REGIONAL SDC 25

BY Developers FOR Developers

Regional SDC Denver
April 30, 2025

On-Prem Object Storage: S3 Ecosystem and Its Role in Various Workloads

Kalyan Gunda & Ugur Kaynar

SNIA®

Limitations of traditional app development

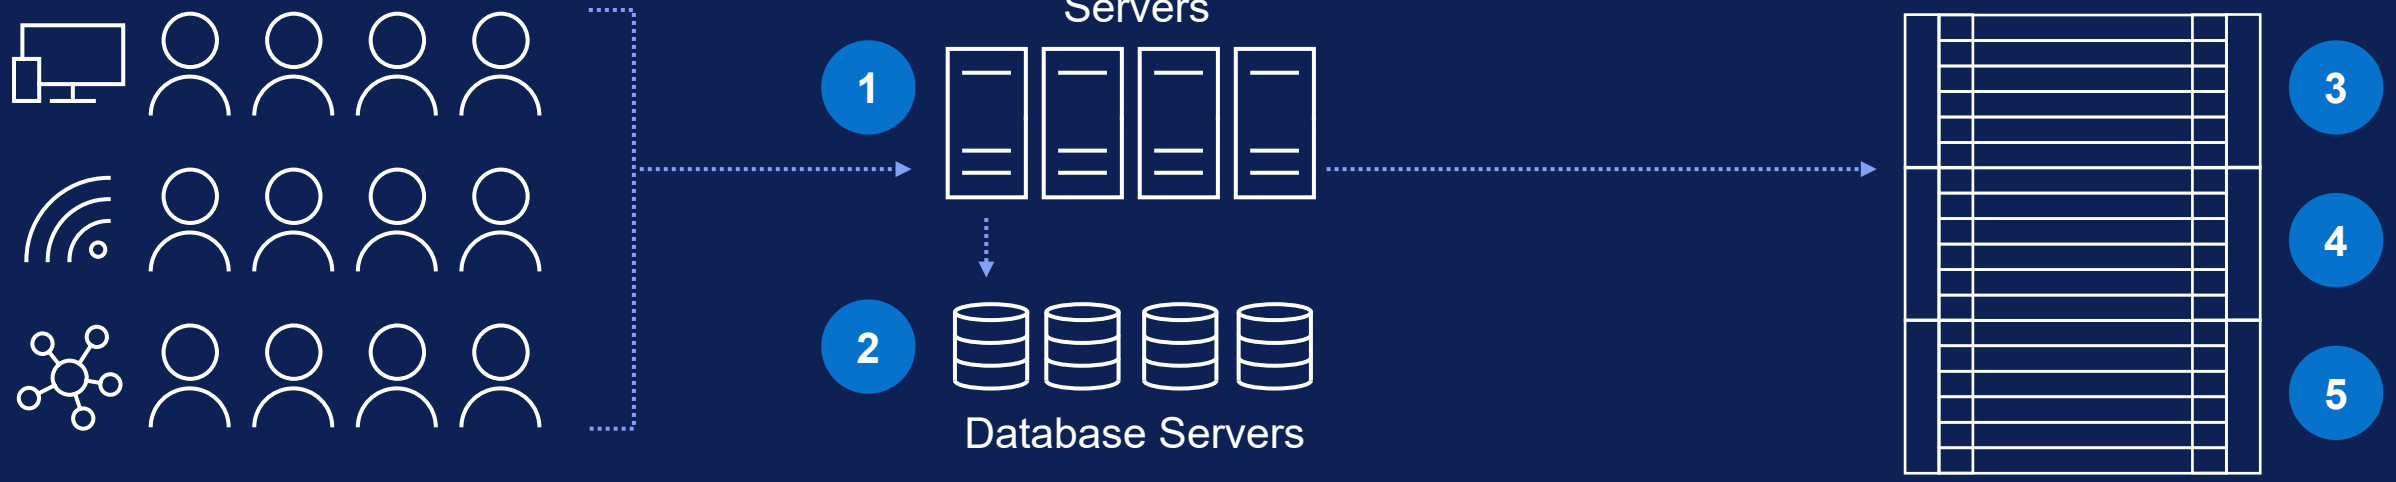

1. Can’t Provision storage directly from the app
2. Cant manage Data Lifecycle directly
3. Not developer friendly (multiple file shares must be managed) and storage must be provisioned
4. Can’t create data copies for DR and other purposes
5. Limited control on managing granular access control via RBAC policies

Today's Infrastructure and Application Administration

Infrastructure admin provides storage access to supports different workloads

- Infra admin provisions and allocates storage resources
- Application admin must coordinate Infra admin/API to provide and manage storage
- Protection of application is driven by storage replication or application specific toolchain.

