SCALITY

PRESENTATION

November, 22nd 2019

WHERE SCALITY FITS IN THE ENTERPRISE IT LANDSCAPE

Scality is a specialist in storing, protecting & managing massive amounts of unstructured data both on-premises and in public clouds

Unstructured data (videos, images, documents, backup, logs, IoT) is the fastest growing part of data overall - IDC estimates we will see 160 Zettabytes of this data by 2025 - this includes file data created by large enterprises

Our Enterprise customers are global banks, insurance companies, travel & reservation providers with storage requirements for one or several unstructured data use-cases

Scality RING and ZENKO are optimal solutions for enterprise private cloud storage, medium-to-long term backups, NAS Archives and to support Innovation for Enterprise Digital Transformation

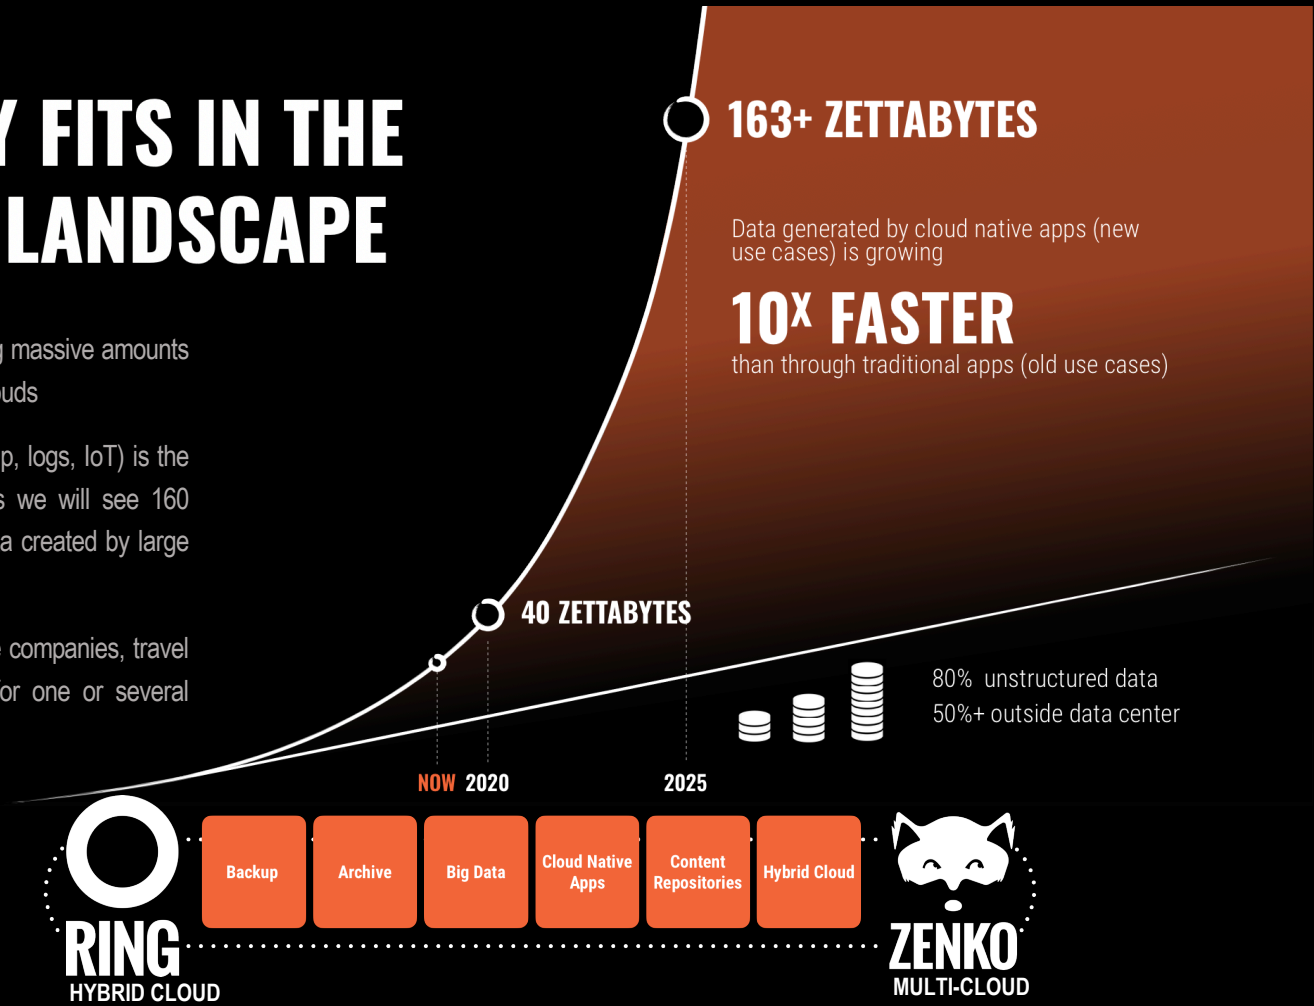

SCALITY

What is Object/Unstructured Storage

The Car Park Illustration

Block file system – defined data sizes – complex management – inflexible upgrades

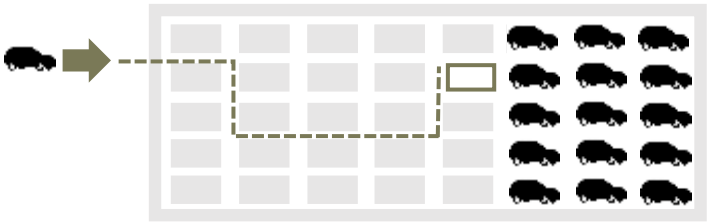

Traditional Appliance

Modern object store system – Petabyte scale – simplified management – incremental, unlimited upgrades – Amazon S3

SCALITY

What is an object?

Data + Metadata + Unique Identifier

Flat Namespace

Object Storage Explained

- **Object storage stores data into containers, called objects**
- **Each object has both data and user defined and system defined metadata (a set of attributes describing the object)**

File Systems

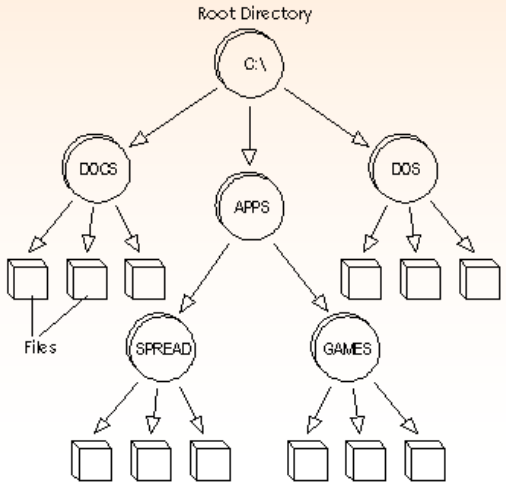

File Systems were designed to run individual computers, then limited shared concurrent access, not to store billions of files globally

Objects

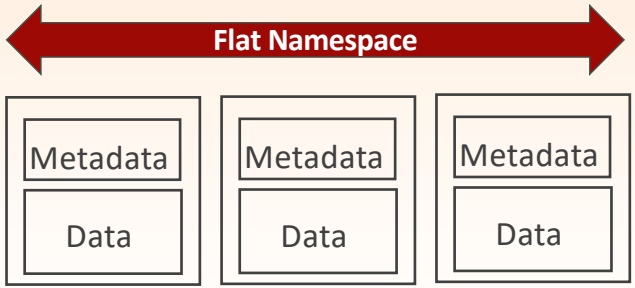

Objects are stored in an infinitely large flat address space that can contain billions of files without file system complexity

SCALITY

GLOBAL PRESENCE

Founded in 2009 · 8 global offices · 20+ nationalities

STRATEGIC ALLIANCES

Hewlett Packard Enterprise

CISCO Partner
Solution Partner

GLOBAL CLIENT BASE

5 of the top 15 banks & 7 of the top 20 Service Providers

EUROPE

amadeus · Ministère de la Défense · orange · Mairie de Paris · Société Générale · RTL · SFR · dailymotion · TF1 · Eurosport · Hospices Civils de Lyon · Renault · SNCF · AB · EDF · Natixis

AMERICAS

Bloomberg Media · rackspace · Penn Medicine Lancaster General Health · Charter Communications · Bank of America · deluxe · HBO · Comcast

AUSTRALIA

SMX Search Marketing Expo · foxtel · T · EndemolShine Group · MediaHub

JAPAN

KDDI · BIGLOBE · DMM.com · @nifty · Nagoya University

SCALITY LEADS

2019 GARTNER MAGIC QUADRANT

FOR DISTRIBUTED FILE AND OBJECT STORAGE

For the fourth consecutive year, Scality is recognized as a Leader

2019 GARTNER CRITICAL CAPABILITIES

FOR OBJECT STORAGE

Scality RING receives highest score for hybrid cloud storage use case | first or second score across all five use cases

2018 IDC MARKETSCAPE WORLDWIDE

OBJECT-BASED STORAGE

IDC MarketScape: Worldwide Object-Based Storage 2018 Vendor Assessment places Scality in the leader category once again

SCALITY

RING IS FOR MASSIVE SCALE UNSTRUCTURED DATA

RING is for customers who have 100's of Terabytes to 100's of Petabytes of unstructured data requiring highly-reliable, cost-effective storage

Backups, images, video, medical images & scans, documents, emails, maps, data lake (log files, IoT sensor data) and more

Unstructured data can be stored as files or object data, depending on the application

