IBM Storage Ceph

IBM Storage Ceph Object

Daniel Parkes
IBM Storage Ceph Product Manager


ceph

IBM Storage Ceph Object Storage Overview Agenda

01 Introduction to Ceph

02 Introduction to Ceph Object Storage

03 Ceph Object Performance at Scale

04 Ceph Object new features in 8.0

05 Ceph Object ZeroTrust LakeHouse

06 Ceph Object Looking into the future

07 Question & Answers

IBM Storage Ceph

Enterprise-grade distributed, universal, software defined storage solution.

A single, efficient, unified storage platform for object, block, and file storage

Enterprise support and services, certified updates and service level agreements for use in production environments.

IBM Storage ready nodes are a ready to run solution for IBM Storage Ceph, including hardware and software.

Applicable for all known IBM Storage Ceph workloads.

Enabling businesses for the quickest time to market in a convenient way, to run IBM Storage Ceph in production.

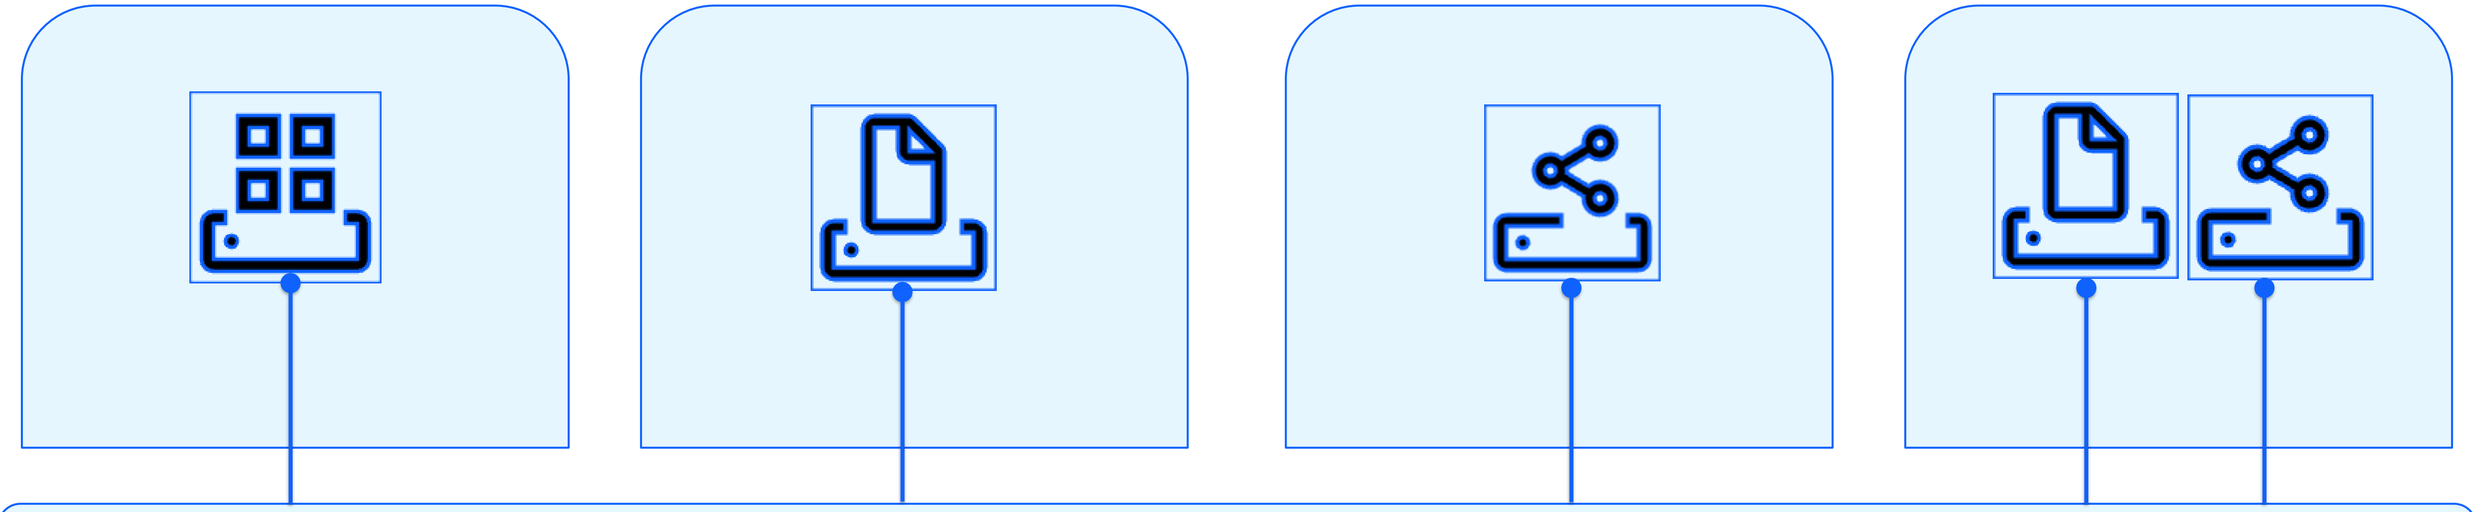

IBM Storage Ceph - Unified Software Defined Storage Platform

Block Storage

Data storage that divides data into fixed-size blocks.

Accessed through Ceph RBD and accessible via native RBD driver or NVMe/TCP for non-Linux clients.

File Storage

File storage organizes data into files, directories and subdirectories.

Accessed through CephFS native client or NFS for non-Linux clients.

S3 Object Storage

Object storage organizes data in discrete units named objects.

Each object consists of the data itself, metadata and a unique identifier Accessed via S3 API.

Unified File & Object

Access the same data resources by file and/or object simultaneously.

In example, ingest data by file access, through NFS, and analyze that same data through S3 API, used by analytics.

IBM Storage Ceph Consumption options

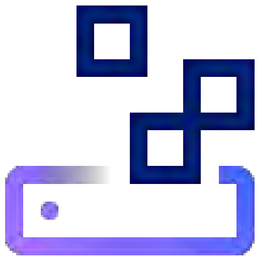

Storage Ceph

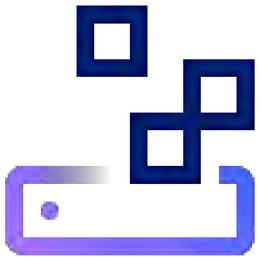

Storage Ready Node

IBM Storage Ceph Software only

Install and run IBM Storage Ceph on industry-standard x86 server hardware of a company preferred hardware vendor.

IBM Storage Ceph on Ready Nodes

A flexible way for organizations to quickly build an on-prem cloud-scale software-defined storage solution with software and hardware support provided by IBM.

IBM Storage Ceph as-a-Service

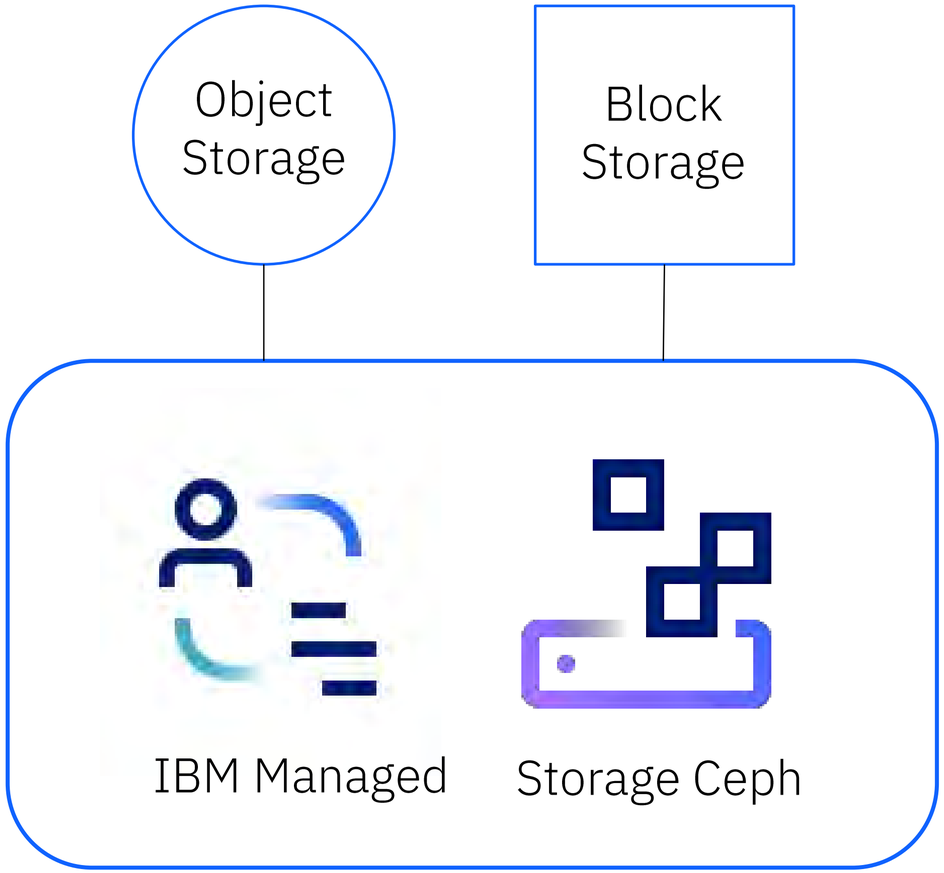

IBM Storage Ceph as-a-service

IBM Storage Ceph as-a-Service is flexible, easy to use, cost-efficient, and based on service level agreements.