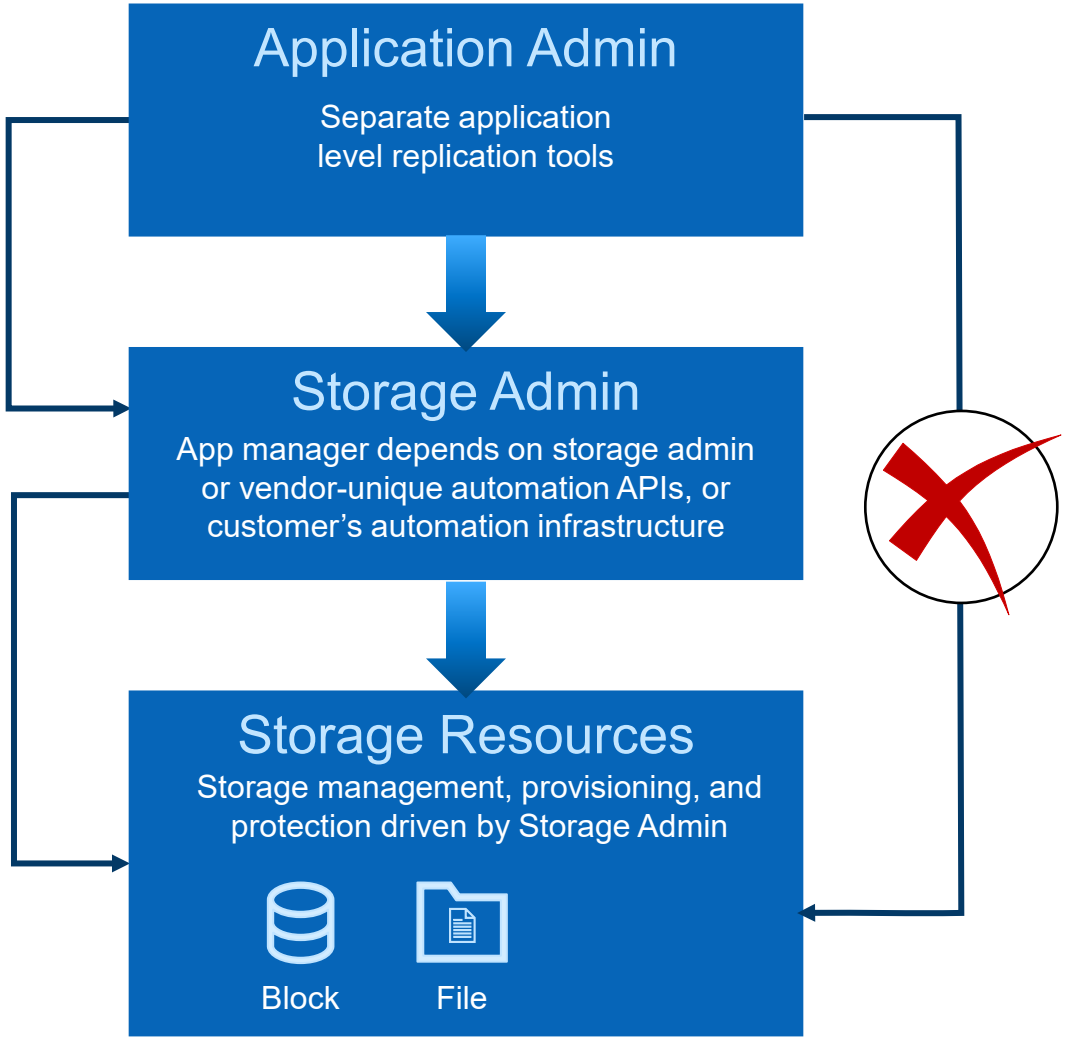

Block and File protocols do not provide standard, vendor independent interfaces for application admins to directly provision and manage storage and protection features

REGIONAL SDC

What is object storage?

TRANSACTION UNITS
Objects – Immutable data

PROTOCOLS
REST over HTTP/S

METADATA SUPPORT
User-defined metadata; System metadata

WELL SUITED FOR
Static Data; Billions;

BIGGEST STRENGTH
Application driven Storage management;
Scalability and distributed access

LIMITATIONS
Not suited for frequently changing transactional data

REGIONAL SDC

Modern Application S3 Storage Provisioning Workflow

Application admin self provisions S3 storage access and data service policies

- Infrastructure admin creates StorageClass
- Application admin deploying applications provides S3 storage access info StorageClass.
- Protection of application is driven by S3 policy and replication API's