RING software is hosted on industry leading servers

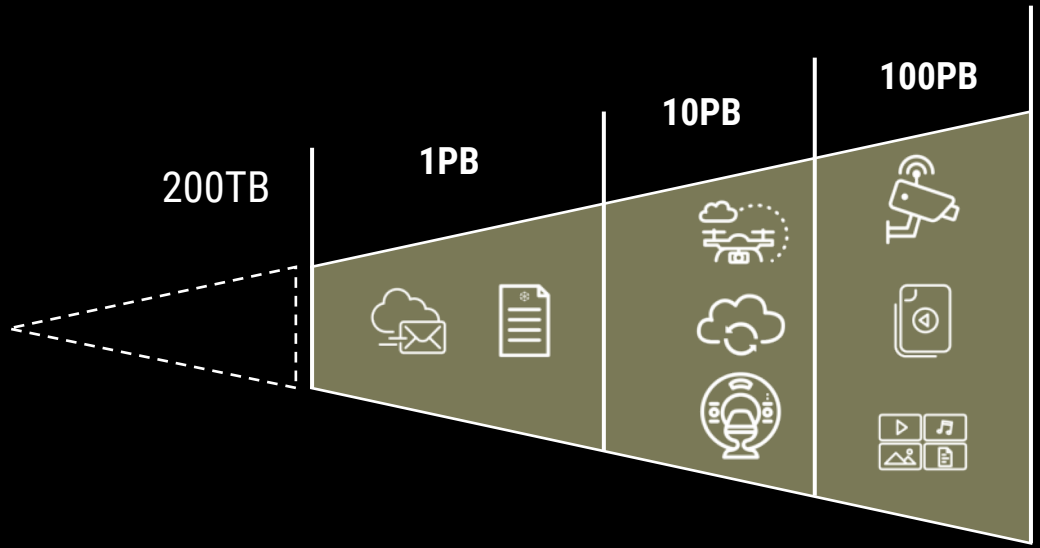

SCALITY

Scality Use Cases

Private/Hybrid Cloud

Backup & Archive

Content Distribution & Nearline Archive

Medical Imaging Archives

Public Cloud Consumer Services & Email

Video Surveillance

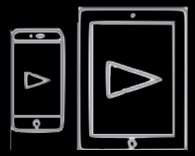

Enterprise File Sync & Share

Hadoop Datalake

Deep Learning

READY-TO-DEPLOY ISV PARTNERS

Backup and Archive

Medical Imaging

Big Data

Content Distribution

Cloud Native App

CCTV

https://www.scality.com/isv-compatibility-list/

SCALITY RING SOFTWARE

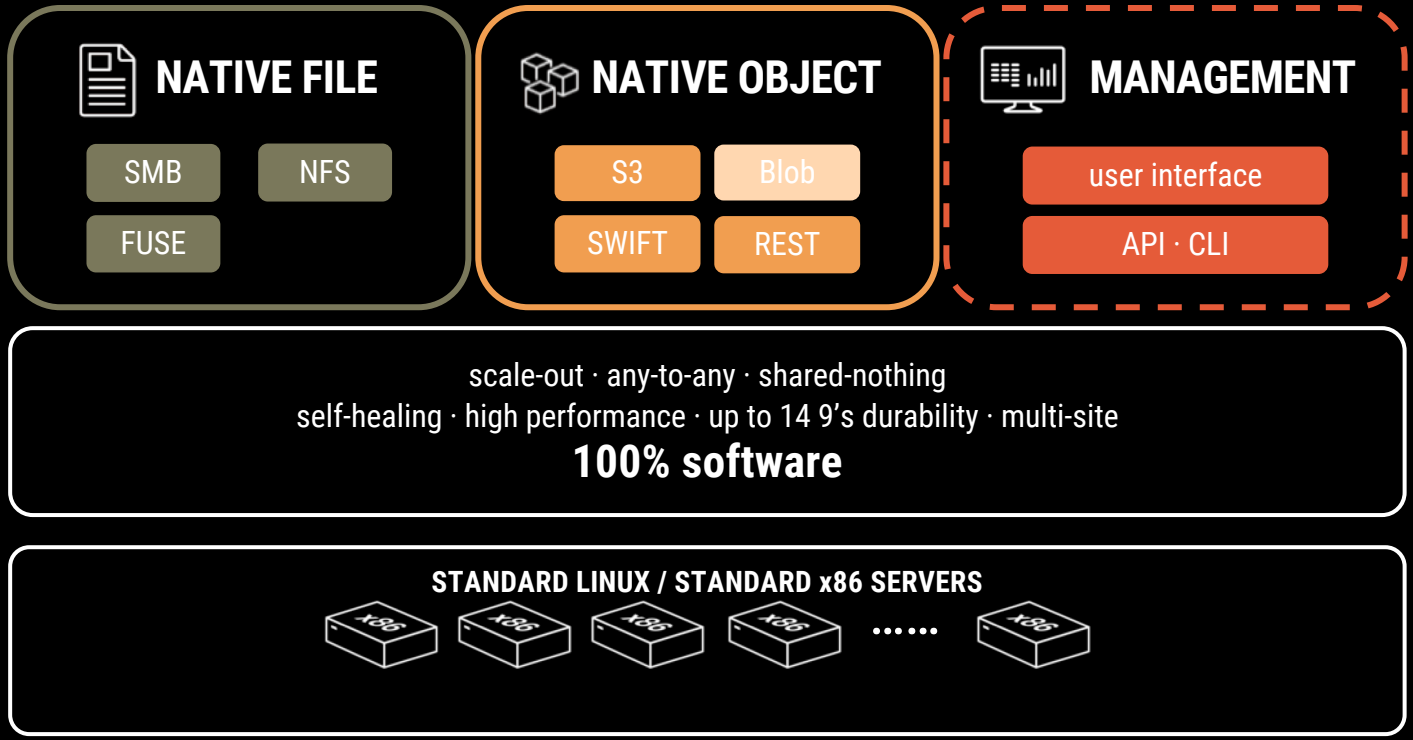

RING DATA PROTECTION 3 SITES

REPLICATION ERASURE CODING MULTI-SITE

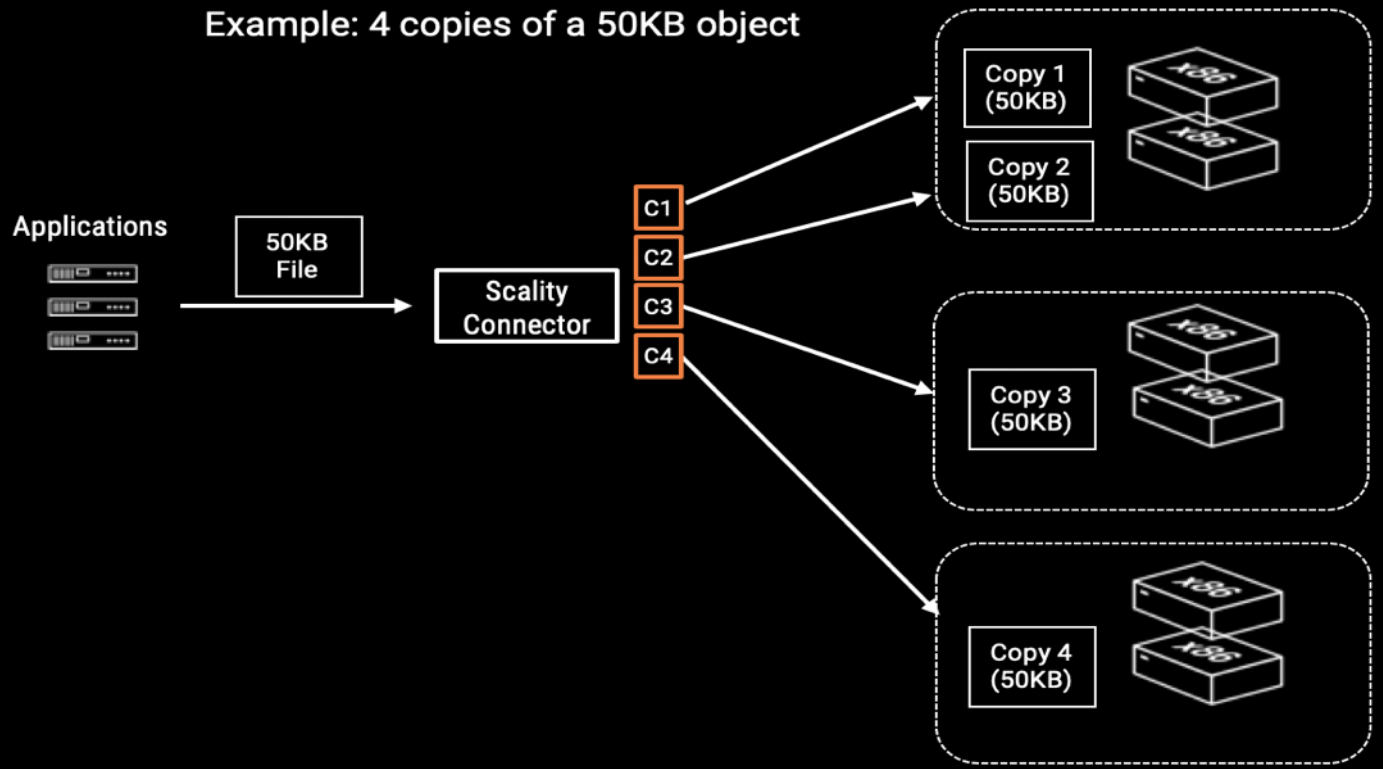

SCALITY

RING DATA PROTECTION 3 SITES

REPLICATION **ERASURE CODING** MULTI-SITE

Example: EC (7,5) on a 7MB object
Provides three-disk failure protection with ~71% overhead

SCALITY

RING DATA PROTECTION

REPLICATION ERASURE CODING **MULTI-SITE**

Single Site RING

Mirrored RINGs

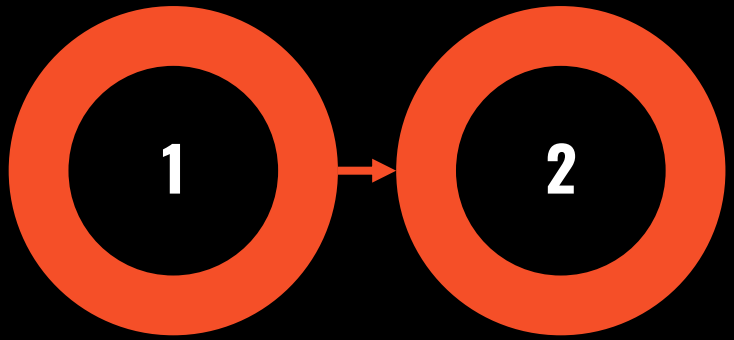

Stretched RING

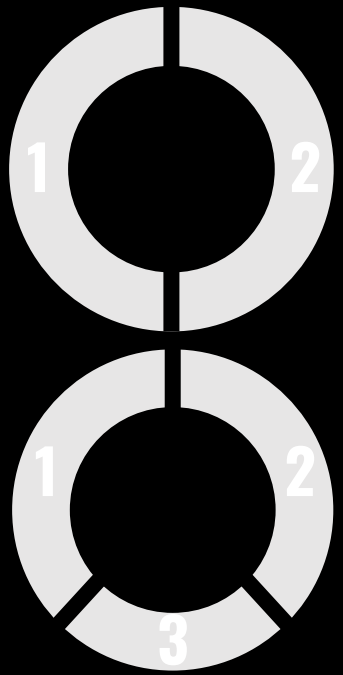

SCALITY

RING SCALE OUT FILE SYSTEM MULTI-SITE

MIRRORING | 2-SITE WITH WITNESS | 3-SITE STRETCHED

asynchronous replication of data from site 1 to site 2 for DR
configurable per volume · RPO = minutes · delta-only replication
read-only access on DR site

SCALITY

RING SCALE OUT FILE SYSTEM MULTI-SITE

MIRRORING | 2-SITE WITH WITNESS | 3-SITE STRETCHED

synchronous operations across 2 full sites + a quorum/witness site
any volume *belongs* to one site only · any site can *host* volumes
1 site failure tolerance with RPO = 0
better durability and storage efficiency than 2-site mirroring

SCALITY

RING SCALE OUT FILE SYSTEM MULTI-SITE

MIRRORING 2-SITE WITH WITNESS 3-SITE STRETCHED

synchronous operations across 3 sites · any 1 site failure tolerance · any 1 volume *belong* to one site · any site can *host* volumes · RPO = 0
best durability and storage efficiency combination of all multi-geo models

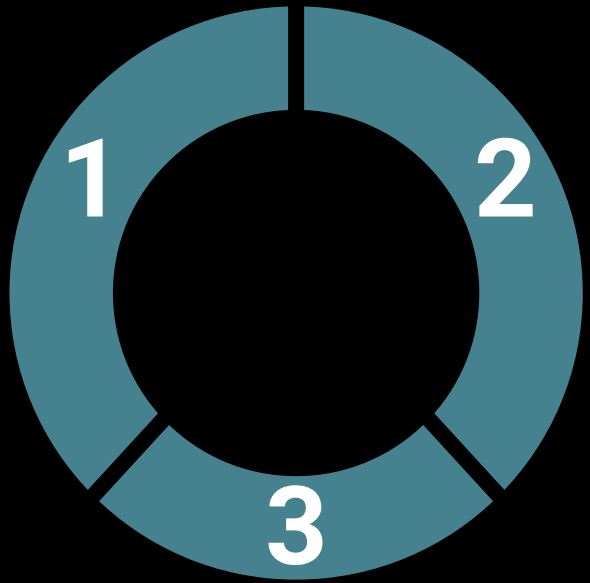

SCALITY

RING SCALE OUT FILE SYSTEM

SMBv3, NFSv4, Linux FUSE access
IPv6 support for NFS & SMB with native encryption
unlimited amount of volumes and files
distributed POSIX metadata · stateless connectors
quotas · file versioning · WORM

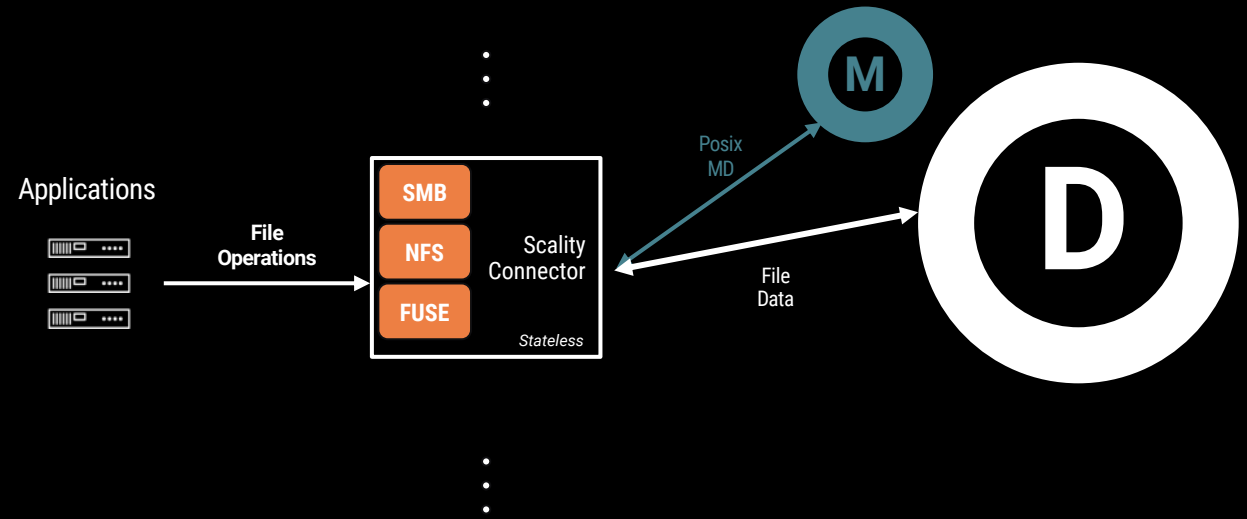

RING SCALE OUT FILE SYSTEM FEATURES

FILE VERSIONING WORM FILE DATA THROUGH S3 API

create, delete and truncate = version
kept in *recycle* directory · per minute granularity · customizable ACLs
configurable version retention policy
configured per volume

RING SCALE OUT FILE SYSTEM FEATURES

FILE VERSIONING WORM FILE DATA THROUGH S3 API

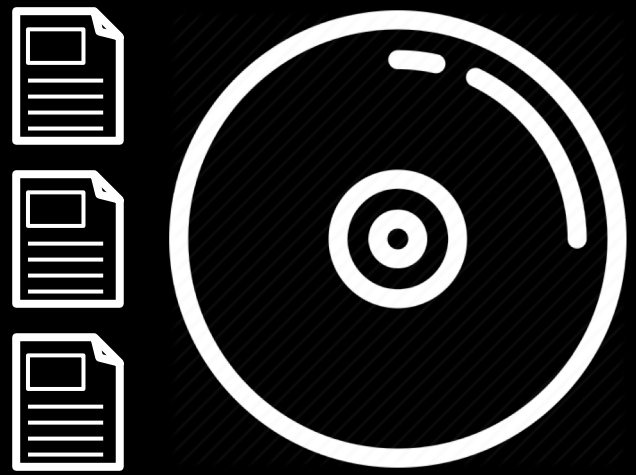

data becomes read-only automatically after a set period of time
time starts counting at file creation · configurable trigger period
WORM compliance · protects against identified threat like ransomware
configurable per volume