Have Ceph functionality, outsource installation, management and maintenance to IBM.

IBM Storage Ceph benefits

IBM Storage Ceph as-a-service easily integrates with other tools and platforms, streamlining workflows and improving efficiency.

IBM clients always have access to latest features and security patches without disruptions or business downtime.

IBM Storage Ceph Added Value

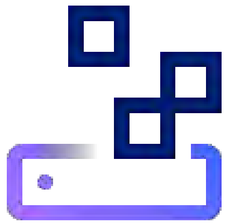

Enterprise Support

24x7x365 support

For urgent issues, including security response patches from IBM Storage Ceph experts.

Enterprise Services

Professional implementation, upgrade, migration and health-check services from IBM Expert Labs
In-house training capabilities.

Enterprise Security Quality Engineering

Penetration testing, Threat modeling, Code scans, Open-source license clearance, HA+DR testing

Extended Lifecycles (2+1+3)

2 years Regular support
1 year Extended Support with critical bug fixes
3 years Extended Support w/o bug fixes

Advanced Support Option

Management and integration

Consolidated monitoring with Storage Insights
Platform and solutions integration and certifications:
VMware, WatsonX, WORM Compliances

Storage Ready Nodes

IBM qualified and supported hardware for running IBM’s Storage Ceph cluster platform.

IBM Storage Ready Nodes are a simple, cost-effective way to run IBM Storage Ceph.

IBM Storage Ceph Object Storage Overview Agenda

01 Introduction to Ceph

02 Introduction to Ceph Object Storage

03 Ceph Object Performance at Scale

04 Ceph Object new features in 8.0

05 Ceph Object ZeroTrust LakeHouse

06 Ceph Object Looking into the future

07 Question & Answers

IBM Storage Ceph Object

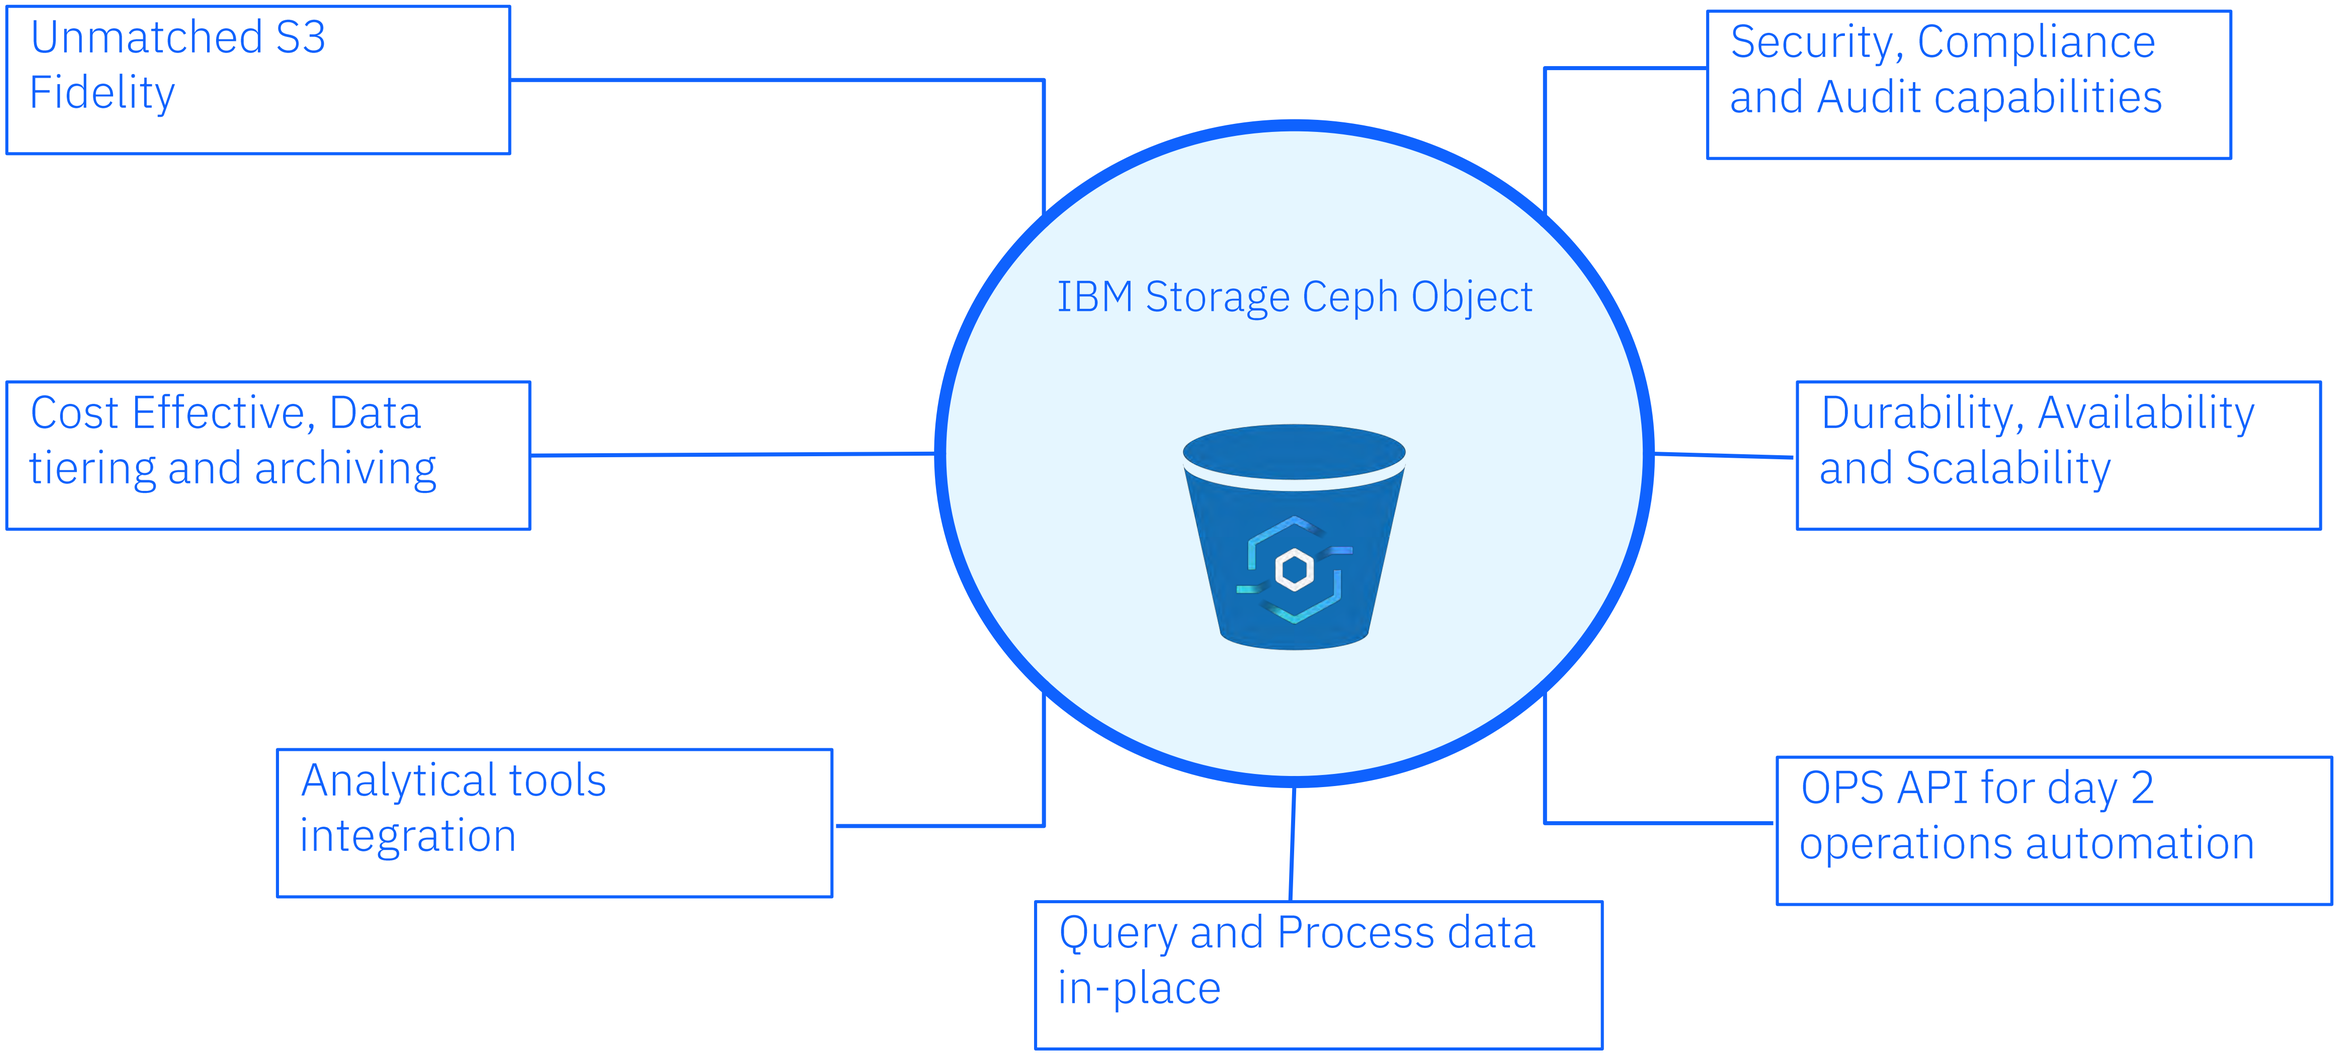

IBM Storage Ceph Object Storage Overview Agenda

01 Introduction to Ceph

02 Introduction to Ceph Object Storage

03 Ceph Object Performance at Scale

04 Ceph Object new features in 8.0

05 Ceph Object ZeroTrust LakeHouse

06 Ceph Object Looking into the future

07 Question & Answers

IBM Storage Ceph Horizontal Scaling and Performance

Architecture designed for Massive Scale

Designed for unlimited Scale. Scale out within a single cluster for capacity and performance

Client to Cluster Architecture

Ceph doesn’t have a centralized metadata server. Client’s directly read from and write to all OSDs in the cluster

Increase Performance

Adding new nodes to Ceph will not only increase the capacity but will also significantly improve the performance.