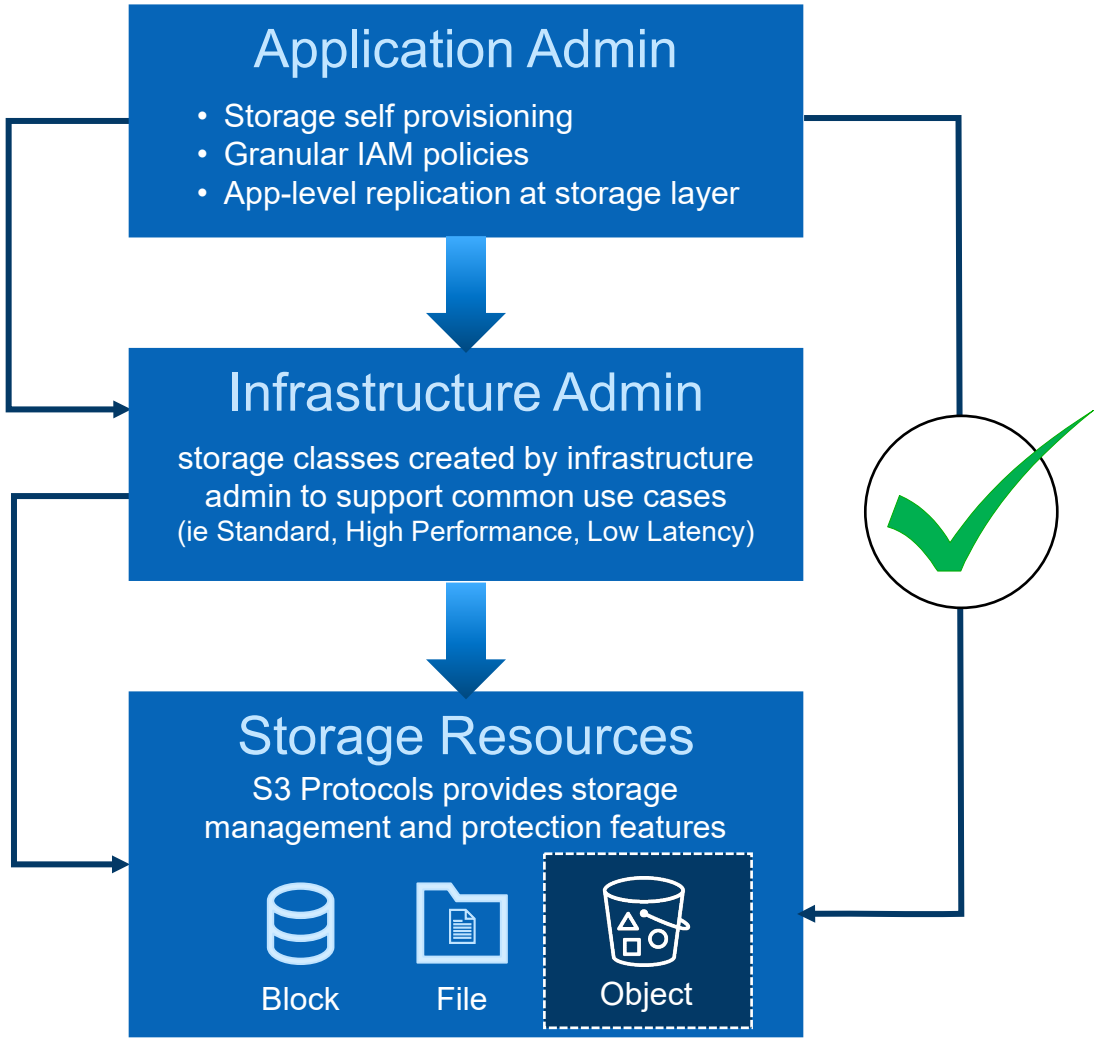

S3 compliant access provides standard storage management and protection API's

REGIONAL SDC