RING SCALE OUT FILE SYSTEM FEATURES

FILE VERSIONING WORM FILE DATA THROUGH S3 API

enables S3 read-only access to a scality scale out file system volume · great when switching applications from file to S3 · helps with migration and maintaining access to legacy file data

RING S3 CONNECTOR

AWS S3 and IAM API compatible
unlimited amount of buckets and objects
distributed scale-out metadata · stateless connectors
bucket encryption · object versioning · WORM · utilization API ·
bucket location control · lifecycle management

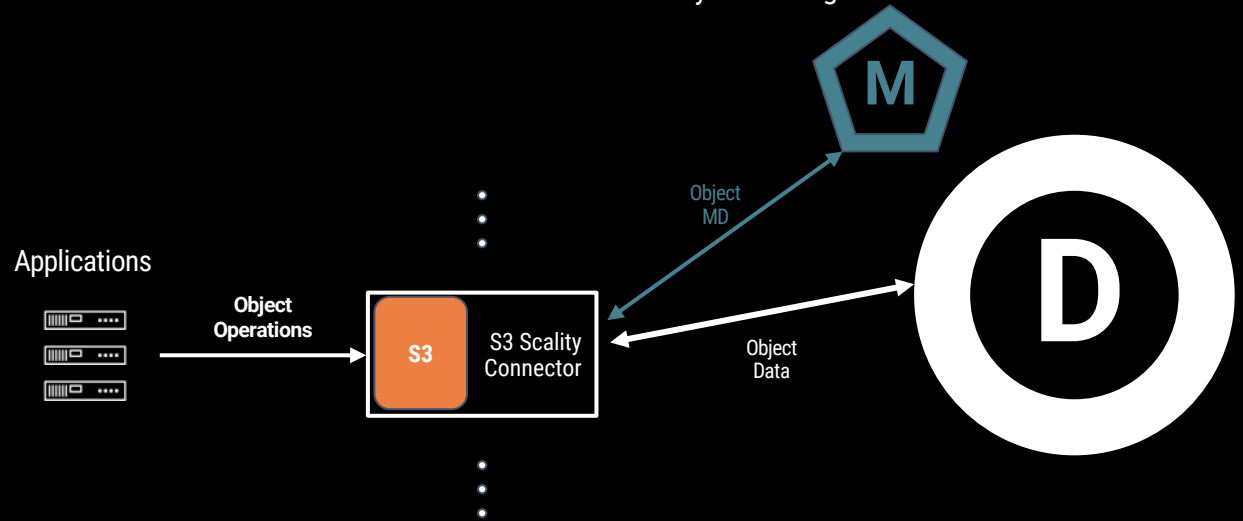

SCALITY

RING S3 CONNECTOR

S3/IAM API support Utilization API

Bucket Encryption Versioning & WORM Bucket Location Control Bucket Lifecycle Management

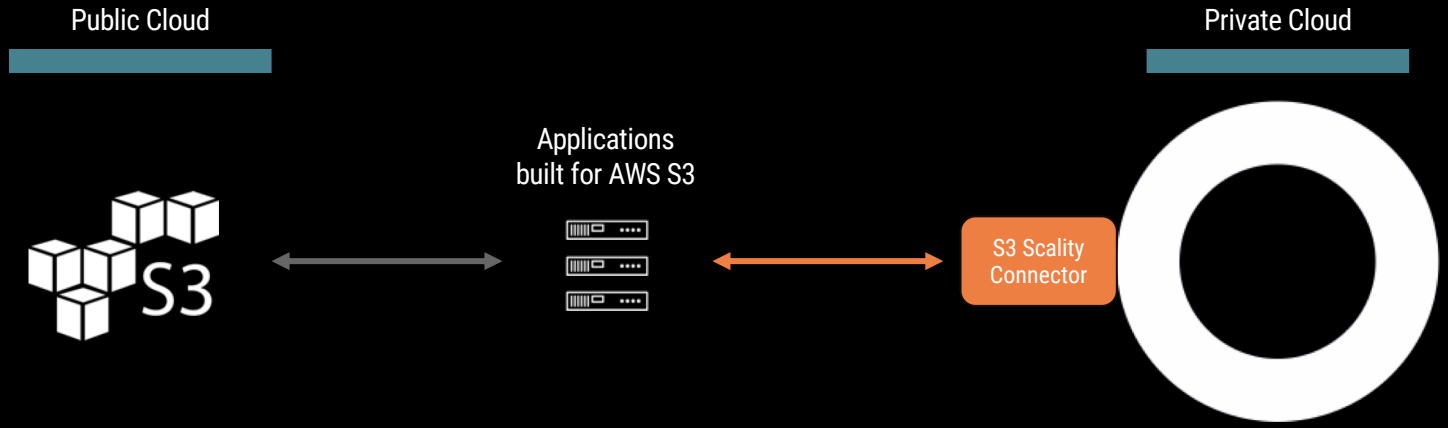

plug and play S3 endpoint replacement for your existing AWS S3-built applications
comprehensive S3 storage API support
AWS identity and access management (IAM) support

RING S3 CONNECTOR

S3/IAM API support Utilization API

Bucket Encryption Versioning & WORM Bucket Location Control Bucket Lifecycle Management

bucket and account level metrics
storage capacity and number of objects · bandwidth (bytes transferred) & number of S3 operations per unit of time

accessible via API calls, or through the scality UI · integrated with its own IAM policies for access control

RING S3 CONNECTOR

Docker based S3 Metadata Vault S3 Utilization API

Bucket Encryption Versioning & WORM Bucket Location Control Bucket Lifecycle Management

Bucket-Level Encryption

- Special header used at Bucket creation
- Extension to the Amazon S3 specification
  - All objects in bucket are encrypted, no special calls necessary
  - AWS SSE encrypts at the object level, through special calls from the application

Requires external Key Management Service (KMS)

- The KMS is responsible for managing encryption keys based on a passed Bucket ID
- KMS is called for PUT (encrypt) and GET (decrypt) object operations
- Key Management System (KMS) integration via KMIP 1.2

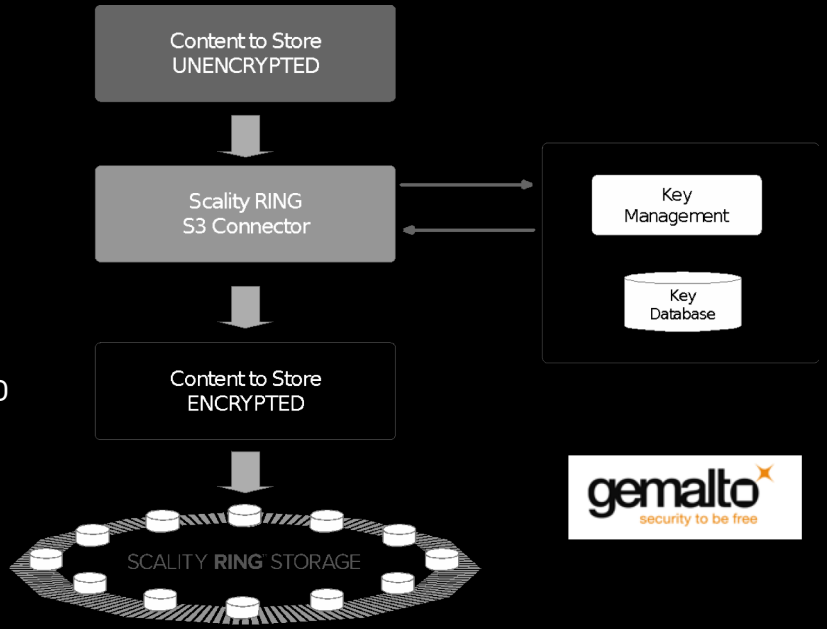

RING S3 CONNECTOR

Docker based S3 Metadata Vault S3 Utilization API

Bucket Encryption Versioning & WORM Bucket Location Control Bucket Lifecycle Management

bucket-level
updates (PUT) or deletes create a new version
reads (GET) by default get the latest version of an object · reads can reference an older version
can't be disabled, but can be suspended
specific versions can be deleted · lifecycle expiration can automatically delete old versions of objects

WORM can be enabled via policies (no deletes)

RING S3 CONNECTOR

Docker based S3 Metadata Vault S3 Utilization API

Bucket Encryption Versioning WORM **Bucket Location Control** Bucket Lifecycle Management

Bucket Location Control

- Location Constraint at the Bucket level
- When creating a Bucket, an option called *LocationConstraint*, can be used to specify the physical location of a Bucket
- A RING is installed at multiple locations with the S3 Connector, e.g. RING-West and RING-East
- Each site has its own location constraint
- A PUT Bucket1 LocationConstraint: RING-West will create a Bucket whose objects will be located on the RING-West RING only
- A PUT F1.txt s3://Bucket1 will create the object F1.txt on the RING-West RING only, in Bucket1

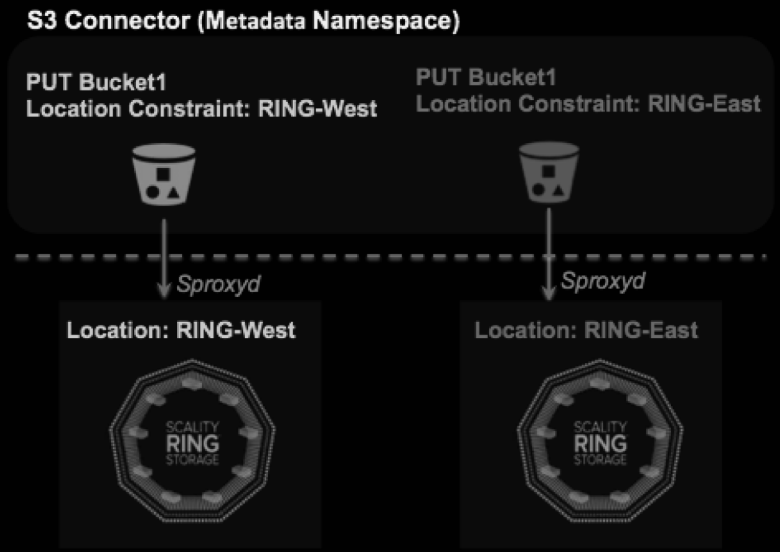

PERFORMANCE GUIDELINES

Number of Fully Provisoned Servers (24x HDD)	Useable TB (12TB HDD)	Throughput Maximum Values (MBps)			
		S3 Write	S3 Read	File Write	File Read
Single Site					
3	472	2086	3202	2039	1706
6	1063	4255	5540	3246	3627
2 Sites Stretched					
6	696	2659	3462	2029	2267
8	929	2947	3837	2249	2513
12	1393	4432	5771	3381	3779
3 Sites Stretched					
6	976	3723	4847	2840	3174
9	1464	4653	6059	3550	3967
12	1942	6205	8079	4734	5290

Test System Configuration

File testing consists of 5x NFS Connectors

S3 testing consists of a single cluster of 5x S3 Connectors

S3 and File maxiumum Write values based on 100% write operations with 1GB files and 50 workers

S3 and File maxiumum Read values based on 100% read operations with 1GB files and 50 workers

Network Connectivity 2 x 10 GbE Per Server including links between sites

SCALITY

Multi-Cloud Data – Challenges

How do I distribute and retrieve data from many different locations in the most efficient way possible?

How can I reduce storage costs?

How do applications and users access data stored in different locations?

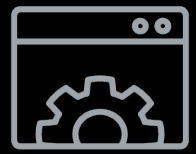

Where is my data?

