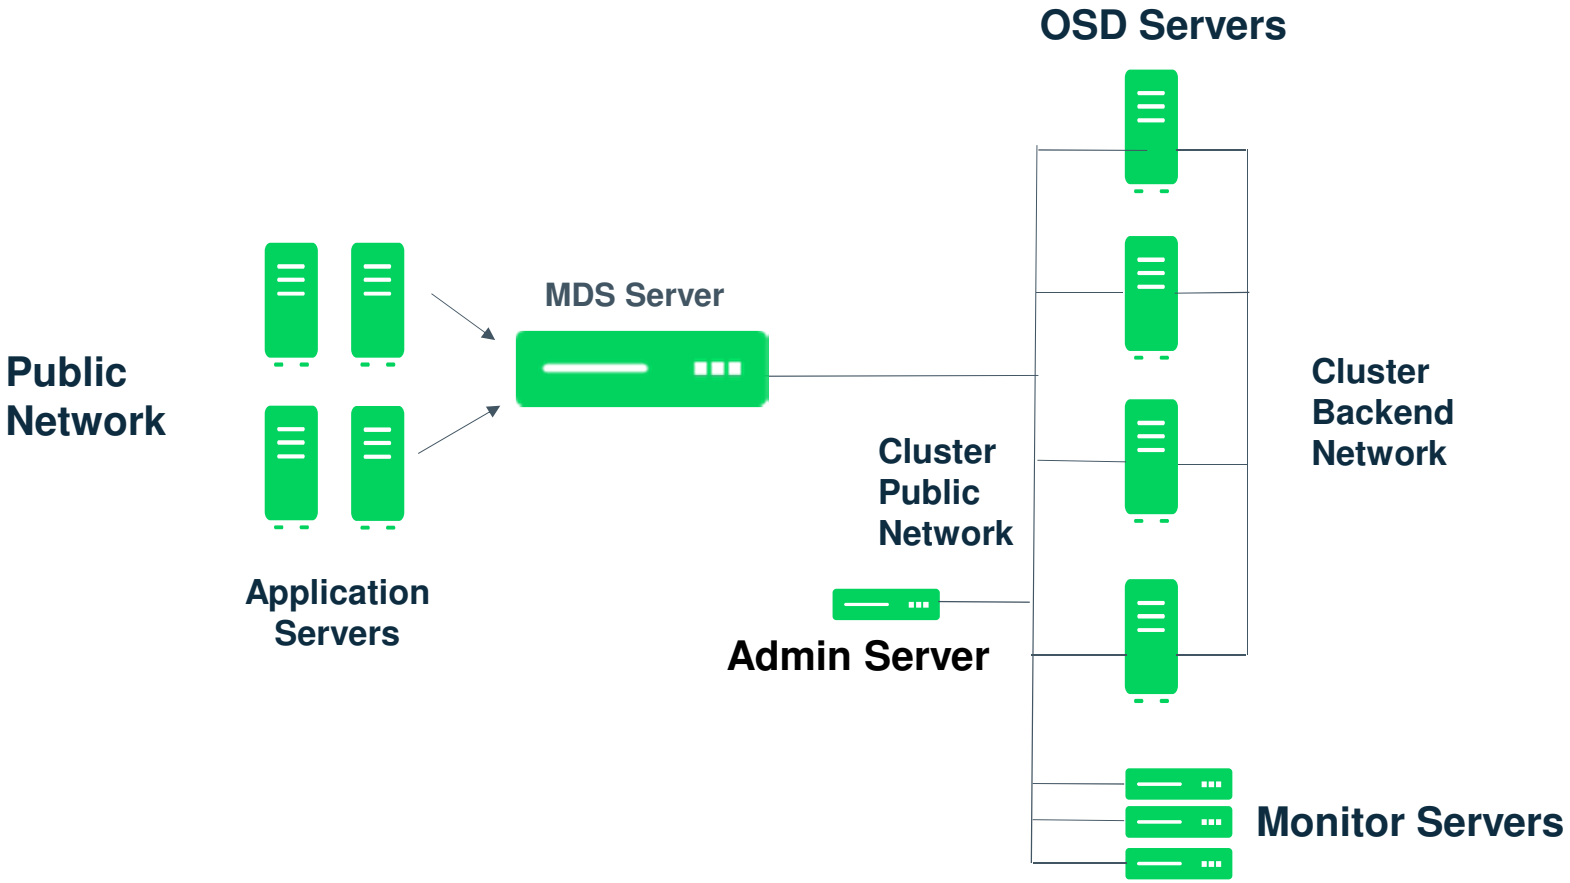

# RBD

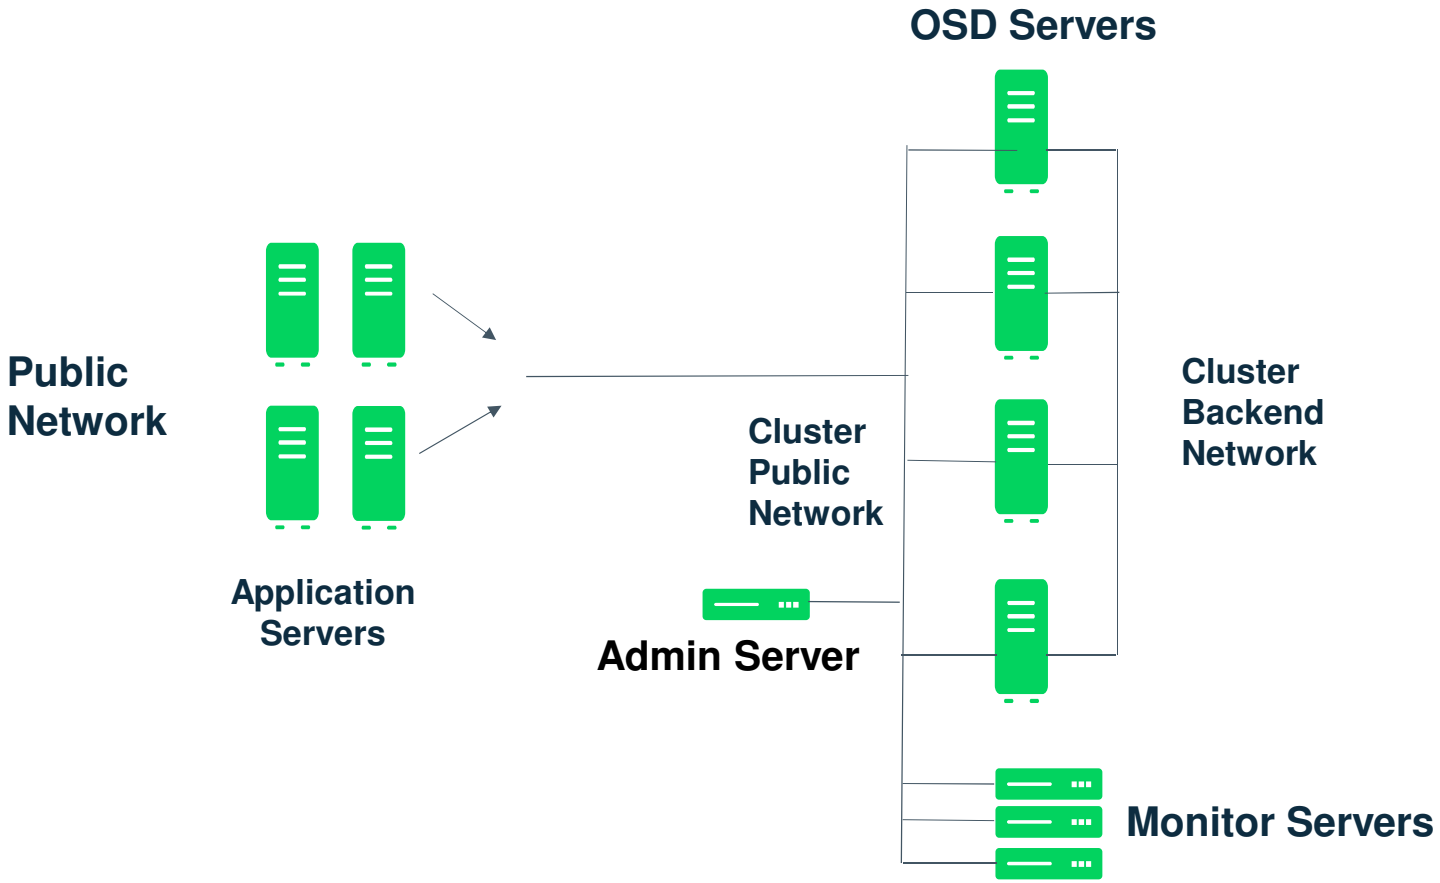
OSD Servers
Public Network
Application Servers
Cluster Public Network
Admin Server
Cluster Backend Network
Monitor Servers