Object Storage | Traditional workloads

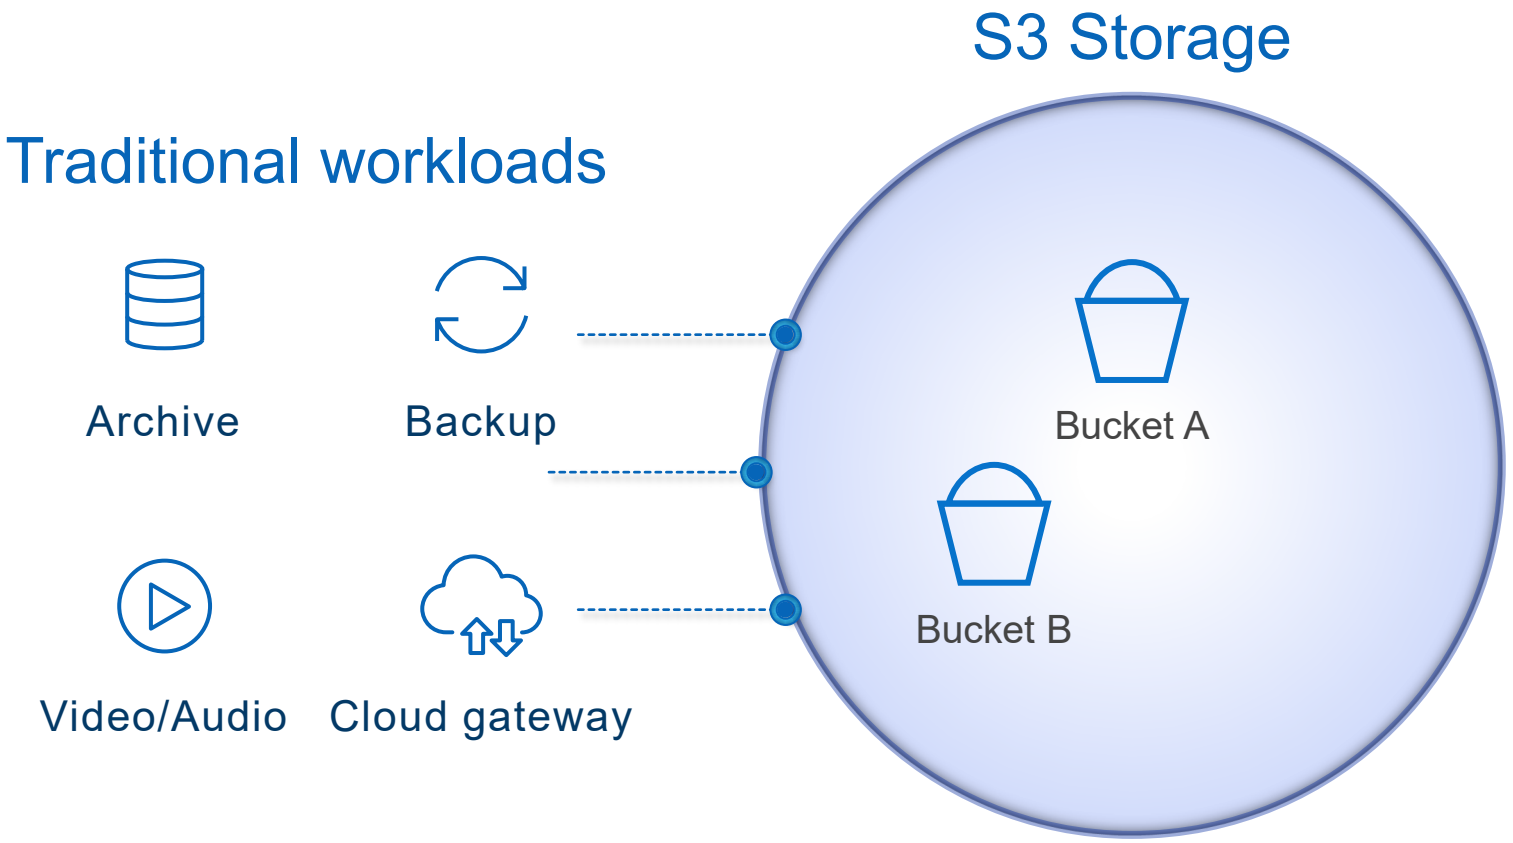

SCALABILITY | SIMPLICITY | METADATA | APIs

REGIONAL SDC