Multi Cloud Data Controller Capabilities

Application Interfaces	Amazon S3 API
Storage Locations	Scality RING, AWS S3, Azure Blob Storage, Google Cloud Storage, Digital Ocean, Wasabi, Ceph, ECS, NAS (SMB/NFS)
Data Mobility Features	1-1 replication, 1-to-Many replication, Lifecycle expiration, Lifecycle transition
Metadata	System and user-defined metadata; Retrieve objects on metadata attribute values; Global search via an API & GUI
Management with a GUI	 Hosted management portal for easy, point-and-click management

eXtended Data Management

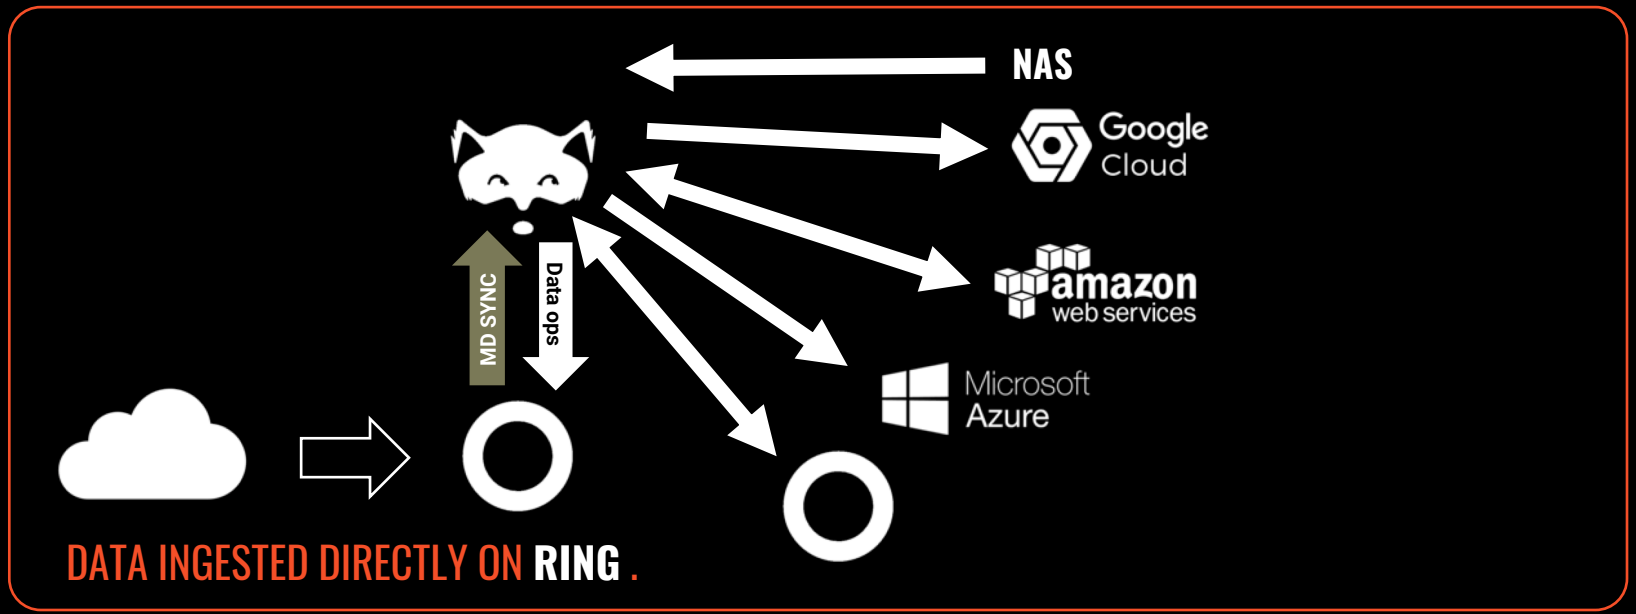

XDM is updated Asynchronously to be able to manage data across the different clouds. XDM is NOT on the data path: We recommend 3 servers based Architecture (possible loss of 1 servers). 5 servers (possible loss of 2 servers) and 1 server (if no HA requirement at all ex: Lifecycle Expiration only) are supported configurations.

HYBRID CLOUD

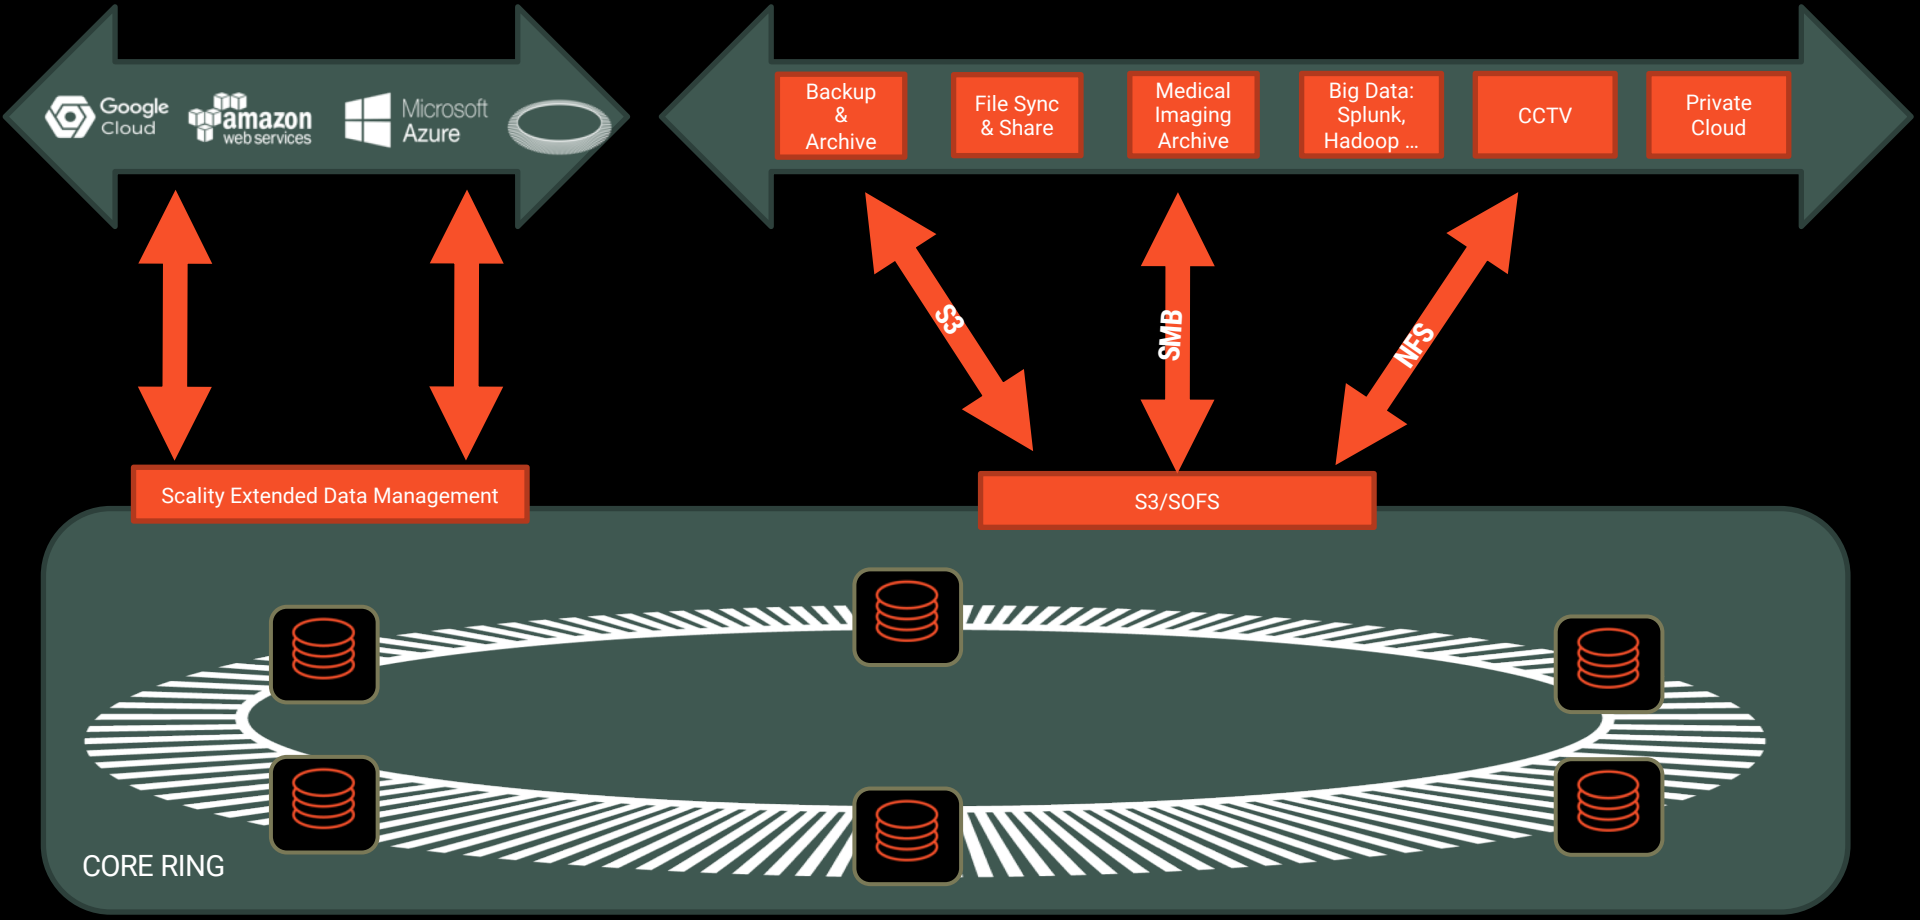

HEALTHCARE & RECHERCHE

SCALITY

Cas d’usage dans la recherche et les hopitaux

- PACS – Cas le plus large

(Scanner, Radiology)

- Digital Pathology – Nvx enjeux

(Microscope OEM with Philips)

- VNA (asset Management pour le contenu)

25% des organismes vont s’équiper d’une VNA

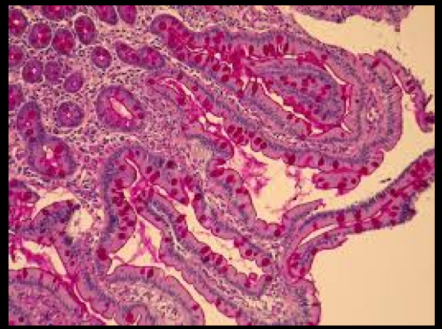

- Archivage Génomique
- Backup, Autres Applications, CCTV

USE CASE: BIG DATA & ANALYTICS

Machine Learning

Autonomous vehicle

Forensic Analysis

Disk/Firmware images, Bodyworn cameras

Log Analysis

Satellite logs

3D models & Imaging

3D Virtual copy, Satellite imaging

Hadoop DataLake

Risk Analysis for FSI

SCALITY

Hadoop Challenges

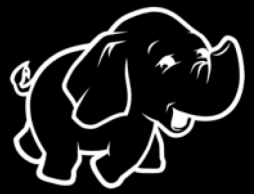

GROWING ONE SINGLE LARGE CLUSTER

- Availability issues & Downtime
- Hadoop is not multi-tenant: All jobs competing for the same resources
- Rigid scaling: Compute vs Storage ratio remains static

DEPLOYING MULTIPLE SMALL CLUSTERS

- Still can't scale storage independently from compute
- Datasets potentially need to be copied on each cluster
- No authoritative data source

GROW COMPUTE INDEPENDENTLY FROM STORAGE

- S3 storage with Scality: On-Prem with RING and Multi-Cloud with Zenko
- Single authoritative data source
- 100% data availability
- Workload isolation