# SMB / CIFS

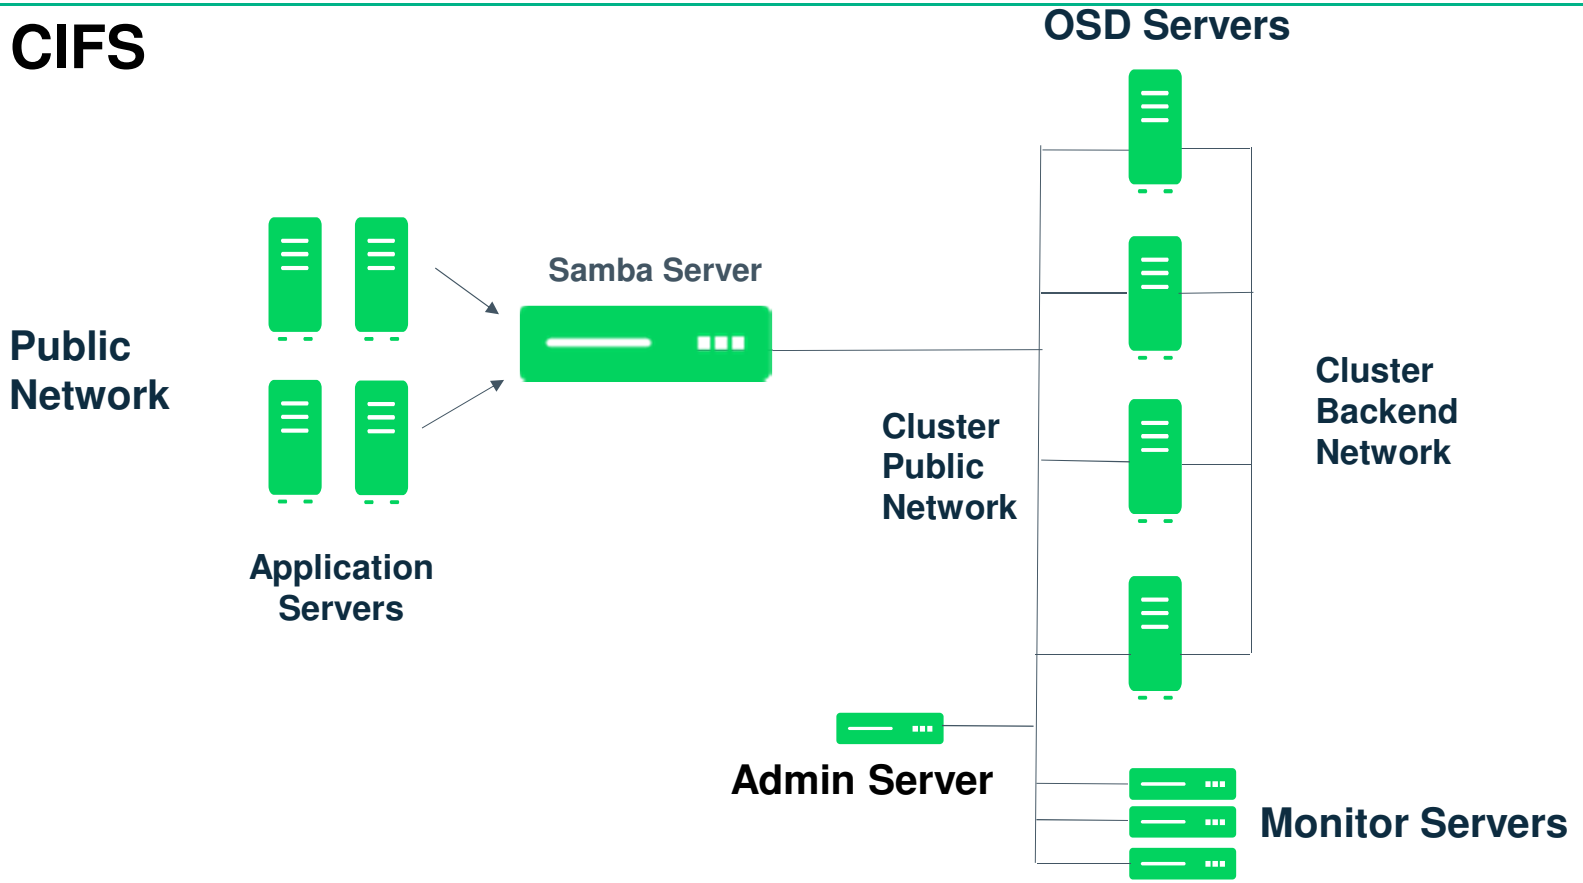


„For the Samba gateway (which will be tech preview in SES5) we'll offer AD domain membership via Winbind.

# SUSE Enterprise Storage Minimalkonfiguration für POC's


4 Object Storage Nodes + 1 Management Server

– 4 x SUSE Enterprise Storage Object Storage Nodes mit:

- Minimum 2 x 10 Gb Ethernet (1 x Storage Backbone, 1 x Client Network)
- Minimum 32 OSD HDDs im Cluster
- Dedizierte HDD für Betriebssystem
- Minimum 1 GB RAM pro TB Bruttokapazität pro Storage Node
- Minimum 1.5 GHz pro OSD pro Storage Node