Transparent Scale-in/out

Ceph will transparently re-balance data as you add more nodes/OSDs to the cluster.

You can also dynamically Scale-in, reducing the number of OSDs

IBM Storage Ceph Horizontal Scaling and Performance

Enterprise Ready.

Demonstrated scalability to more than 10 billion objects in a six-node, 4.5 PB IBM Storage Ceph cluster.

One Billion Objects in a single bucket.

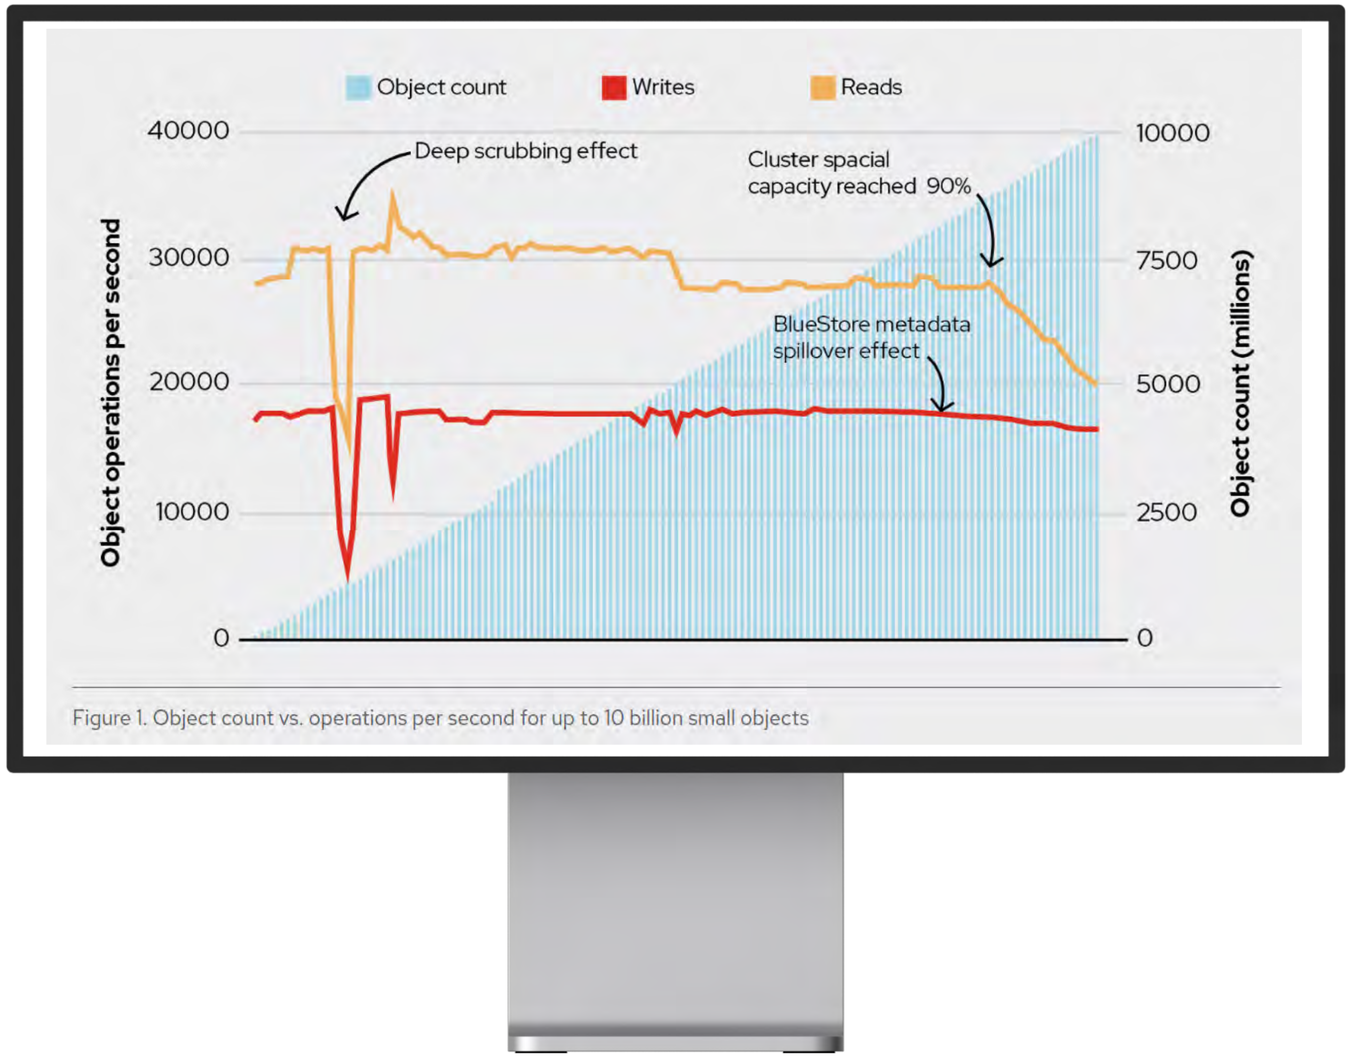

Figure 1. Object count vs. operations per second for up to 10 billion small objects

IBM Storage Ceph Horizontal Scaling and Performance

Increase your total cluster throughput by scaling-out the number Rados Gateway daemons in the cluster

IBM Storage Ceph Horizontal Scaling and Performance

Increase total Throughput as you horizontally scale your cluster by adding new nodes.