Persistent Storage for Cloudera CDH

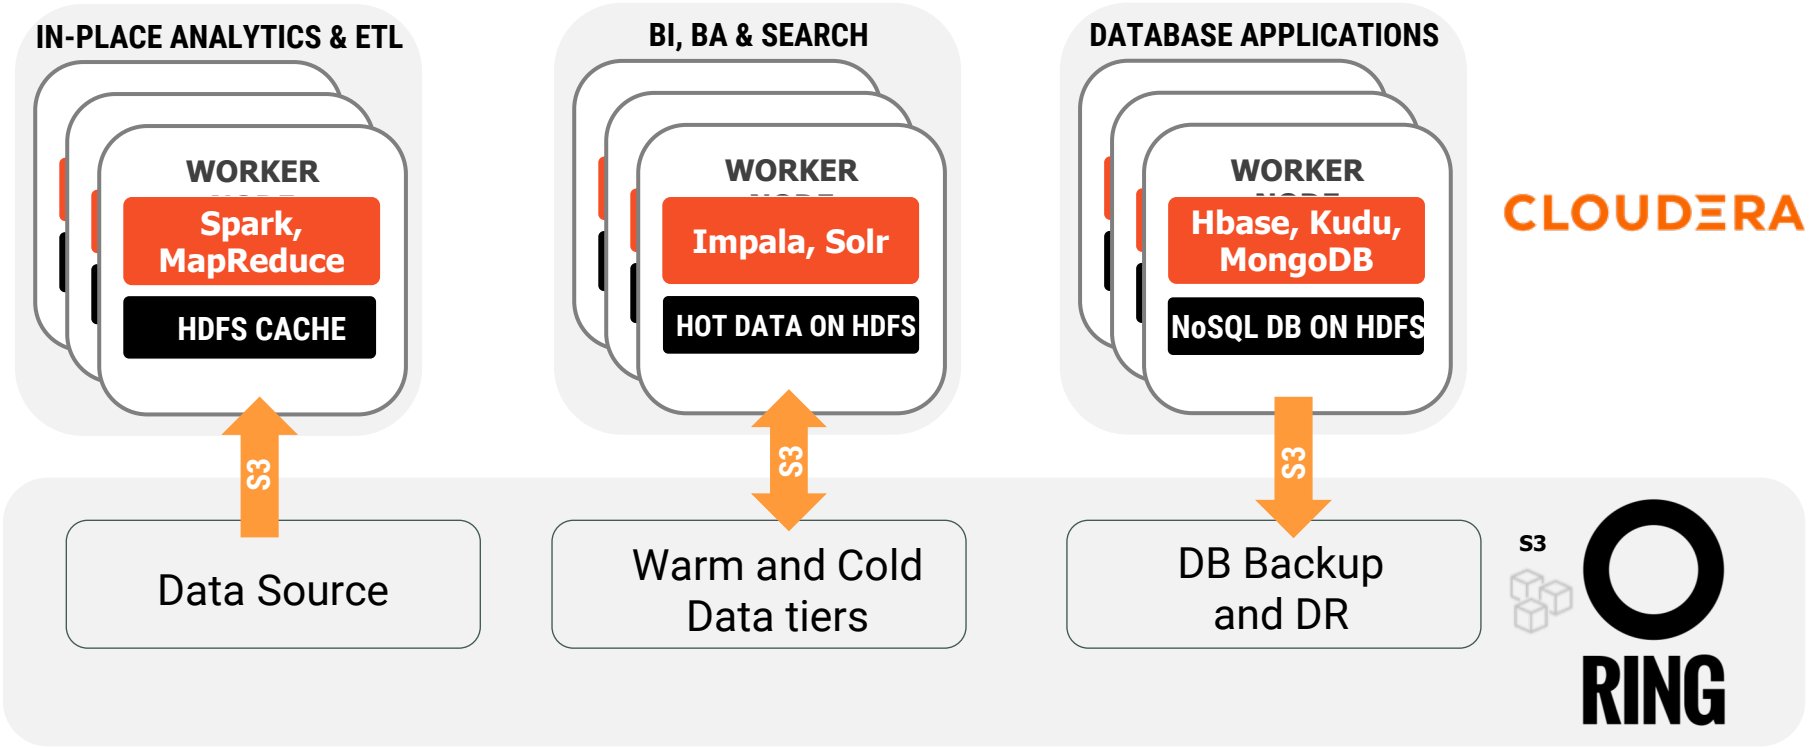

SCALITY

Splunk SmartStore Object Storage

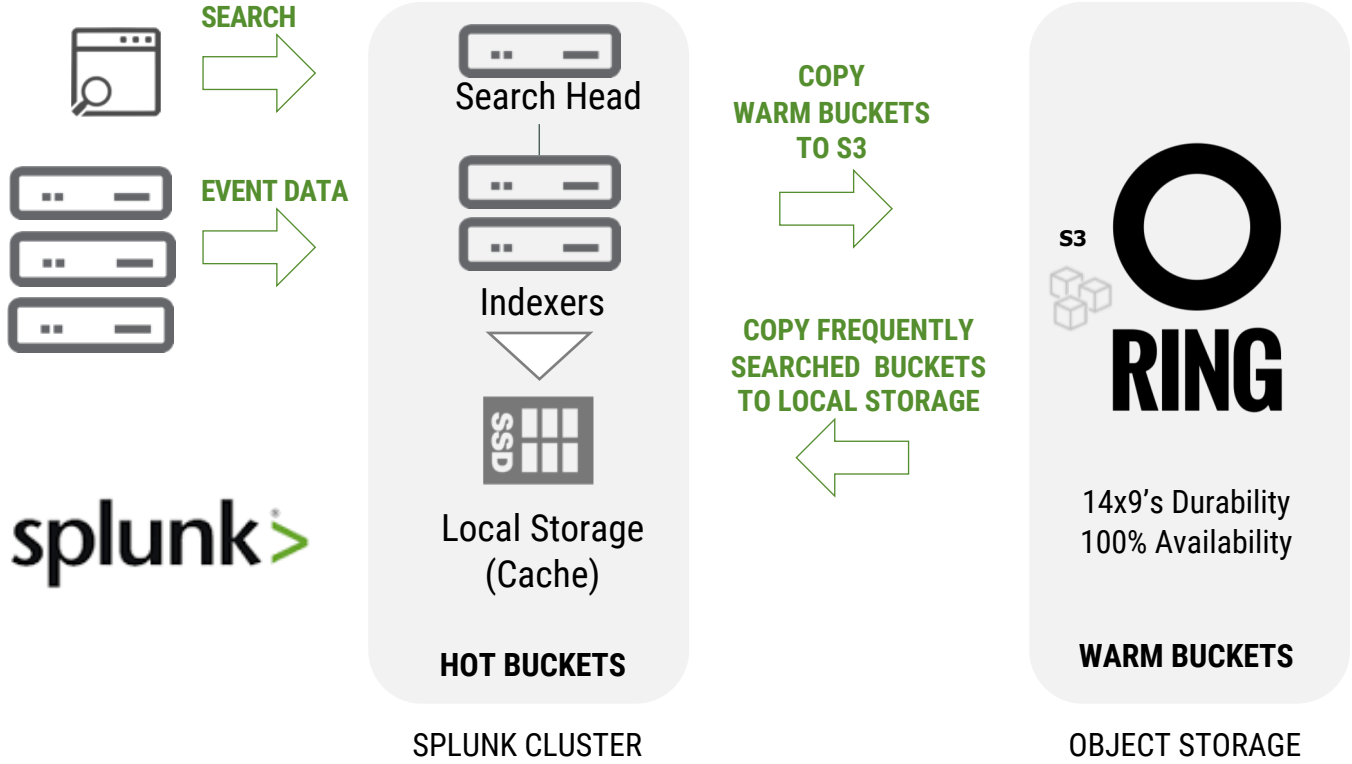

RING Numbers

amadeus
Your technology partner

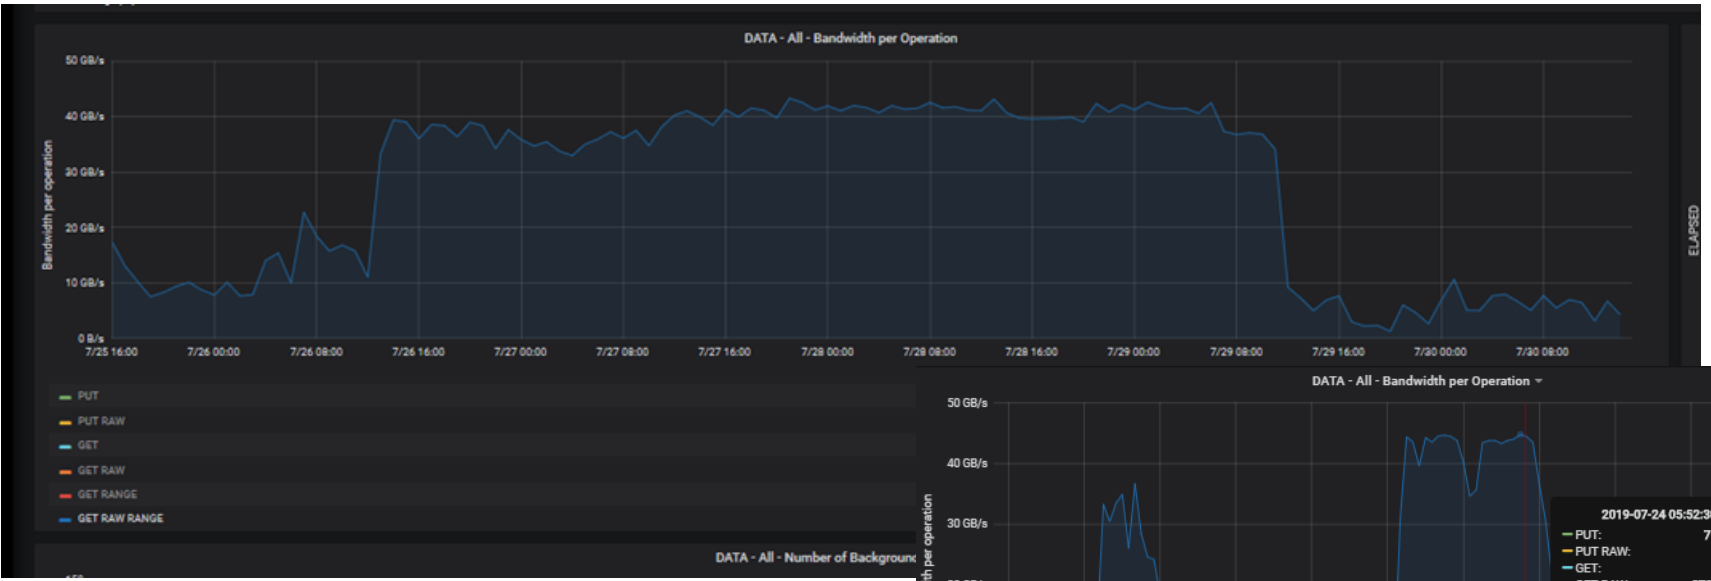

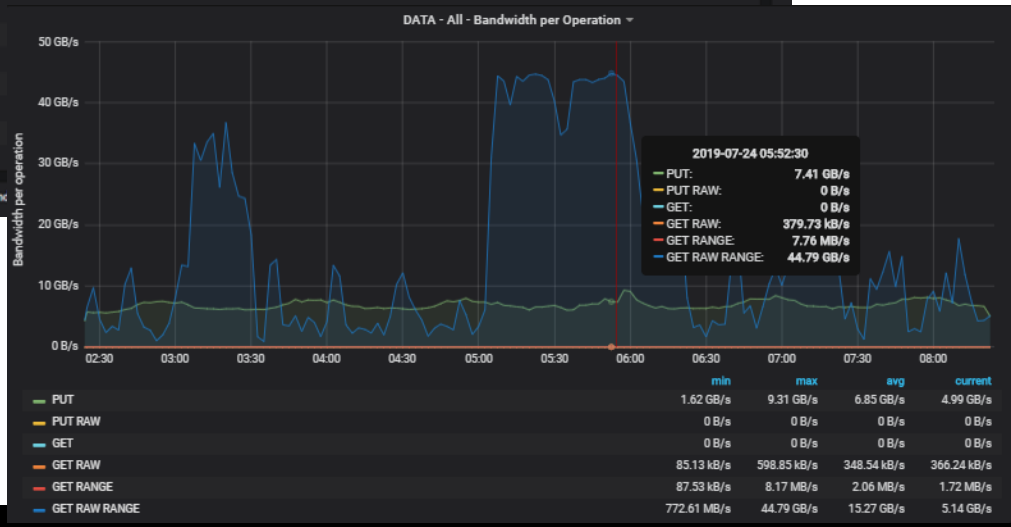

36 Apollo 4200 with 24 HDDs, with about 150 Splunk indexers using Fuse

USE CASE: BACKUP

Backup landscape ...