– Separater Management Server

- Minimum 32 GB RAM, Minimum 4 Cores, 2 HDDs (SSD) für Betriebssystem

# SUSE Enterprise Storage Minimalkonfiguration für Produktion

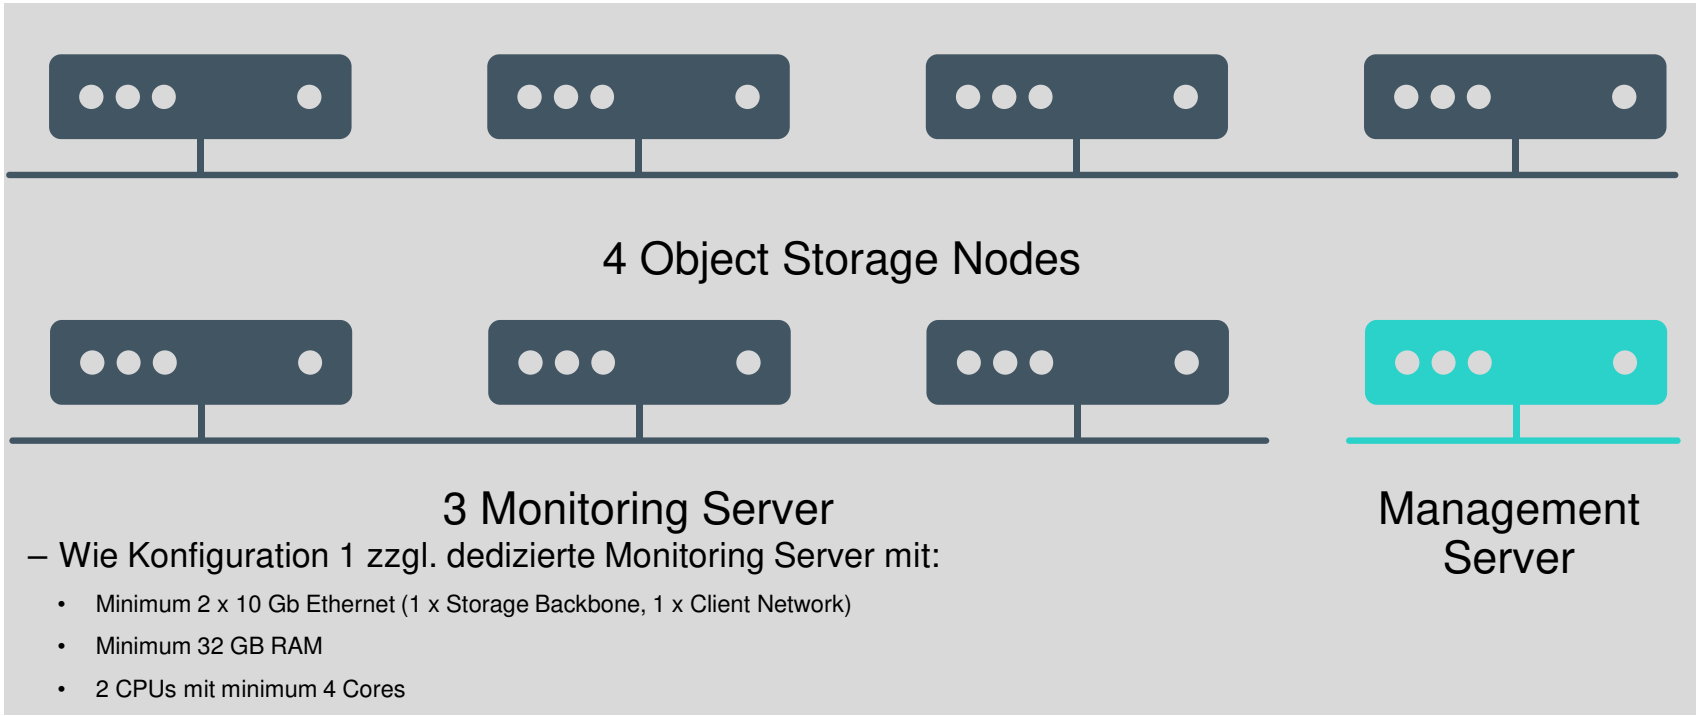


– Wie Konfiguration 1 zzgl. dedizierte Monitoring Server mit:

- Minimum 2 x 10 Gb Ethernet (1 x Storage Backbone, 1 x Client Network)
- Minimum 32 GB RAM
- 2 CPUs mit minimum 4 Cores
- Dedizierte HDD für Betriebssystem

# SUSE Enterprise Storage Produktivumgebung

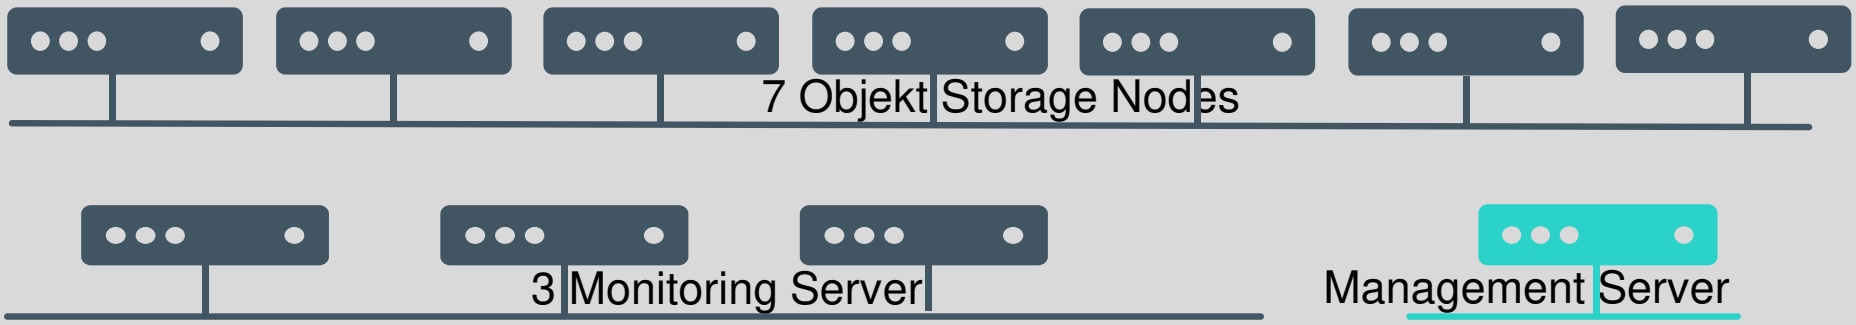


**7 SES Storage Nodes**

- 4 x 10 Gb Ethernet
- 56+ OSDs im Storage Cluster
- 2 SSD im RAID 1 für Betriebssystem
- SSDs für Journal (6:1 Ratio SSD Journal zu SATAs pro Node)
- 1.5 GB RAM pro TB Rohkapazität pro Storage Node
- 2 GHz pro OSD pro Storage Node