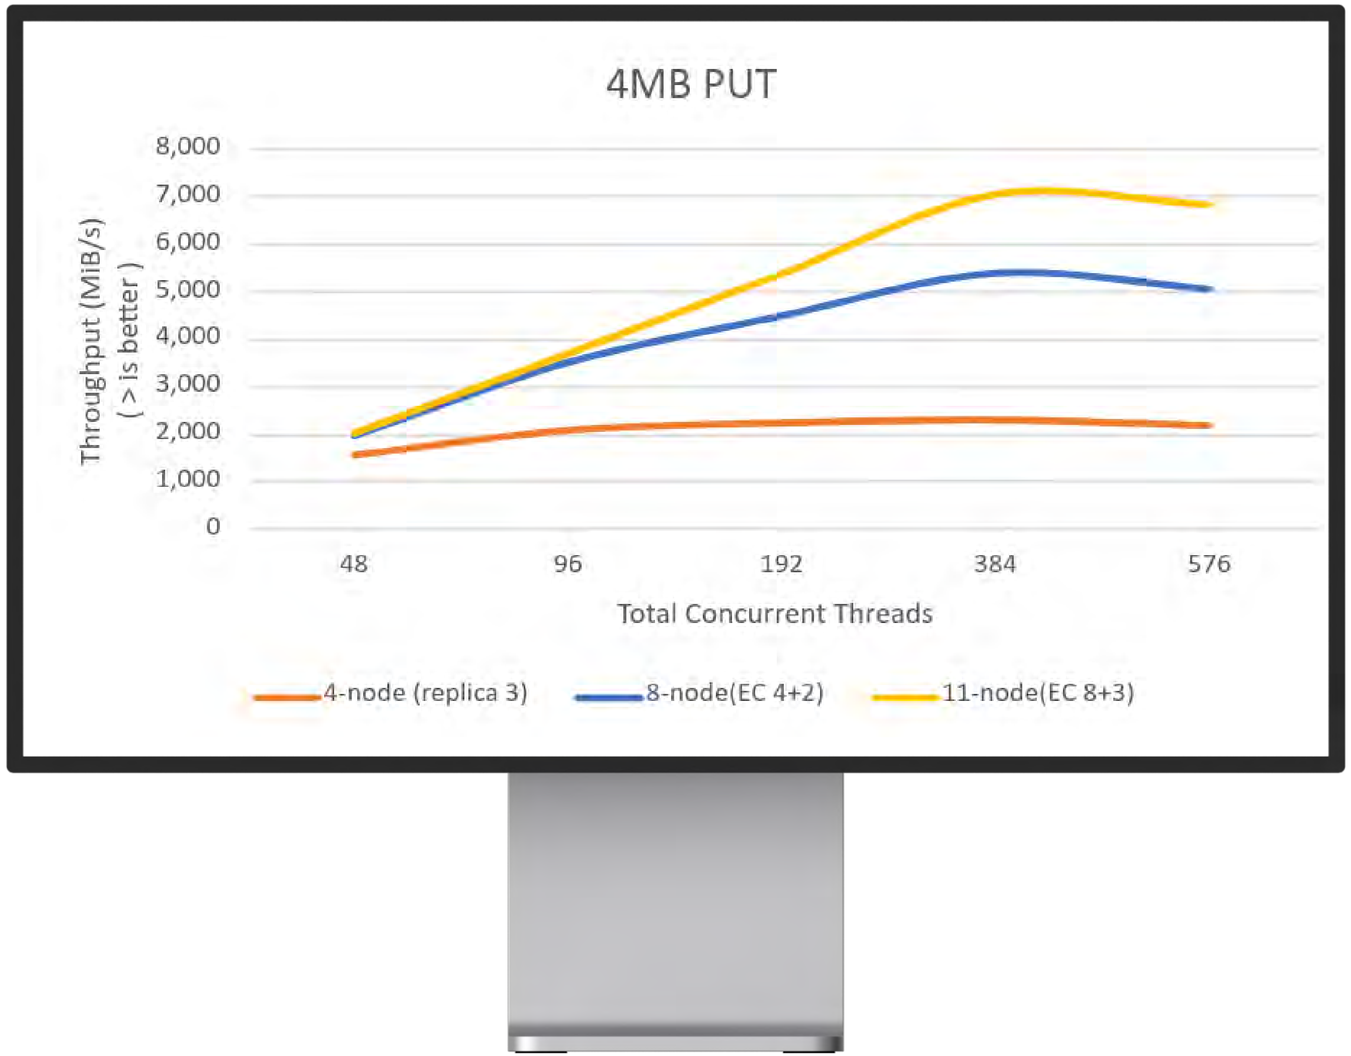

IBM Storage Ceph Horizontal Scaling and Performance

Performing under failure Conditions.

IBM Storage Ceph designed to ingest and serve billions of objects was resilient to both device and node failure, allowing processing to continue at a high level, even under failure conditions.

IBM Storage Ceph Object Storage Overview Agenda

01 Introduction to Ceph

02 Introduction to Ceph Object Storage

03 Ceph Object Performance at Scale

04 Ceph Object new features in 8.0

05 Ceph Object ZeroTrust LakeHouse

06 Ceph Object Looking into the future

07 Question & Answers

IBM Storage Ceph Object S3 Bucket Logging

“Gain deeper insights into application data access with flexible, S3-based Bucket logging.”

Monitor Bucket Activity

Enables end users to configure and manage their application storage access logging per bucket through the familiar S3 API

Use Cases

- Monitor bucket activity
- Detect unauthorized access
- Get insights into the bucket usage
- Enhanced Incremental Backup/Restores

Dedicated Log Bucket

By leveraging Ceph's logging features, users gain actionable insights through logs stored in dedicated buckets

Configured through S3 API

Logging configuration is done at the bucket level and can be
enabled, disabled at any time through the S3 API by the end use

IBM Storage Ceph Object Storage

Unlocking Insights with Trino and Superset: Visualizing Your Application Logs

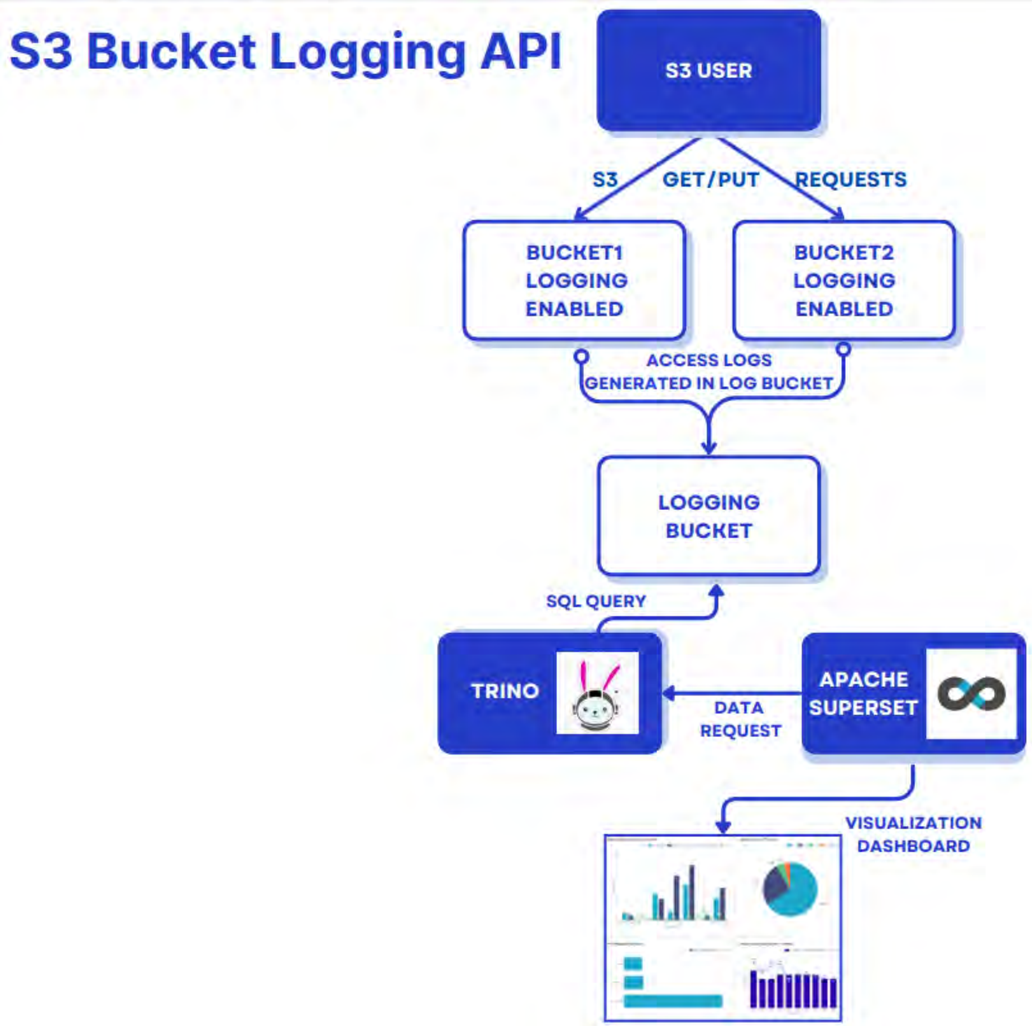

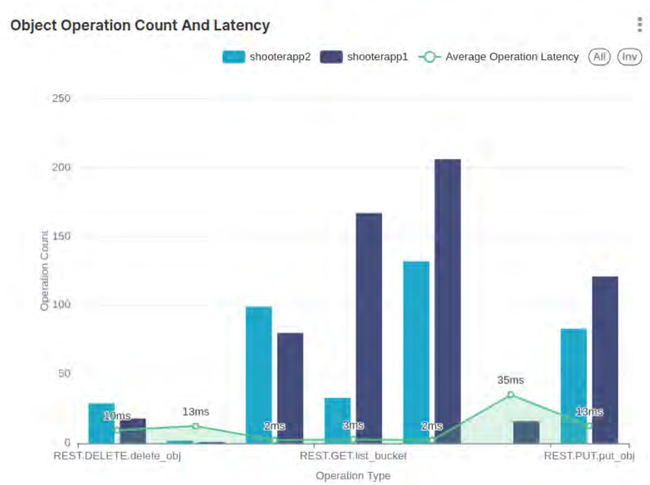

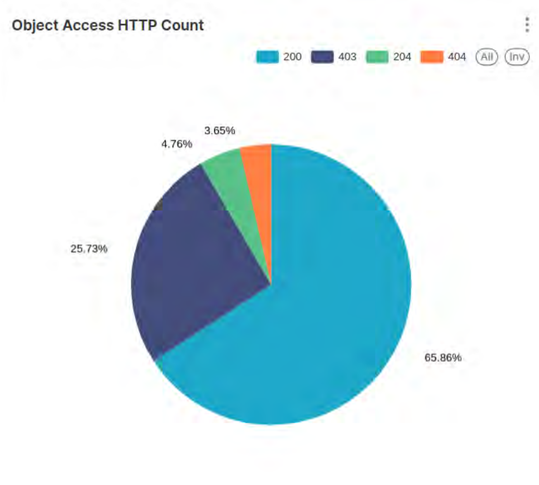

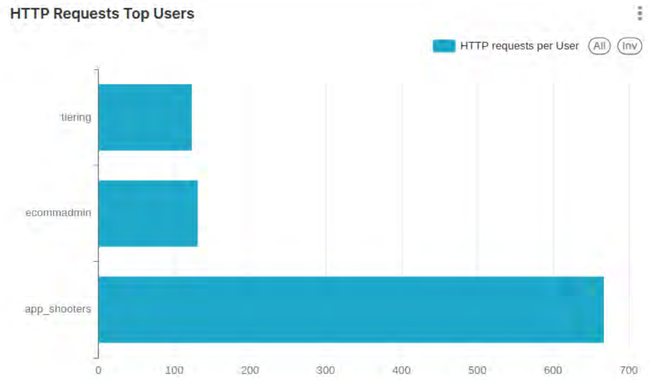

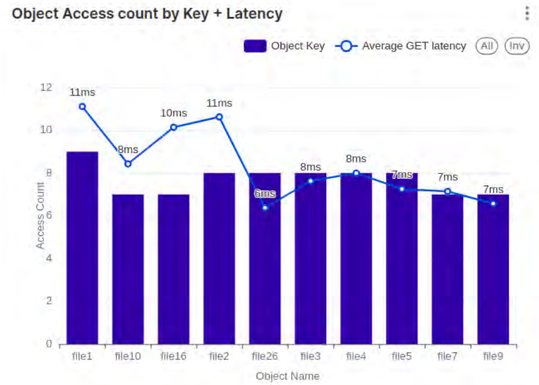

Multitenancy: IAM Accounts

Tenant Level Control

An IAM root user for each tenant/Account can manage users, roles, access keys, quotas, and buckets using the IAM API.