Cost / TB when ISV doesn't dedup Performance

Resiliency concern / Raid-6

Cheap DR
Can't be hacked

Tape Migration
Management cost
Cost of drives

...is evolving

RING

PRIVATE CLOUD

Highly Durable
WAN optimized
Cost / TB

No Deduplication

ZENKO

PUBLIC CLOUDS

Proper Off Site
Wan optimized

Time to restore
Concern on Public Clouds

SCALITY

Veeam Short and Long Term Retention

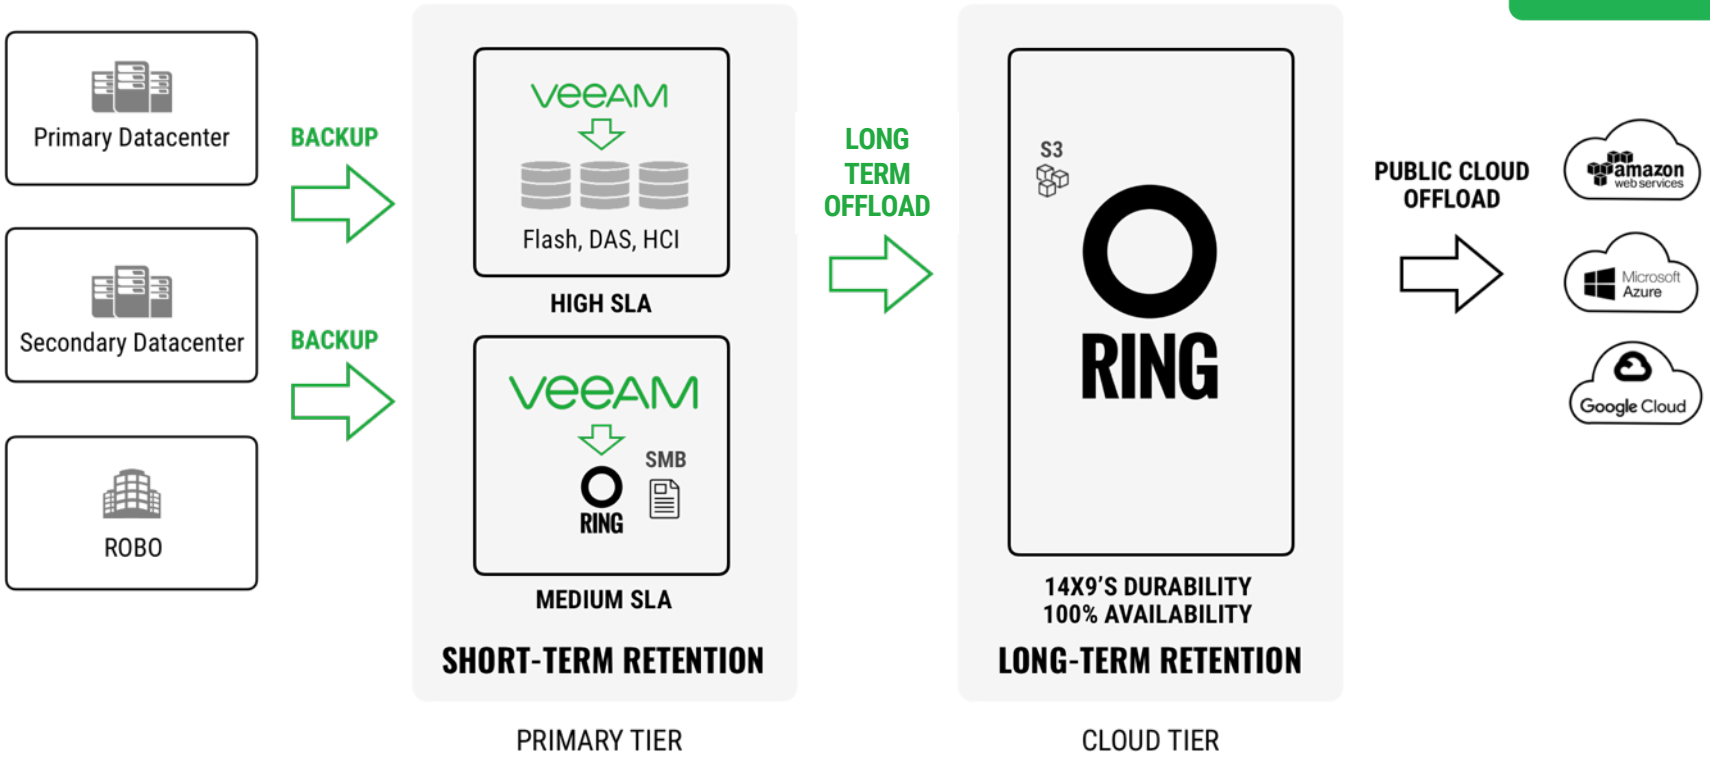

VEEAM READY
Object
Archive

SCALITY

Rubrik DR and Long Term Retention

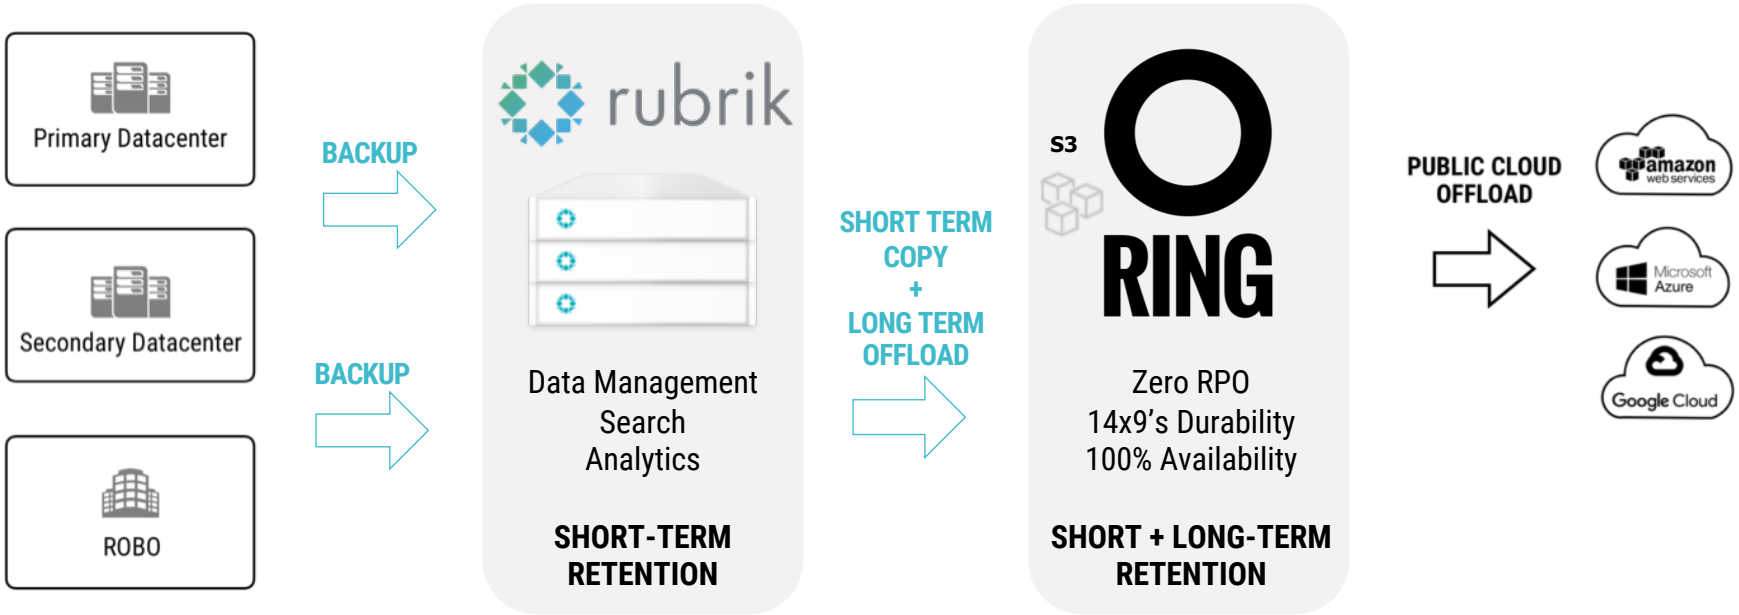

SCALITY

Commvault Short and Long Term Retention

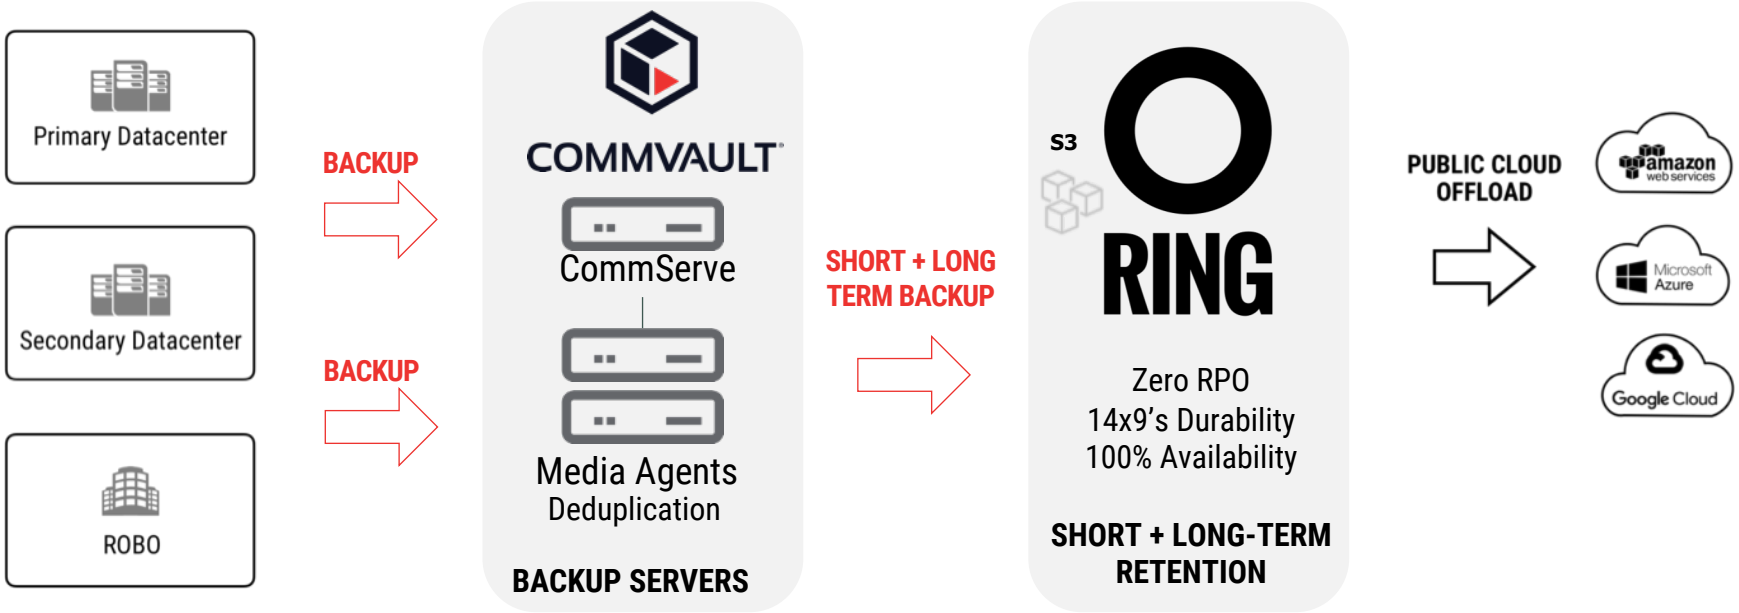

NetBackup Long Term Retention