**Infrastruktur Nodes:**

- **Monitoring Nodes:**
  - Minimum 2 x 10 Gb Ethernet (1 x Storage Backbone, 1 x Client Network)
  - Minimum 32 GB RAM
  - 2 CPUs mit minimum 4 Cores
  - Dedizierte HDD für Betriebssystem
- **Management Node:**
  - Minimum 32 GB RAM, Minimum 4 Cores, 2 HDDs (SSD) für Betriebssystem
- **Gateway oder Metadaten Node:**
  - SES Object Gateway Nodes; 32 GB RAM, 8 core processor, RAID 1 SSDs for disk
  - SES iSCSI gateway nodes 16 GB RAM, 4 core processor, RAID 1 SSDs for disk
  - SES metadata server nodes (one active/one hot standby); 32 GB RAM, 8 core processor, RAID 1 SSDs for disk

https://www.suse.com/documentation/ses4/book_storage_admin/data/cha_ceph_sysreq.html

# SUSE Enterprise Storage Preise

– Basis Konfiguration

– SUSE Enterprise Storage und “limited use” von SUSE Linux Enterprise Server:
  – 4 Storage OSD Nodes (1-2 sockets) und
  – 6 Infrastruktur Nodes

– Expansion Node

– SUSE Enterprise Storage und “limited use” von SUSE Linux Enterprise:
  – 1 SES Storage OSD Node (1-2 sockets) oder
  – 1 SES Infrastruktur Node