Object Storage | expanding to new workloads

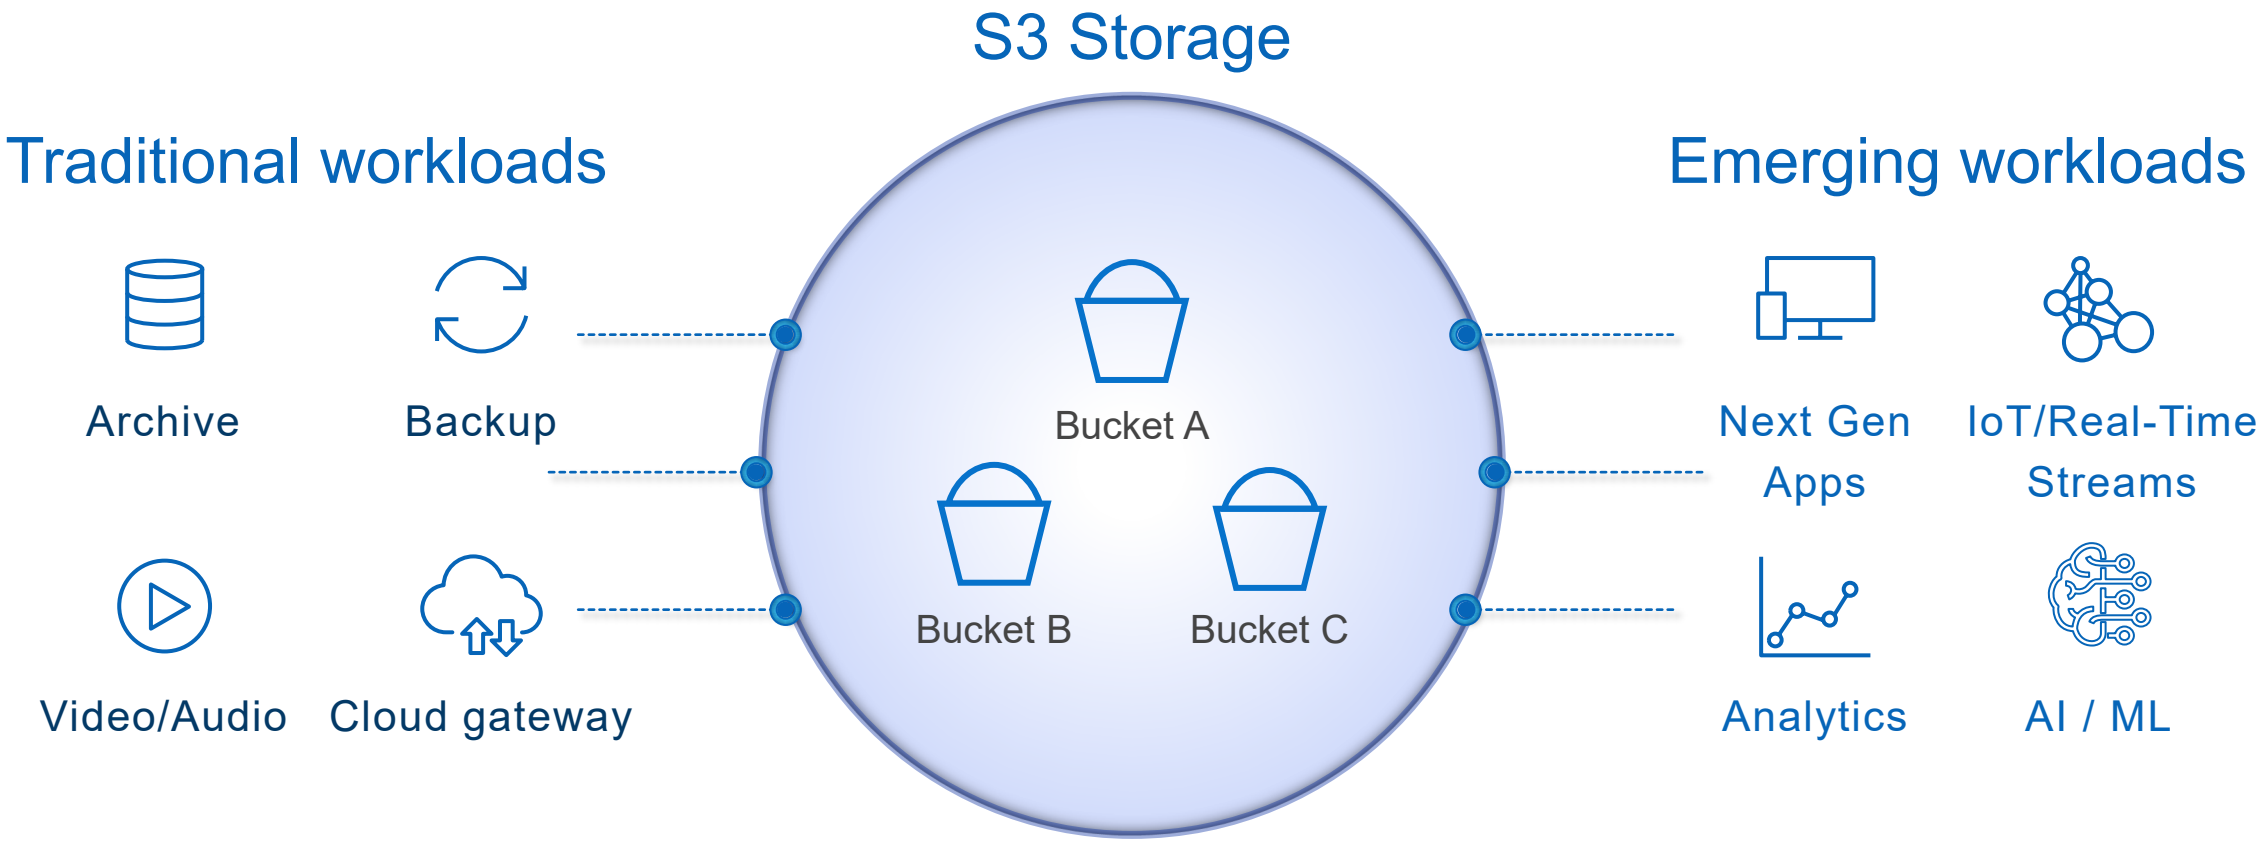

SCALABILITY | SIMPLICITY | METADATA | APIs

REGIONAL SDC