VERITAS Technology Partner

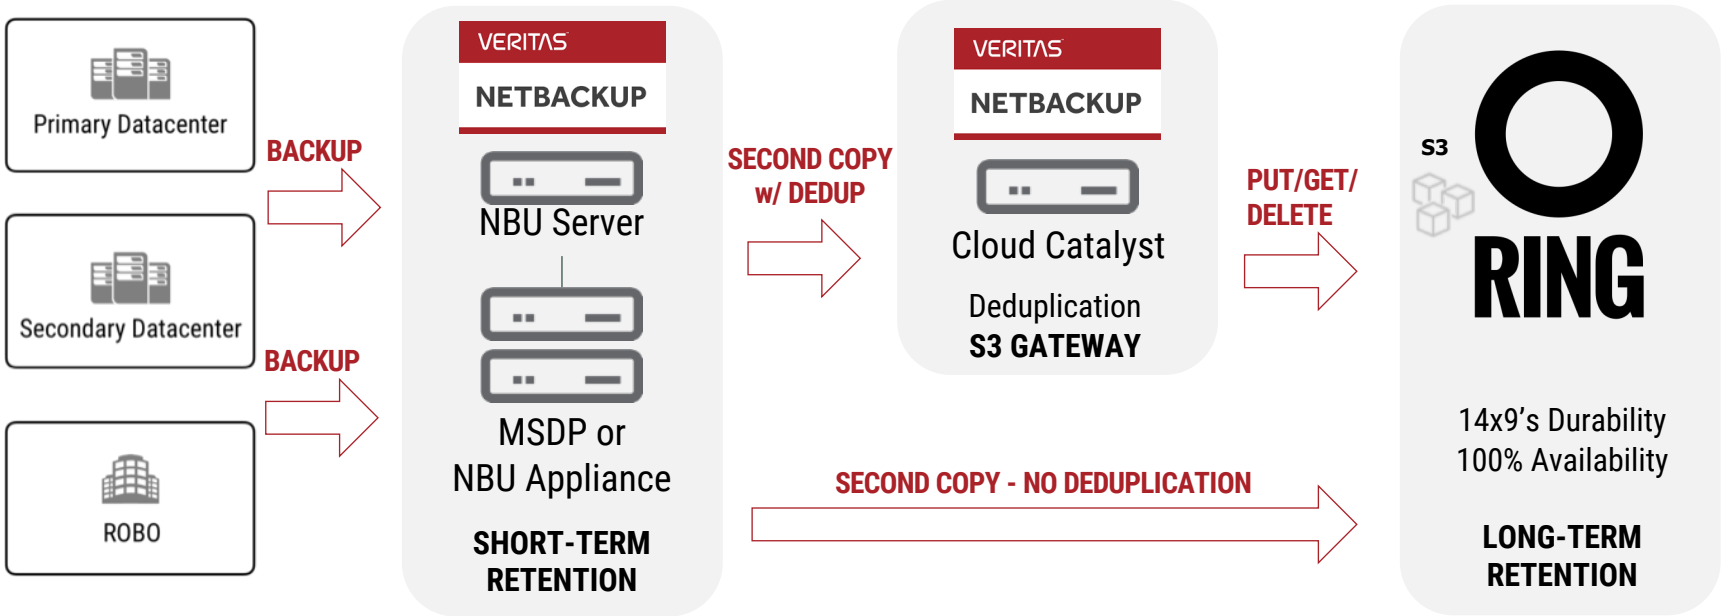

SCALITY

StoreOnce CloudBank Long Term Retention

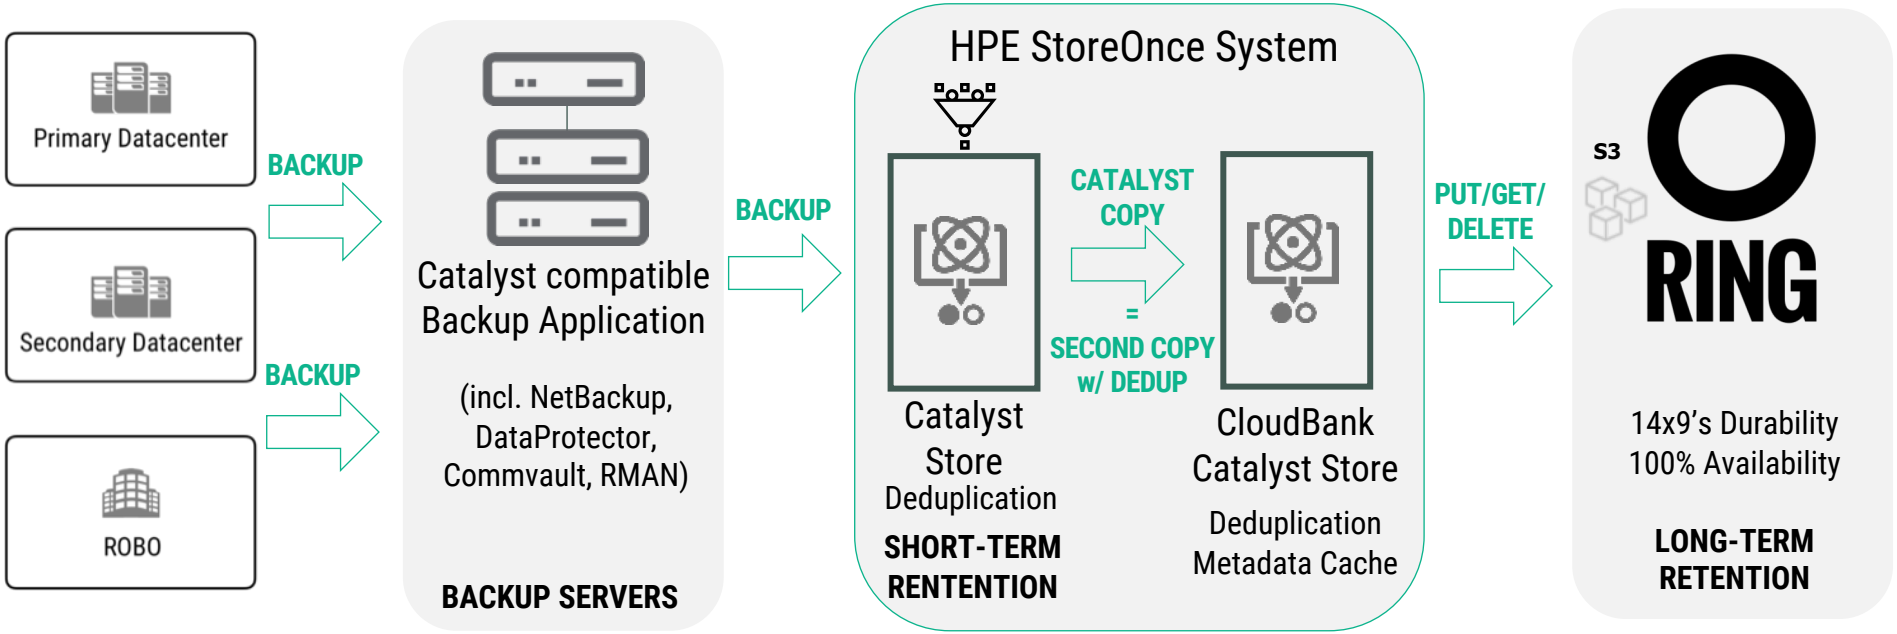

SCALITY

Spectrum Protect
Short and Long Term Retention

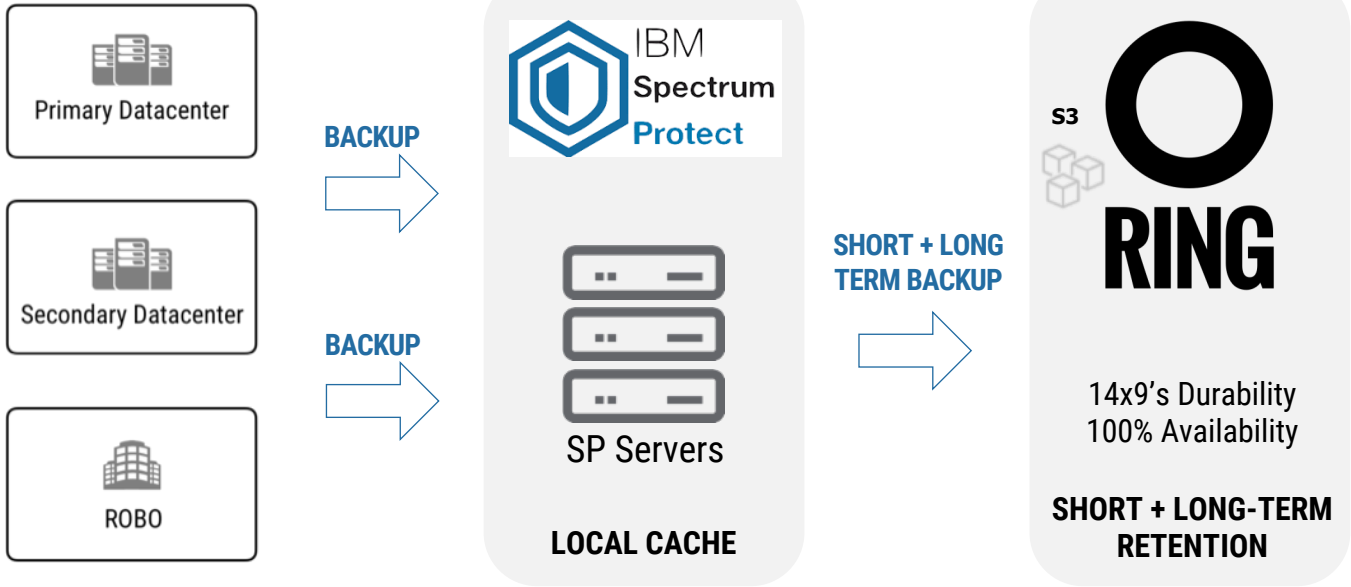

USE CASE: CONTENT REPOSITORIES

SCALITY

CTERA File Server Consolidation

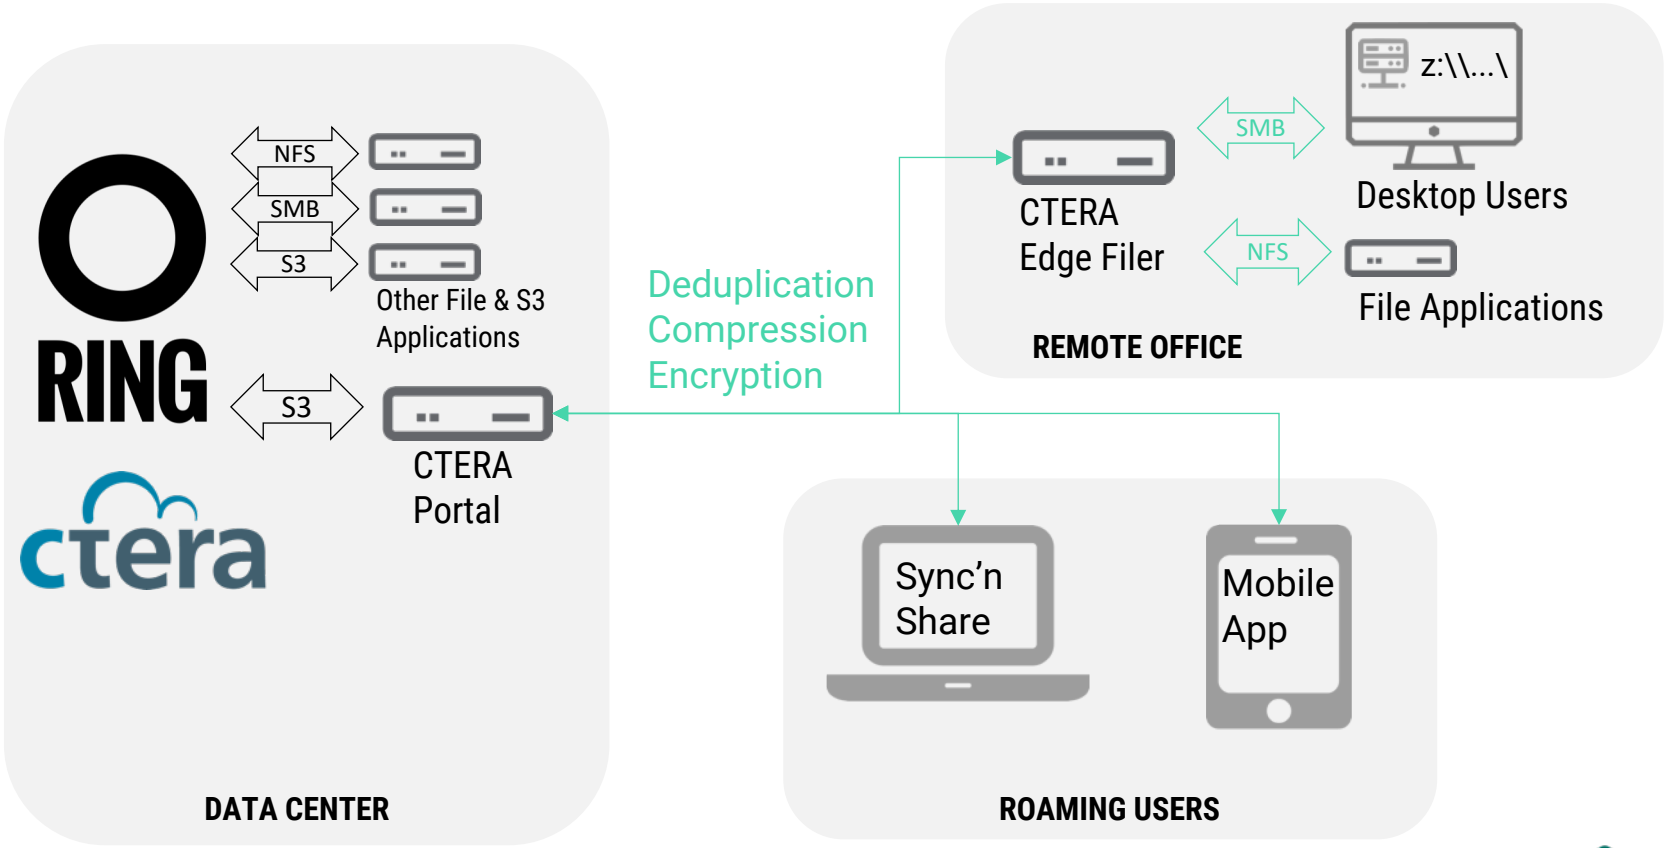

SCALITY

Nasuni File Server Consolidation

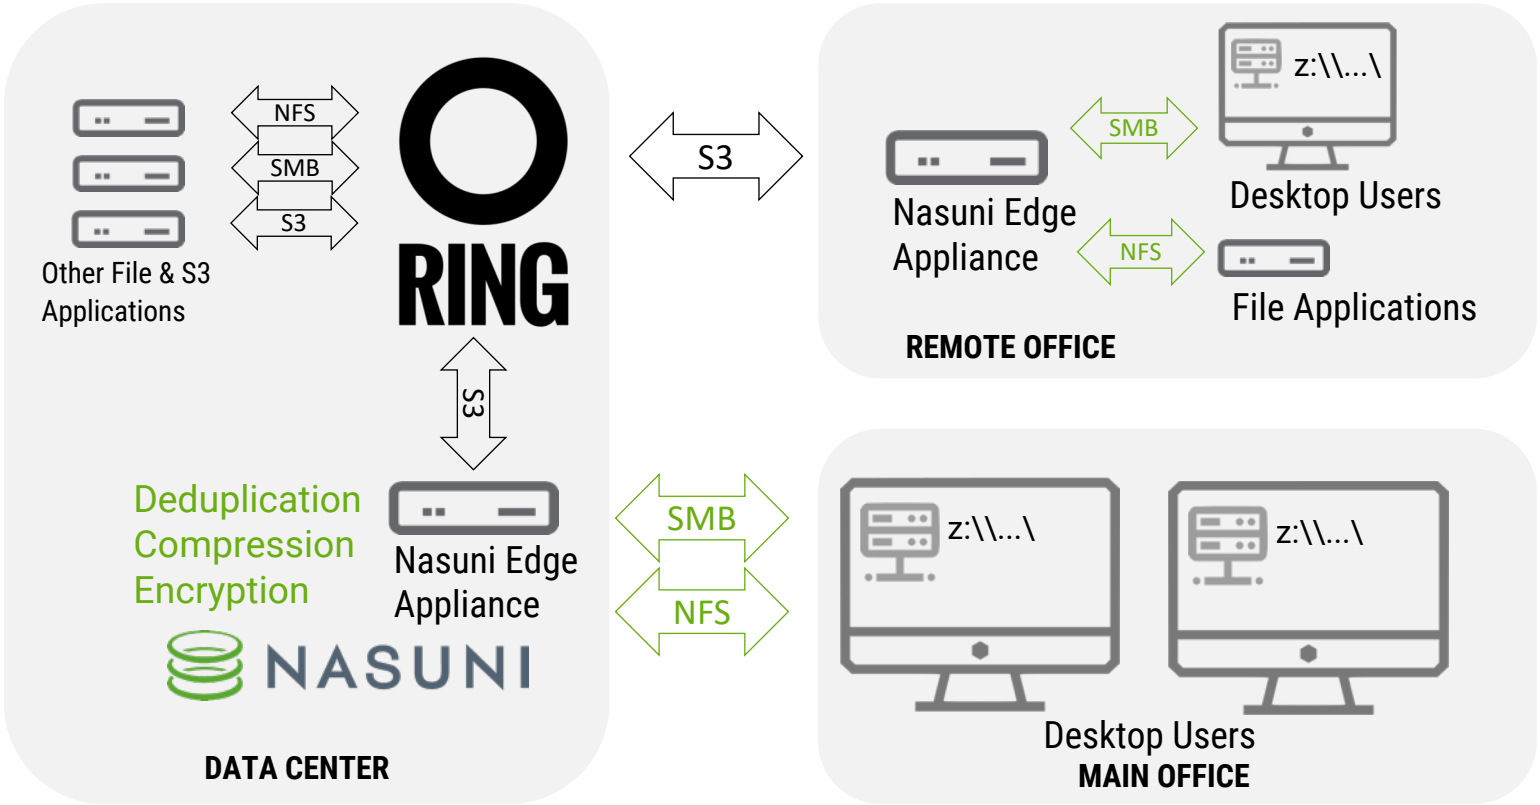

SCALITY

CCTV USE CASE

TRANSFORMING THE CCTV ARCHIVE LANDSCAPE

HIGH PERFORMANCE

Access layer separated from the storage layer allowing one or the other to grow independently, thus optimizing the performance of the solution