# SUSE Enterprise Storage

## Pricing

### Base Configuration

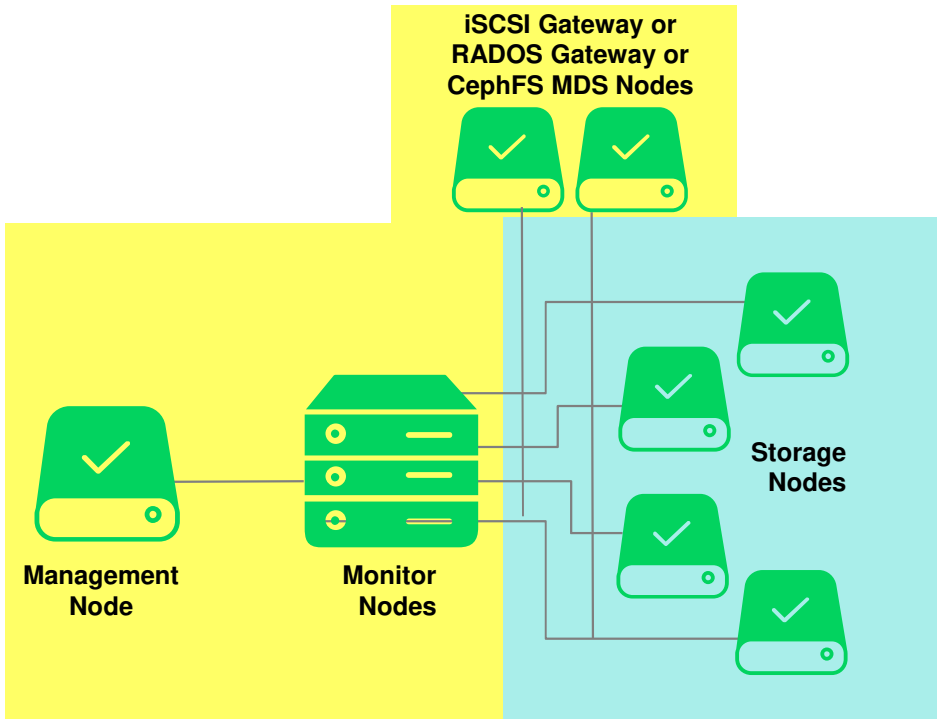


4 Storage Nodes
AND
6 Infrastructure Nodes

### Expansion Node

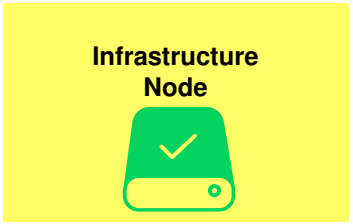


AND

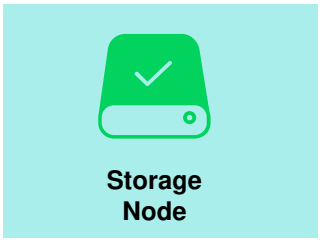


1 Storage Node
AND
1 Infrastructure Node

# Anwendungsfälle

# Anwendungsfälle

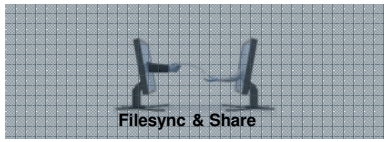




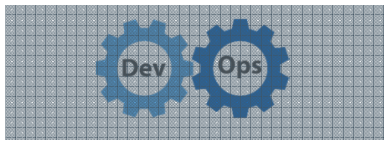


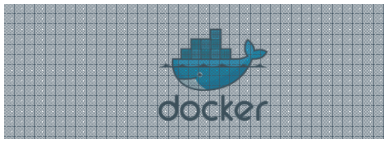




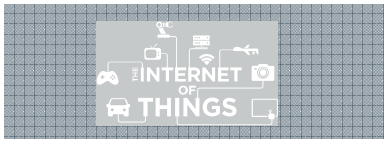


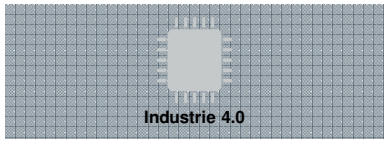



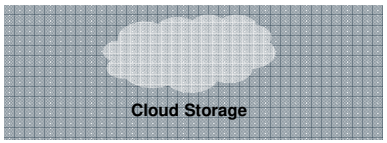



SUSE

# Anwendungsbeispiel - Content Store

- Forschungsabteilungen
  - Meteorologische Daten, Weltraumüberwachung, Satellitendaten
  - IoT Sensordaten, Autonomes Fahren
  - HPC Daten
- Medienindustrie
  - TV, Radio
  - Webdienste (YouTube etc.)
  - Streaming Industrie (Netflix, Sky etc.)

SUSE

# Anwendungsbeispiel - Backup

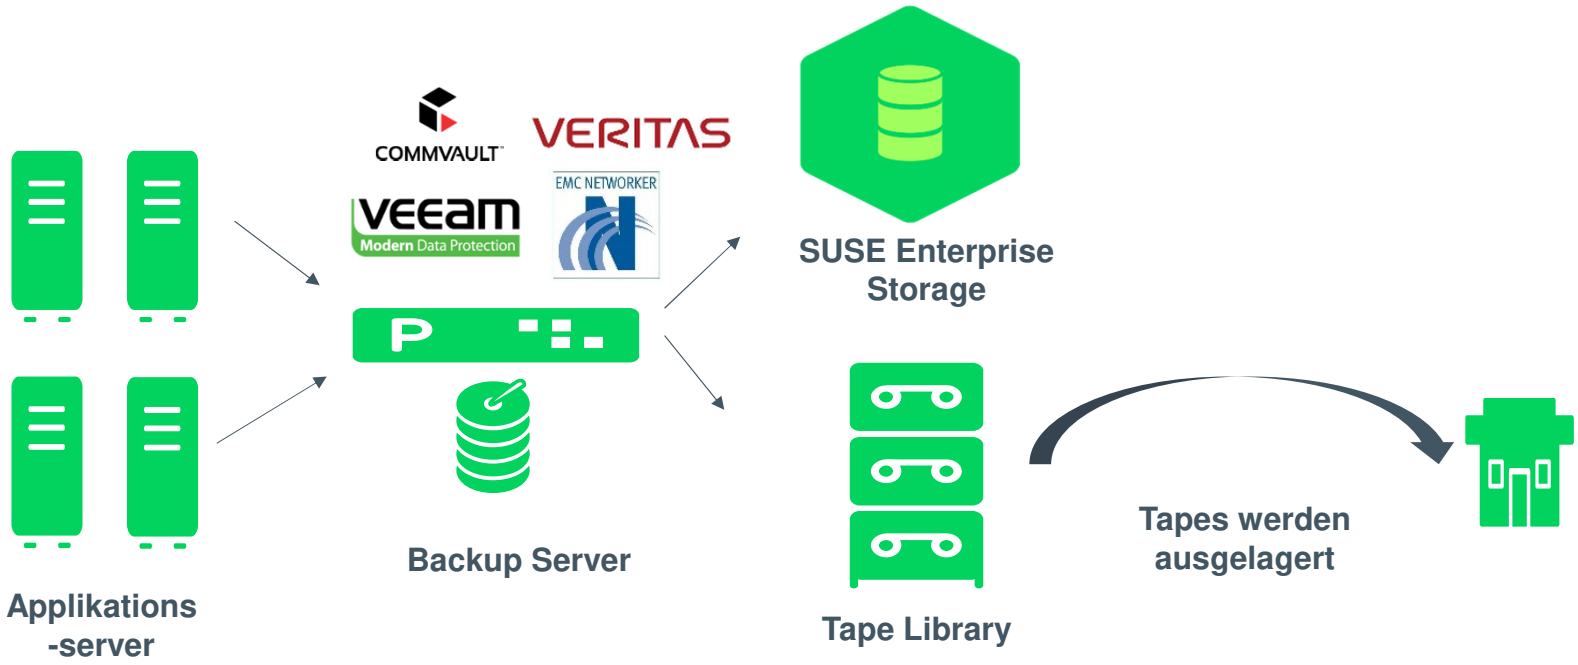

# Anwendungsbeispiel - Videoüberwachung

– Industriegebäudeüberwachung
– Verkehrsüberwachung
– Bodycams der Polizei
– Öffentliche Bereiche (Flughafen, Bahnhof)