Object Storage Ecosystem

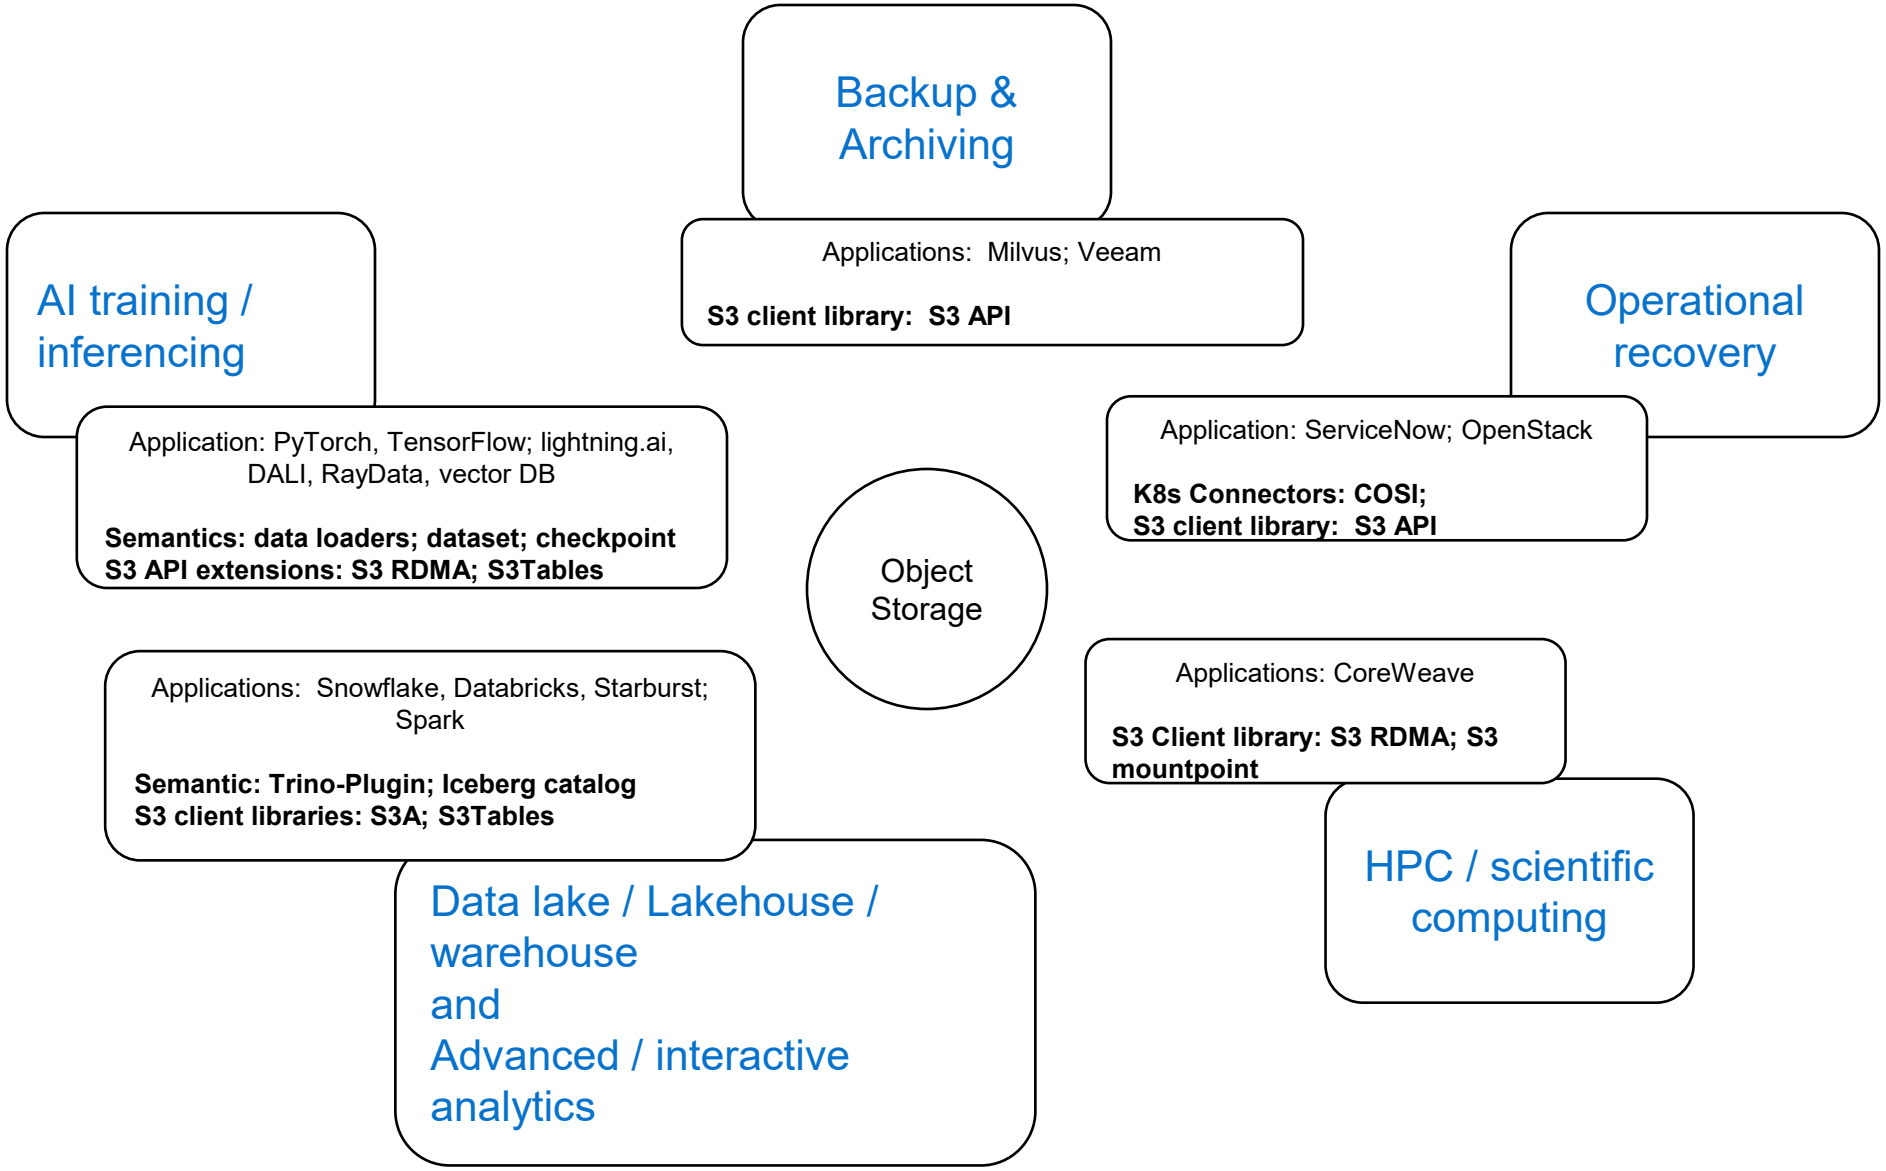

AI Is Not One Workload

Each phase of GenAI has different workloads requirements