Delegate Administration

Storage Ceph admin creates IAM accounts; each tenant root user self-manages his account resources.

Use IAM API

Hand off administrative tasks to individual teams (e.g., dev groups) without exposing all Ceph resources.

Improved Multitenancy

Fully isolate tenants while sharing the same underlying Ceph cluster.

Lower Central Overhead

Enhancing IBM Storage Ceph Object with **IAM Accounts**

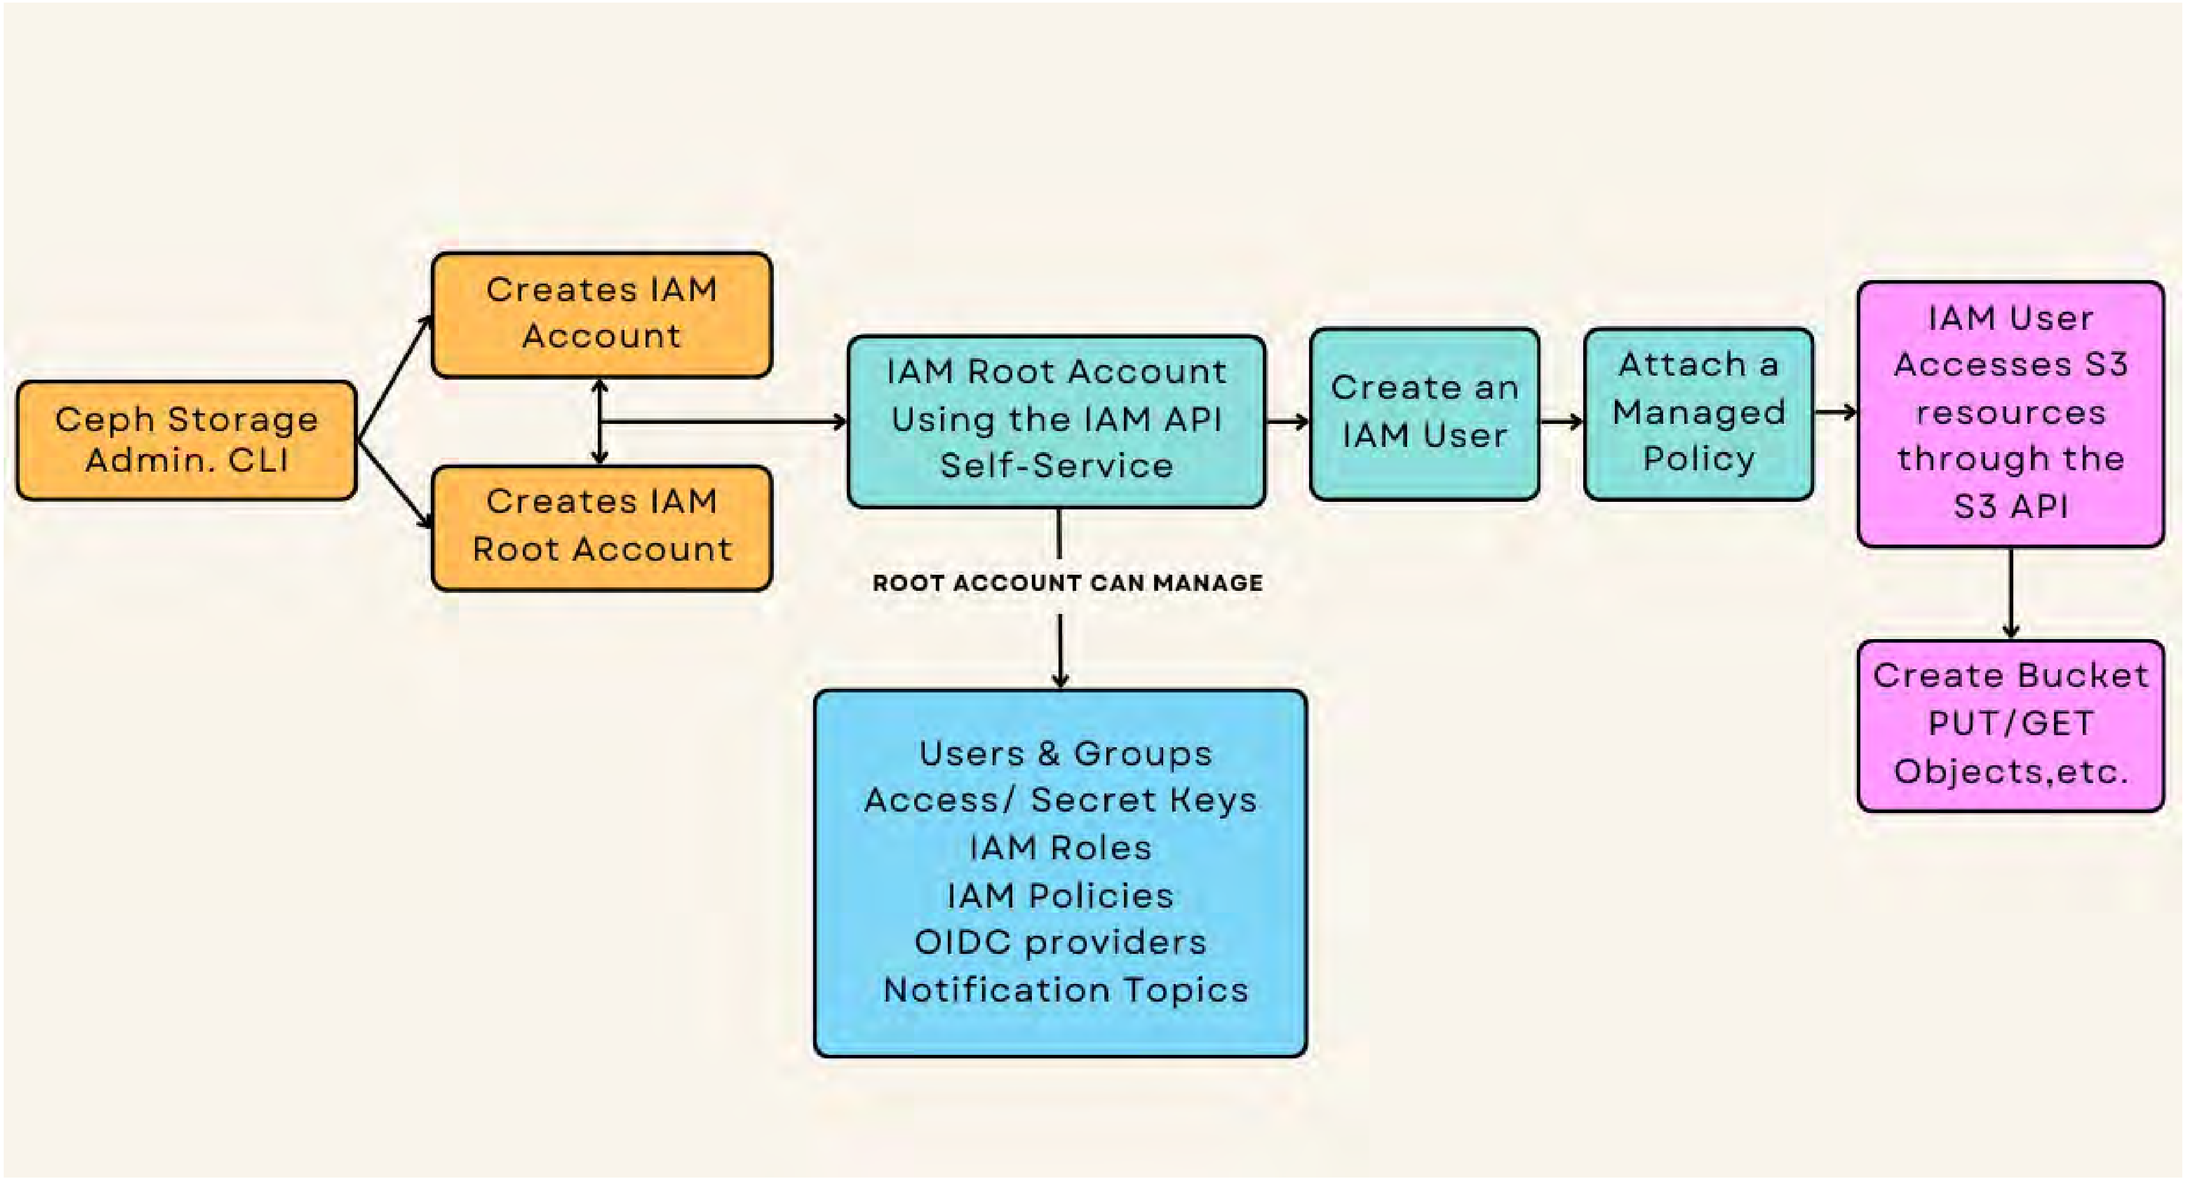

IBM Storage Ceph Object S3 Lifecycle Management

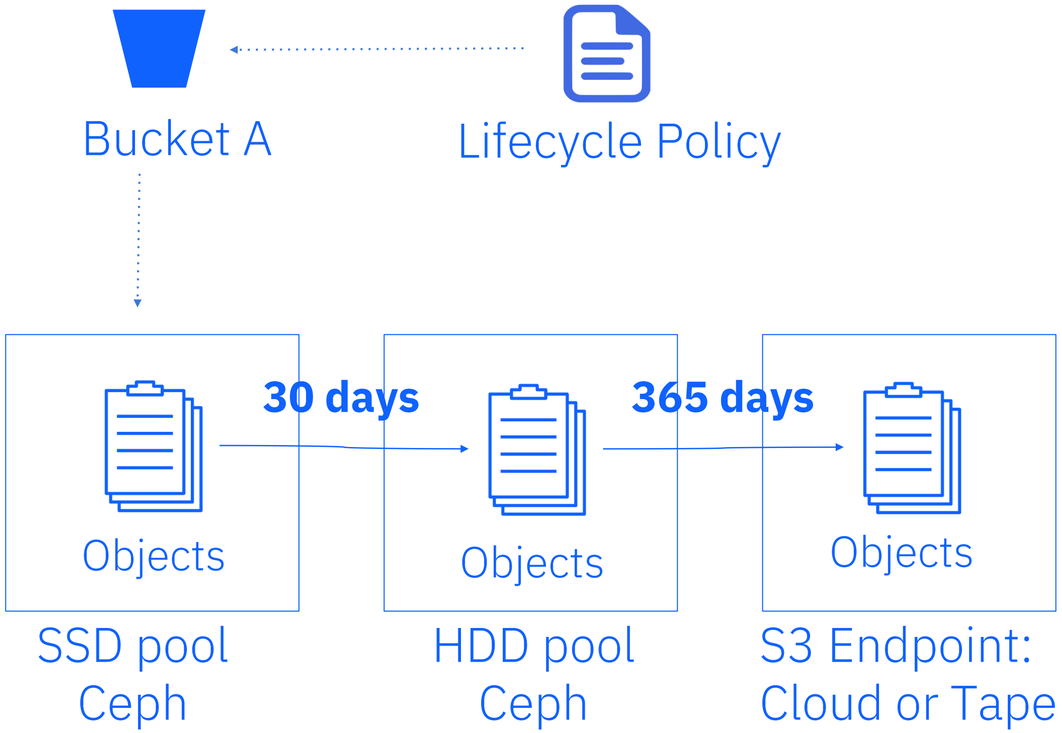

Automatic Transition & Expiration

Easily shift objects between tiers (SSD, HDD, tape, or cloud) after a set number of days.

Granular Policies

We can granularly filter which objects in a bucket are susceptible of LC by Object Prefix or Tags.

S3 Lifecycle Transition can be defined between arbitrary storage classes(Tiers) inside the same cluster or to different S3 compatible endpoints(Including AWS, IBM Cloud, etc)

S3 Lifecycle Expiration includes filters:

-Current/Non-current
-Delete Marker Expiration
-Abort MultiPart Upload
-NewerNoncurrentVersions
-ObjectSizeGreaterThan
-ObjectSizeLess