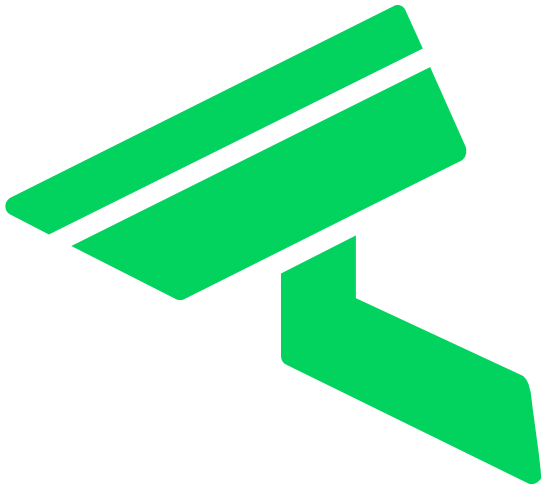

# Anwendungsbeispiel – Objekt / Massenspeicher

– Daten, die konstant zunehmen
– Archivierung
– Finanzdaten, Dokumente,
– Medizinische Daten
– Big Data

# Anwendungsbeispiel - Virtual Machine Speicher

- OpenStack (80% Marktanteil)
- Test-Dev VMs, “normale VMs”
  - Support via:
    - kvm – natives Blockdevice
    - Hyper-V – iSCSI
    - VMware - iSCSI

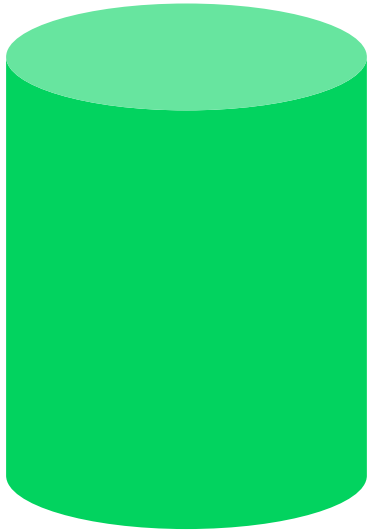

SUSE

# Zusammenfassung

# SUSE Enterprise Storage - Vorteile

– Multiprotokoll:

**Objekt: S3/SWIFT** **Block: Linux/MSFT/Vmware** **File: CephFS**

– Kosteneffizient:

**Je grösser desto günstiger**

– Hochskalierbar und flexibel:

**Unendlich skalierbar bei linearen Kosten**

– Selbstverwaltend:

**Mehr Speicher bei gleicher Anzahl der Administratoren**

– Hochverfügbar:

**Synchrone Spiegel, ansychrone Spiegel, Georedundanz**

– Heterogene Betriebssysteme:

**Heterogener OS-Support**

– Sicher:

**Datensicherheit gewährleistet**

# Die Realität in der IT Abteilung:

- Niedrigere Kosten im Vergleich zu traditionellen Anbietern ✓
- Unified Architektur ✓
- Erweiterung nach Businessanforderung ✓
- Unendlich skalierbar ✓
- Kleine Einstiegskonfiguration möglich ✓
- Basiert auf Standard Serverhardware – kein Lock-In ✓
- Cloud-fähig mit Amazon S3 Schnittstelle ✓