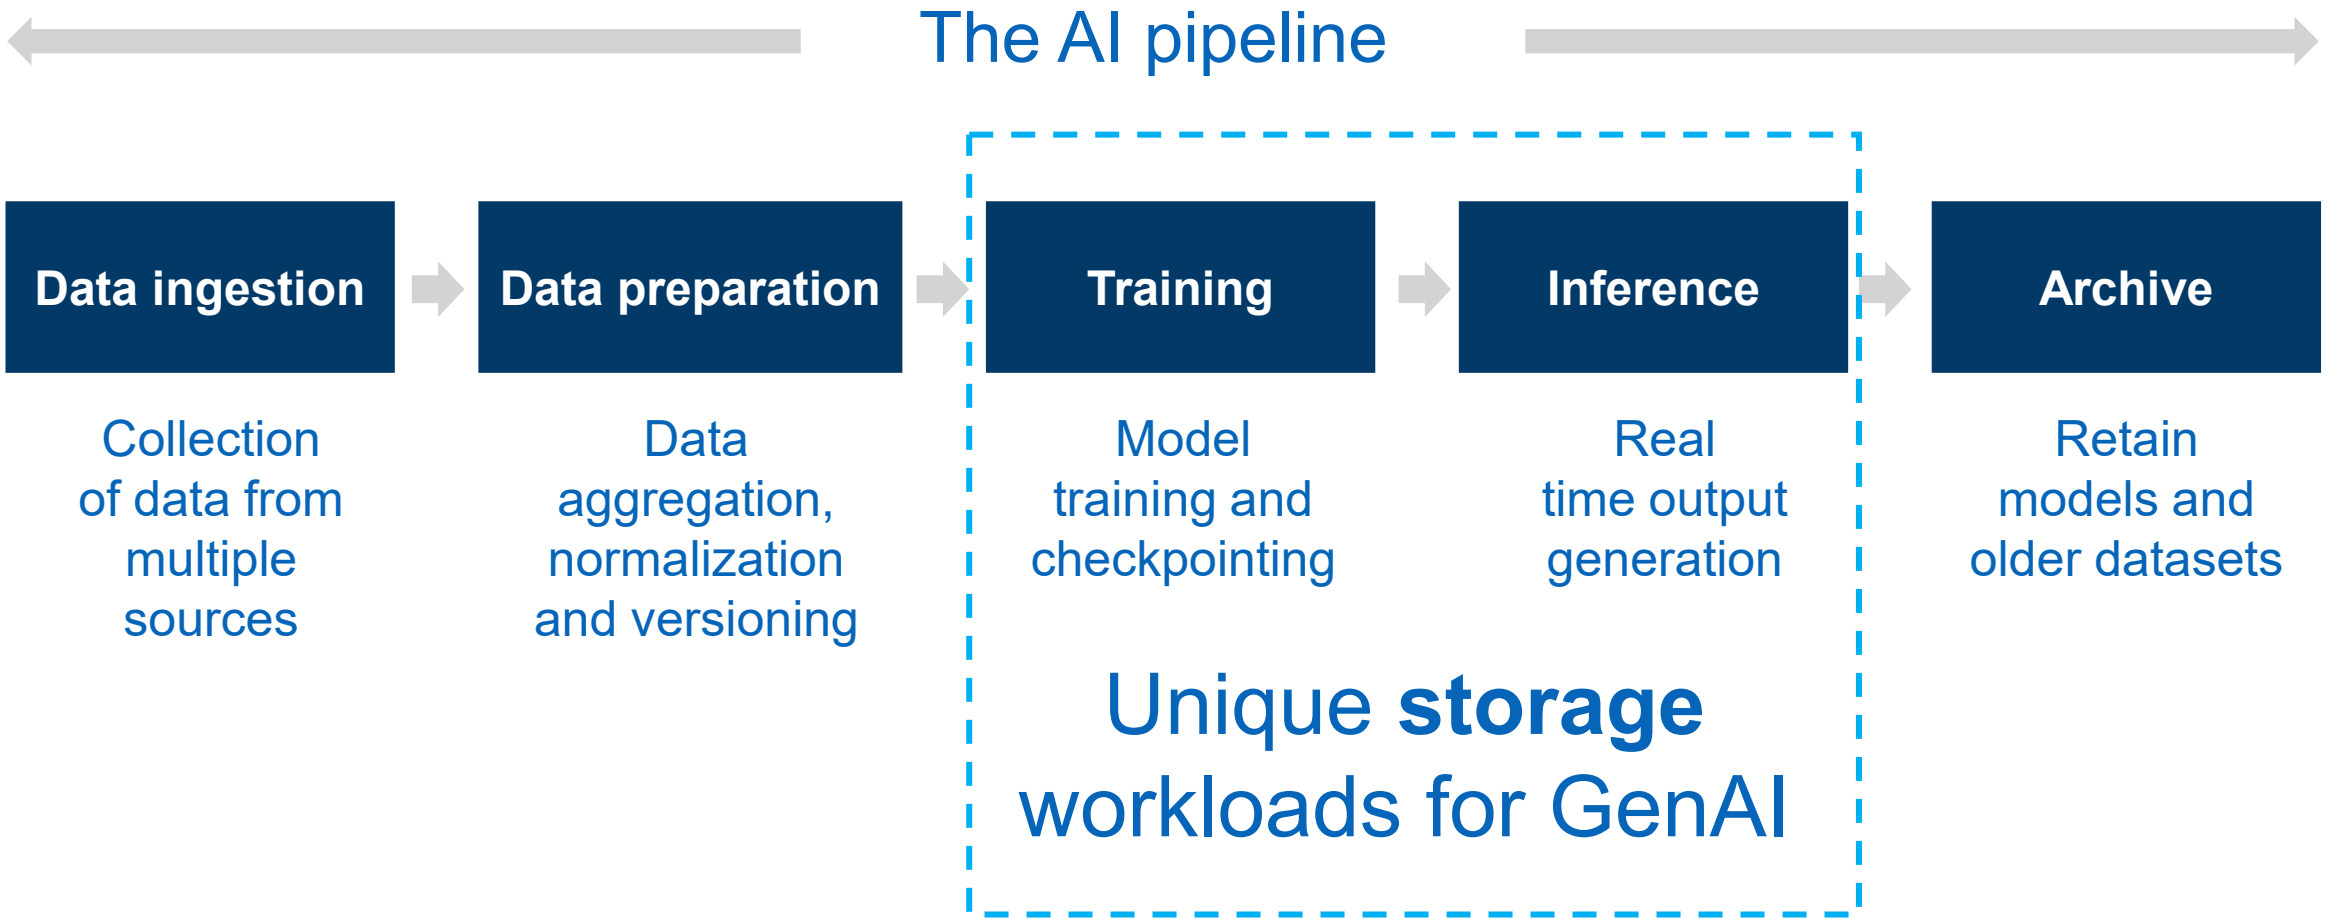

REGIONAL SDC

AI workloads are driving innovation across the entire data path of object storage