IBM Storage Ceph Object

Policy Based Archival

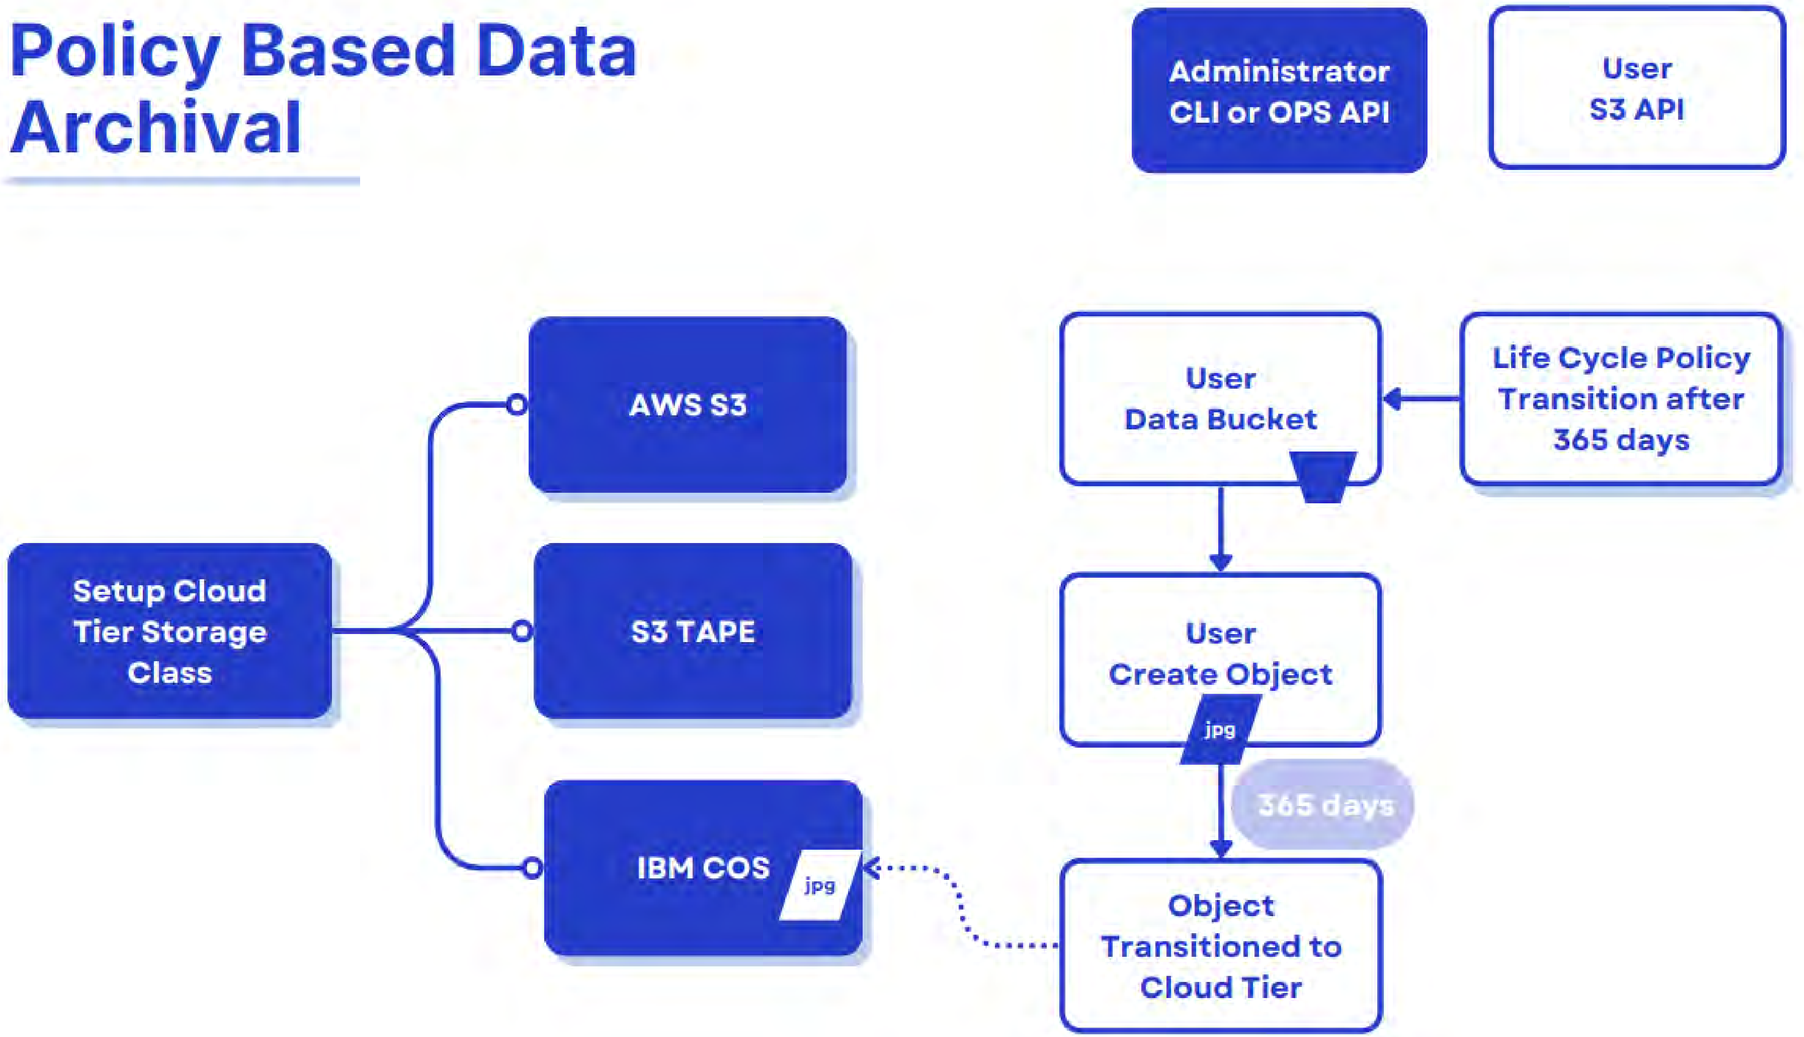

IBM Storage Ceph Object
Policy Based Retrieval

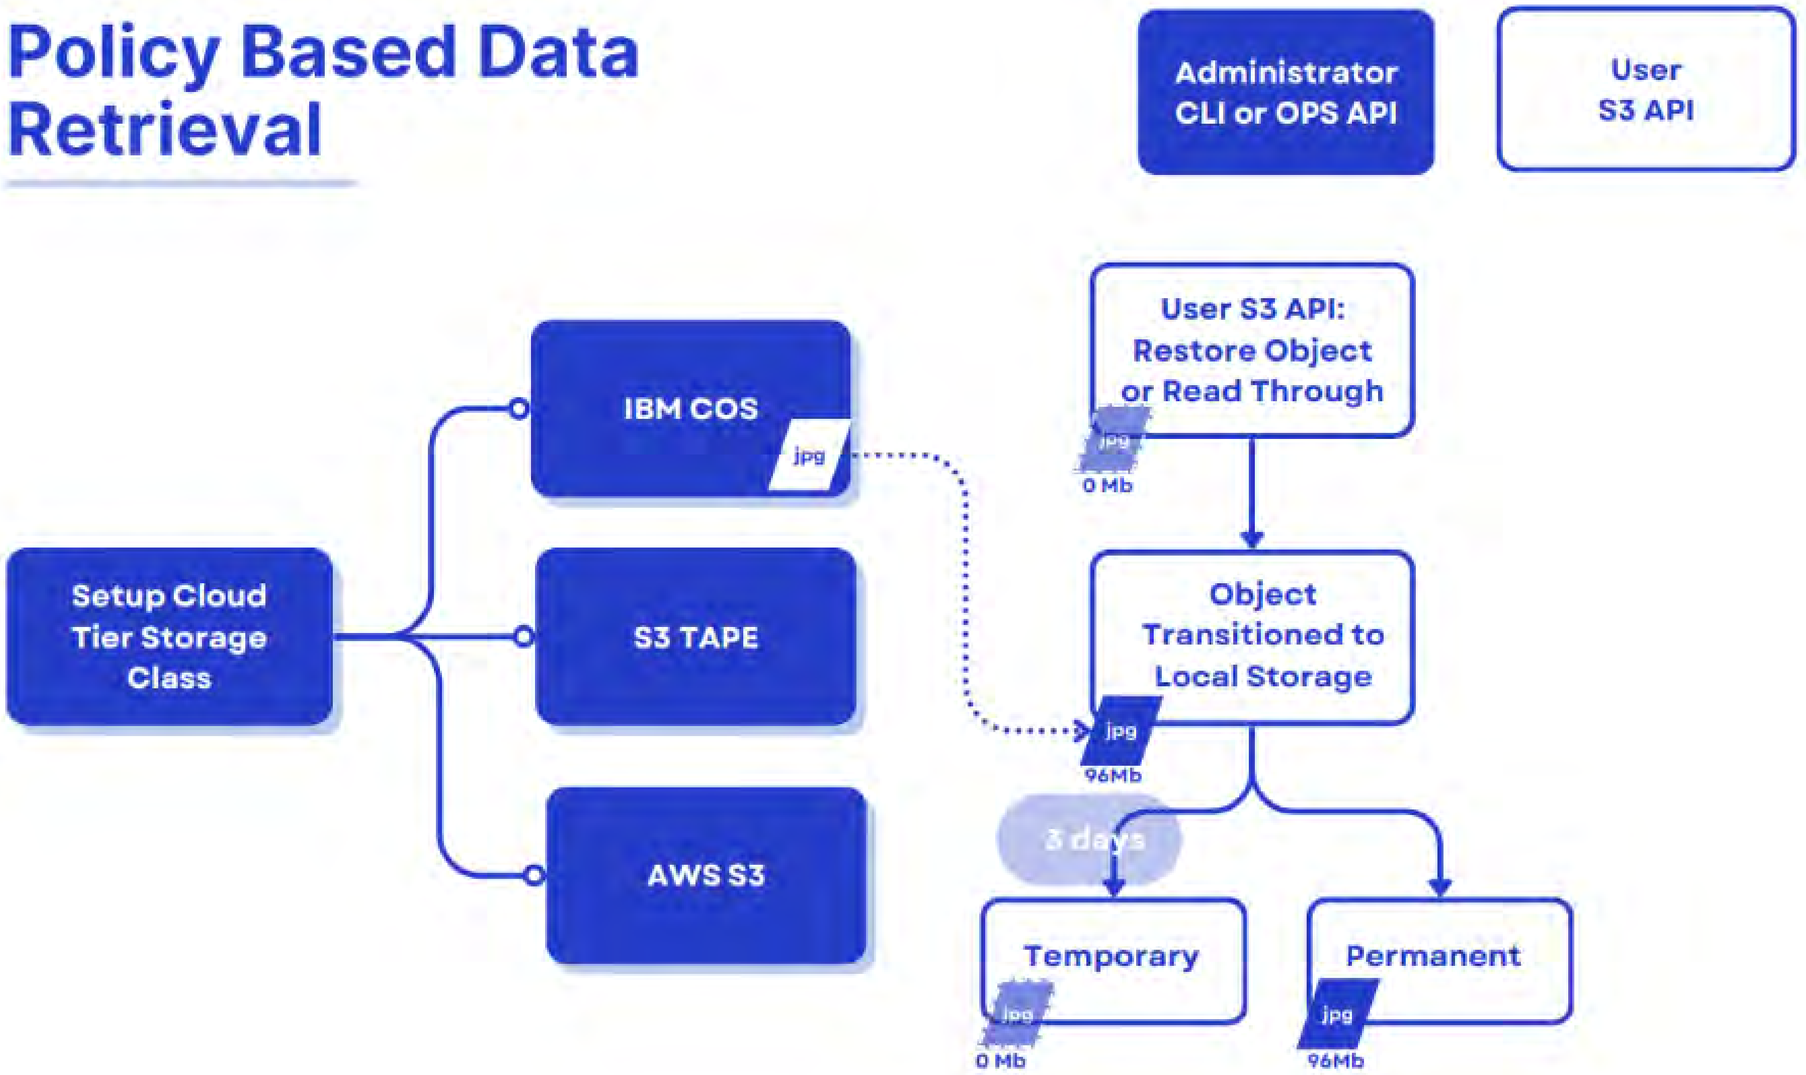

IBM Storage Ceph Object

Other New Features in 8.0

Enhanced Visibility into IBM Storage Ceph Object Multisite Replication with the New HTTP Sync Status Headers

Simplifying IBM Storage Ceph Object Deployments: Production-Ready Ceph Object Gateway with New Cephadm Features

Simplifying Object Multi-site Replication with Ceph Dashboard: Introducing the New 4-Step Wizard

IBM Storage Ceph Object Storage Overview Agenda

01 Introduction to Ceph
02 Introduction to Ceph Object Storage
03 Ceph Object Performance at Scale
04 Ceph Object new features in 8.0
05 Ceph Object ZeroTrust Lakehouse
06 Ceph Object Looking into the future
07 Question & Answers

Lakehouse

“In general, a modern cloud analytics architecture separates compute from storage, with key capabilities pushed to the durable storage component to address enterprise customer requirements, such as governance, security, metadata management, and performance.”

BigLake: BigQuery’s Evolution towards a Multi-Cloud Lakehouse

MODERN DATA LAKEHOUSE

IBM Storage Ceph

Redbooks
ibm.com/redbooks

presto cloudera teradata. APACHE Spark

snowflake watsonx.data Starburst

PostgreSQL redis

hadoop HDFS

kafka cassandra

IBM Storage Ceph

Curated zone
ICEBERG

Staging zone
ICEBERG

Raw zone

- Lifecycle
- IAM Policy
- Bucket Policy
- Notifications
- Object Lock
- Encryption
- Inventory

Core Challenges for Data Lakehouse Authorization

- Database-level access semantics , what users intuitively expect
- Storage-level enforcement , what zero-trust requires
- While maintaining direct paths, what performance and scalability demands

IBM Storage Ceph Object
Polaris + Ceph: Enforcing Table-Level Security via STS & IAM

Table-Aware Access Control

Polaris Catalog manages data at the table level, enabling granular policies that map directly to Ceph RGW's STS and IAM.

Per-user Granularity

This ensures that only authorized users (or applications) can access specific tables

Credential Vending

Polaris dynamically vends temporary session tokens to clients (e.g., Spark jobs), restricting storage access to only those tables (or objects) allowed by the catalog policy.

Scoped Session Policies

Guarantee that compute tasks can interact only with data they're permitted to use—no extra privileges are granted.