GROW WITHOUT LIMITS

Able to add limitless number of storage nodes/disks

Offer both native object and file system protocols

Scale easily and linearly with easy management at multi-petabyte scale

LOW TCO

Allow mix and match standard servers

Grow easily and effortlessly—without limits—by scaling linearly

Support multiple workloads and simply performs

LOWER COSTS

SAFER DATA

EFFICIENT STORAGE

PROVEN PARTNERSHIPS

PERFECT FIT

SCALITY

MILESTONE ARCHITECTURE

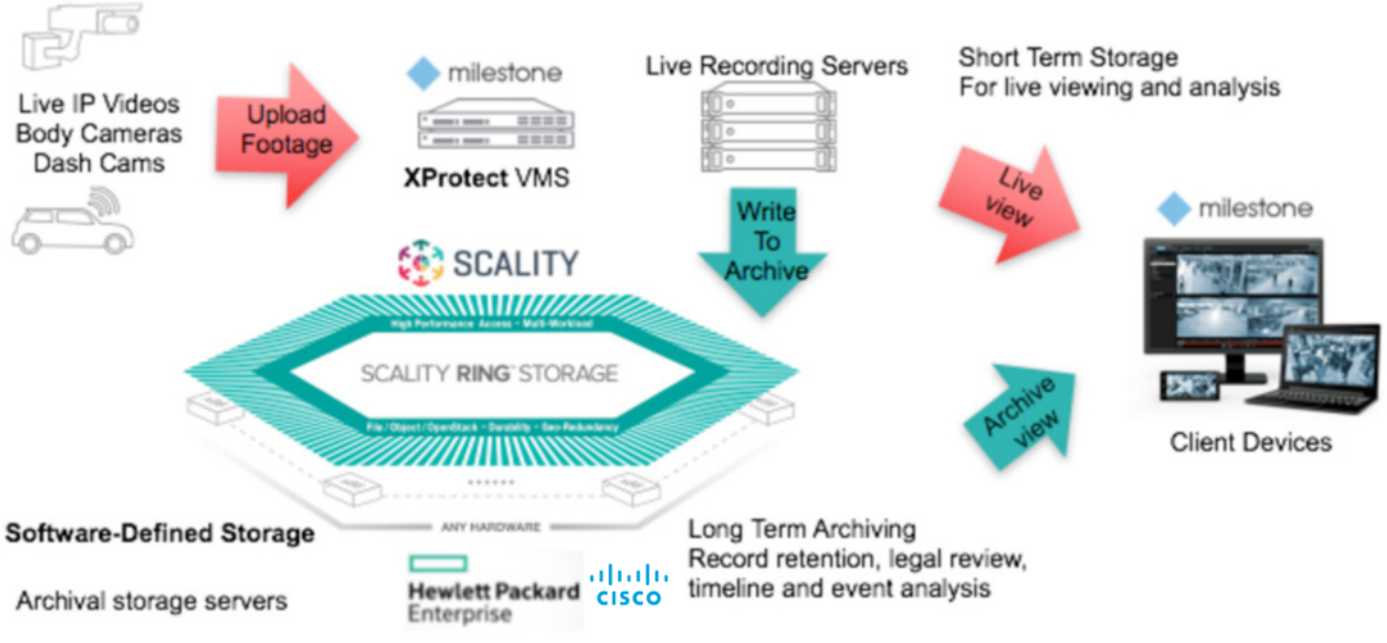

FLIR ARCHITECTURE

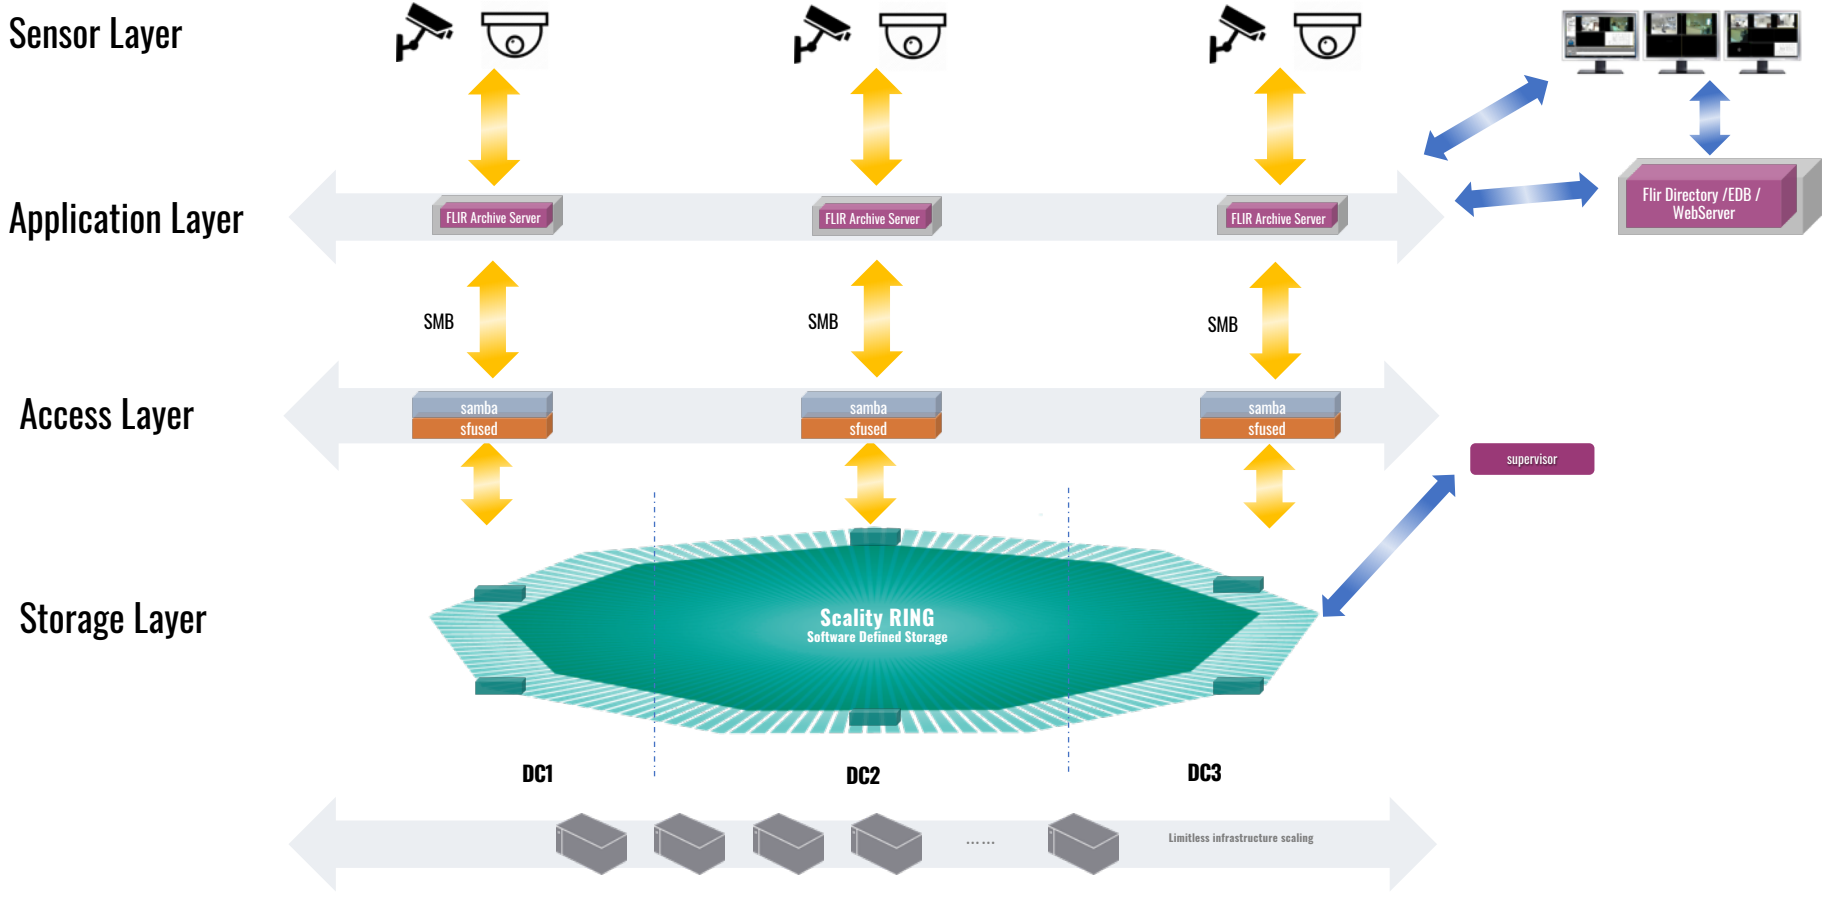

Hardware at edge

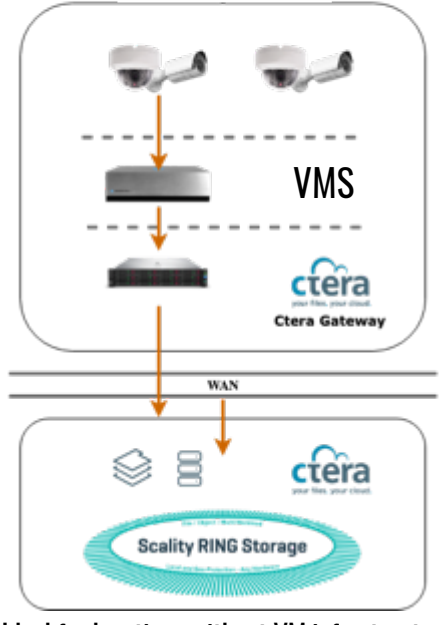

- Ideal for locations without VM infrastructure requiring edge gateway
- Traffic is deduped and compressed

Virtual machine at edge

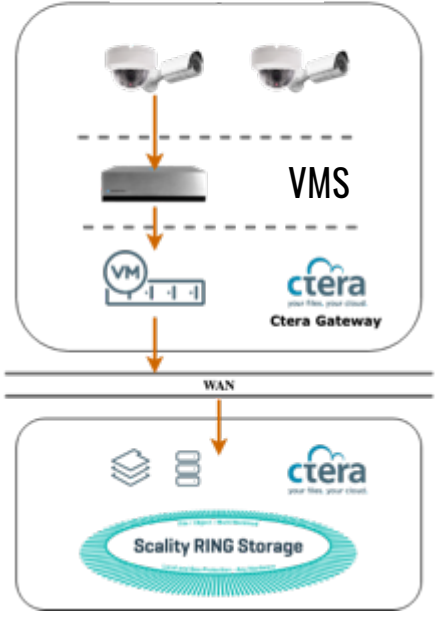

- Ideal for locations with VM infrastructure requiring edge gateway
- Traffic is deduped and compressed

No gateway at edge

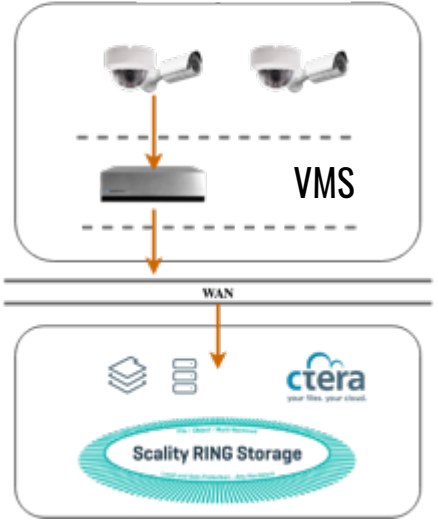

- Ideal for locations when not requiring edge gateways and may be ideal for tiny sites
- Traffic streams raw to data center

nicolas.rivaton@scality.com / +33 6 66 01 90 52

Thank You.

SCALITY

SAN FRANCISCO, USA

149 New Montgomery Street,
Suite 607
San Francisco, CA, 94105

Email: sales.us@scality.com
Telephone: +1 (650) 356-8500
Fax: +1 (650) 356-8501
Toll Free: +1 (855) 722-5489

PARIS, FRANCE

11 rue Tronchet
Paris, 75008

Email: sales.eu@scality.com
Telephone: +33 1 78 09 82 70

WASHINGTON, D.C., USA

43777 Central Station Drive, Suite 410
Ashburn, VA 20147, USA

Email: sales.us@scality.com
Toll Free: +1 (855) 722-5489

BOSTON, USA

50 Milk Street, 16th Floor
Boston, MA 02109, USA

Email: marketing@scality.com

TOKYO, JAPAN

Otemachi Bldg. 4F, 1-6-1, Otemachi
Chiyoda-ku Tokyo, 100-0004 Japan

Email: sales.japan@scality.com
Telephone: +81-3-4405-5400

DUESSELDORF, GERMANY

Speditionstrasse 21
40221 Düsseldorf

Email: sales@scality.com
Telephone: +49 211 88231 723

LONDON, UNITED KINGDOM

20 St Dunstans Hill
London, United Kingdom, EC3R 8HL

Email: sales@scality.com
Telephone: +44 203 795 2434