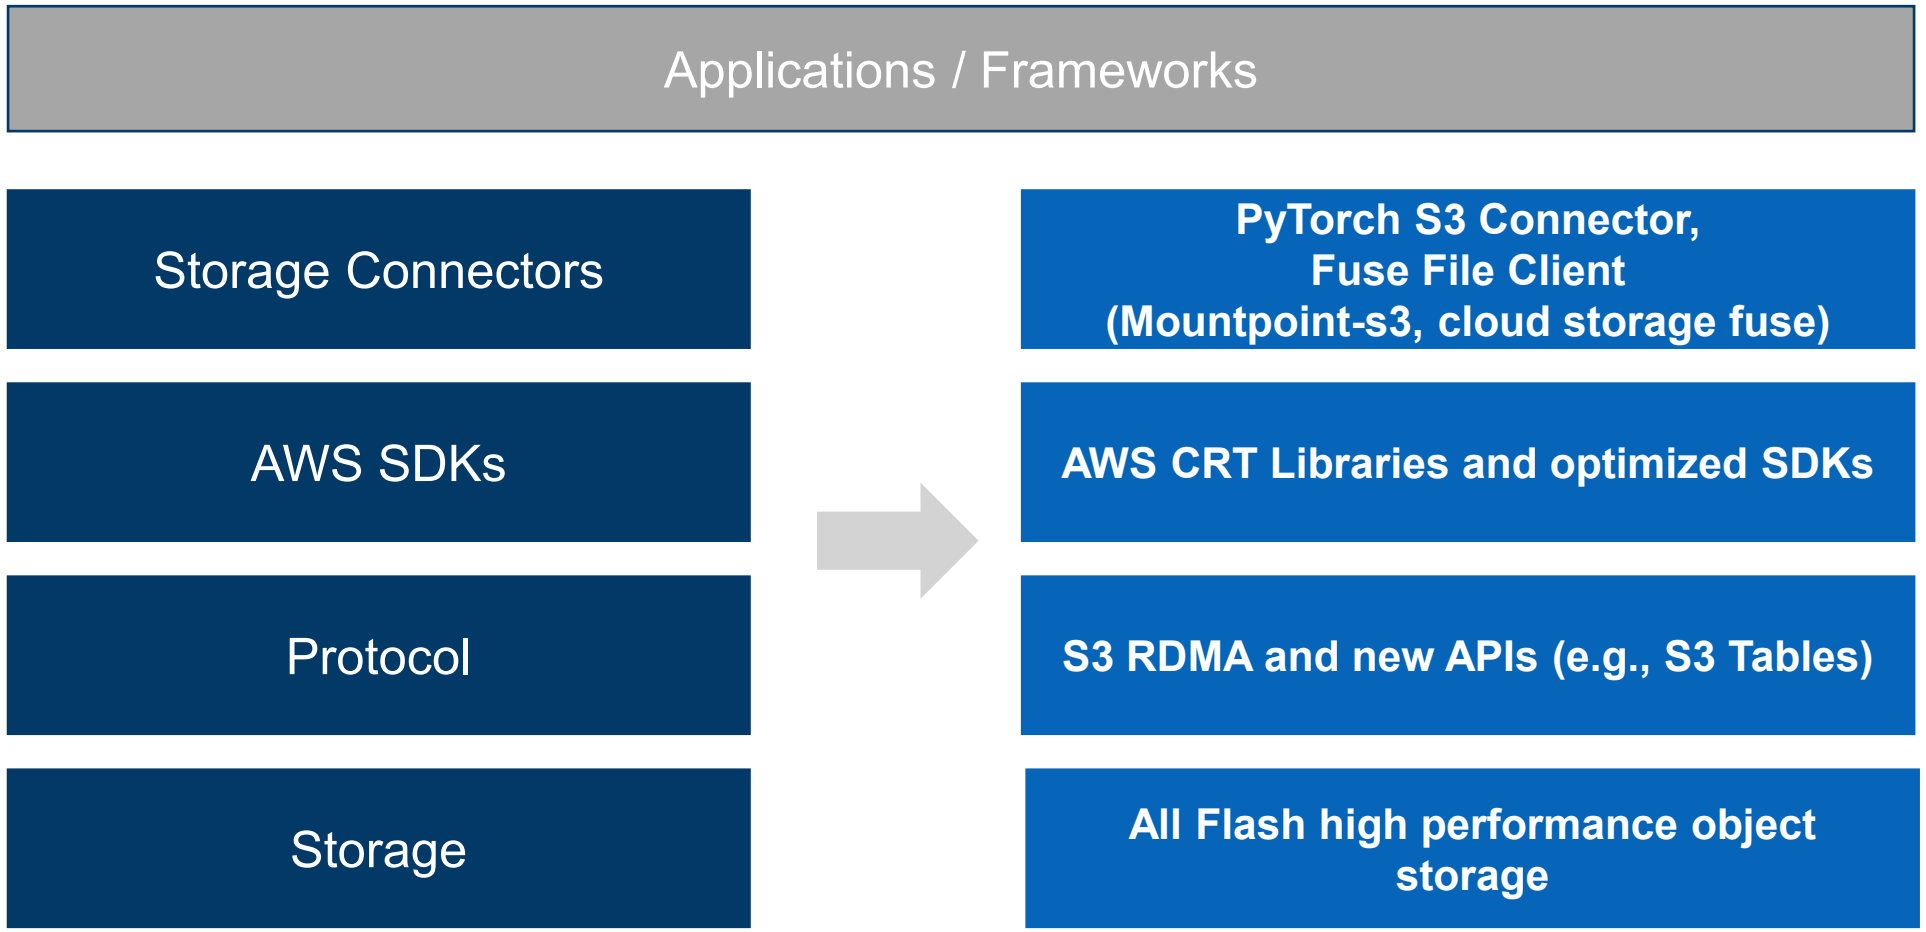

REGIONAL SDC

Architecture of Object Storage Systems

- The semantics of object storage allows;
  - Simplicity
    - Immutability (Write-Once-Read-Many).
      - Minimizes the need for locking mechanisms and reduces complexity in managing concurrent access.
      - Reduces the complexity of managing data consistency
      - Enables lockless architecture – extremely scalable
    - Atomic Operations
    - Relaxed Consistency – Eventual vs Strong Consistency
  - Scalability
    - Decoupled metadata: allows for efficient scaling
    - Flat Namespace: Efficient data distribution & access without bottlenecks
    - Replication: immutable objects drives simpler replication protocol (versioning enable bi-directional replication and global access)
  - Performant:
    - Applications can spawn multiple threads to perform parallel operations.
    - Non blocking IO: Allowing threads to initiate multiple requests without waiting for each to complete
    - Stateless Protocol: Rest APIs (HTTP/s) does not require maintaining a persistent connection between the client and server and reduce the overhead.
    - Immutability enable aggressive caching
    - Relaxed consistency models allow for certain operations to be reordered or delayed, which can improve performance by reducing the need for strict synchronization

REGIONAL SDC