Storage-Layer Security

By enforcing security at the storage layer, organizations avoid misconfigurations or oversights that might occur solely at the compute level

Storage-Layer Security

Access checks happen directly in Ceph, leveraging STS, IAM, and policies—providing stronger defense against unauthorized data access or exfiltration

Catalog credential vending

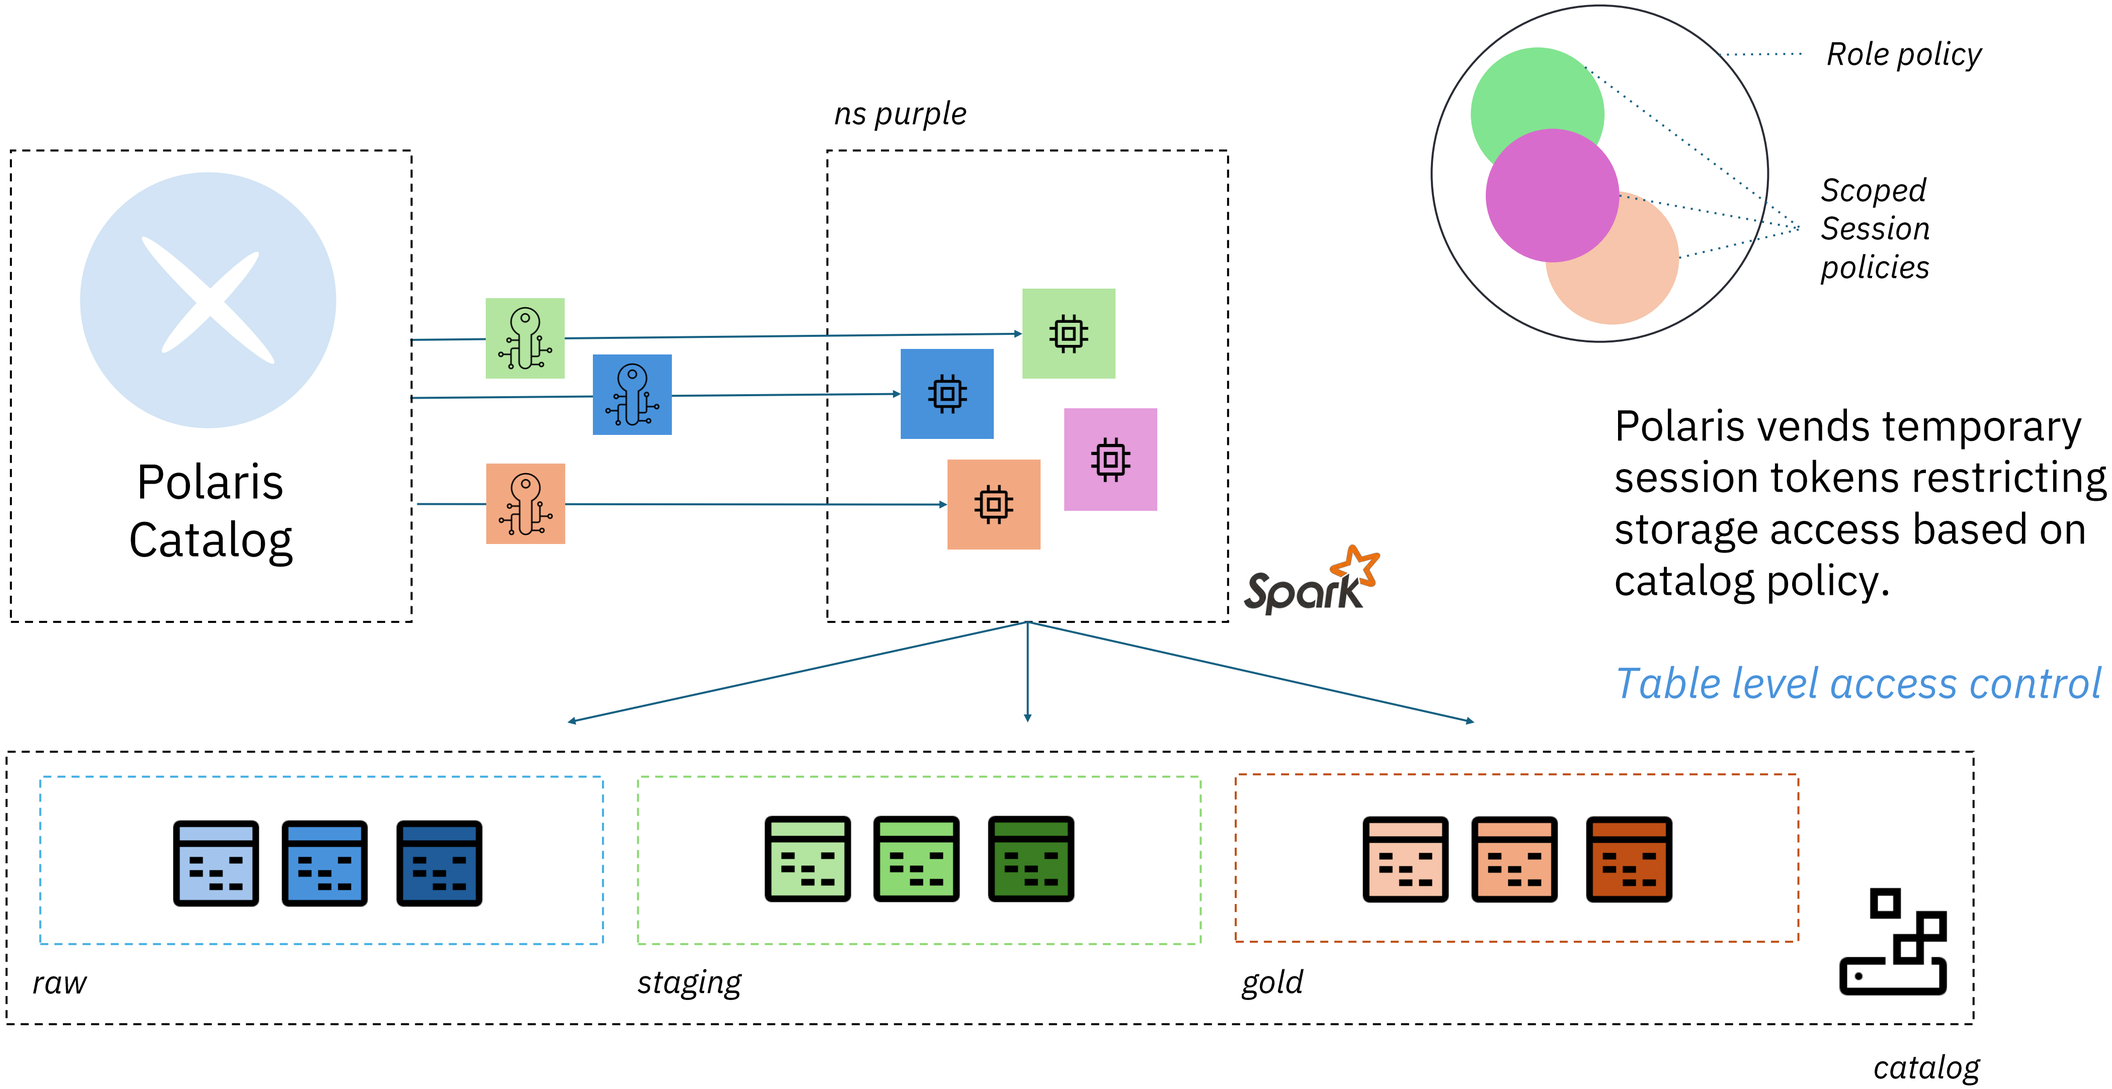

Polaris vends temporary session tokens restricting storage access based on catalog policy.

Table level access control

What is RGW D3N cache ?

- Datacenter-Data-Delivery Network (D3N) uses high-speed storage such as NVMe flash or DRAM to cache datasets on the access side
- D3N improves the performance of big-data jobs running in analysis clusters by speeding up recurring reads from the data lake
- The Rados Gateways act as cache servers for the back-end object store (OSDs), storing data locally for reuse.

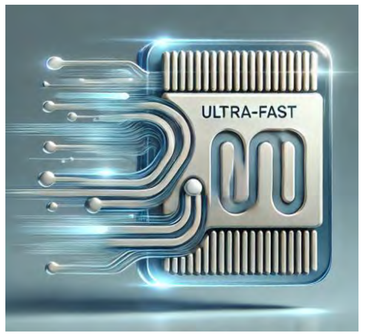

IBM Storage Ceph Object Storage Overview Agenda

01 Introduction to Ceph

02 Introduction to Ceph Object Storage

03 Ceph Object Performance at Scale

04 Ceph Object new features in 8.0

05 Ceph Object ZeroTrust LakeHouse

06 Ceph Object Looking into the future

07 Question & Answers

IBM Storage Ceph Object. Looking forward

Performance

➔Increasing the maximum number of objects per bucket(> 2 Billion)

➔D4N. Predictive caching layer for iterative Analytical/AI workloads

➔Embedded Arrow Flight engine, exposing a Flight endpoint for enhanced pushdown-based in-place analytics

Efficiency

➔Deduplication for large objects(>4MB)

➔Hardware-accelerated Compression with Intel QuickAssist

➔Policy Based Tiering to S3 Tape Glacier Endpoints

IBM Storage Ceph Object. Looking forward

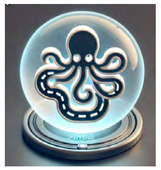

Observability

➔Multisite Replication Observability improvements

➔IBM Storage Ceph Object Open Telemetry Tracing

Multisite Replication

➔Same Zone Group and Cross Zone Group Replication

➔Batch Replication

IBM Storage Ceph Object Storage Overview Agenda

01 Introduction to Ceph

02 Introduction to Ceph Object Storage

03 Ceph Object Performance at Scale

04 Ceph Object new features in 8.0

05 Ceph Object ZeroTrust LakeHouse

06 Ceph Object Looking into the future

07 Question & Answers

IBM Storage Ceph Resources

IBM website

https://www.ibm.com/products/ceph

Seismic

IBM Storage Ceph page on Seismic

Box

IBM Storage Ceph enablement resources on Box

Slack

IBM Storage Ceph Slack channel #ceph-help

Reference guide

IBM Storage Ceph quick reference guide

Ceph community

IBM Storage Ceph open-source community website

Thank You

Q & A. Time for Questions?