SUSE

# Wohin entwickelt sich der Markt?

Enterprise Storage Revenue and Milestones

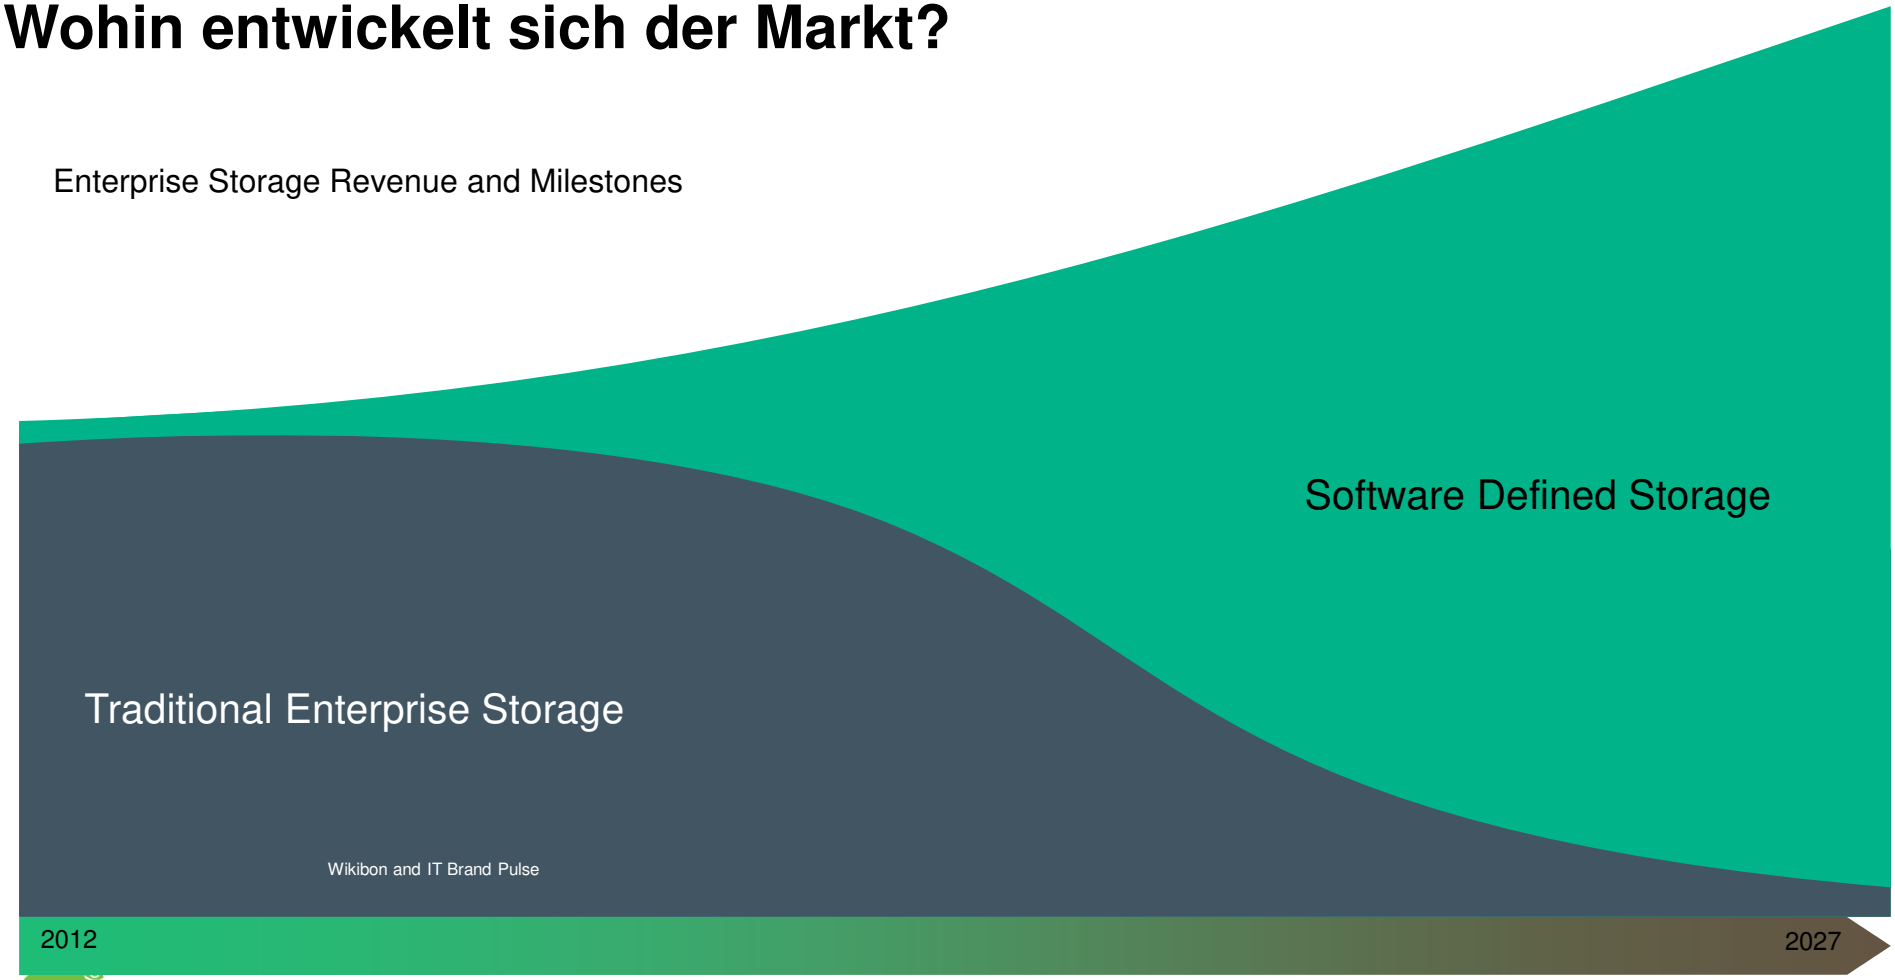

# Alles wird Software-Defined…

## Bleiben Sie am Ball und setzen Sie auf SUSE


Quelle: http://www.runnersworld.de/laufevents/haile-gebrselassie-verzichtet-auf-wm-start-in-berlin.115667.htm

## und nicht auf das falsche Pferd...

Quelle: http://www.holiday-hungary.de/puszta/ungarische-puszta.htm

# Installation

# SUSE Enterprise Storage Deployment

**Stage 0**
The **preparation**— during this stage, all required updates are applied and your system may be rebooted.

**Stage 1**
The **discovery**— here you detect all hardware in your cluster and collect necessary information for the Ceph configuration

**Stage 2**
The **configuration**— you need to prepare configuration data in a particular format

**Stage 3**
The **deployment**— creates a basic Ceph cluster with mandatory Ceph services

**Stage 4**
The **services**— additional features of Ceph like iSCSI, RADOS Gateway and CephFS can be installed in this stage. Each is optional

**Stage 5**
**T**he removal stage. This stage is not mandatory and during the initial setup it is usually not needed

# Thank You