



STORAGE DEVELOPER CONFERENCE

SNIA ■ SANTA CLARA, 2015

Bridging On-premise File System and Cloud Storage

Pankaj Datta
EMC – Isilon Storage Division

Acknowledgements

To the amazing cloud ops team...

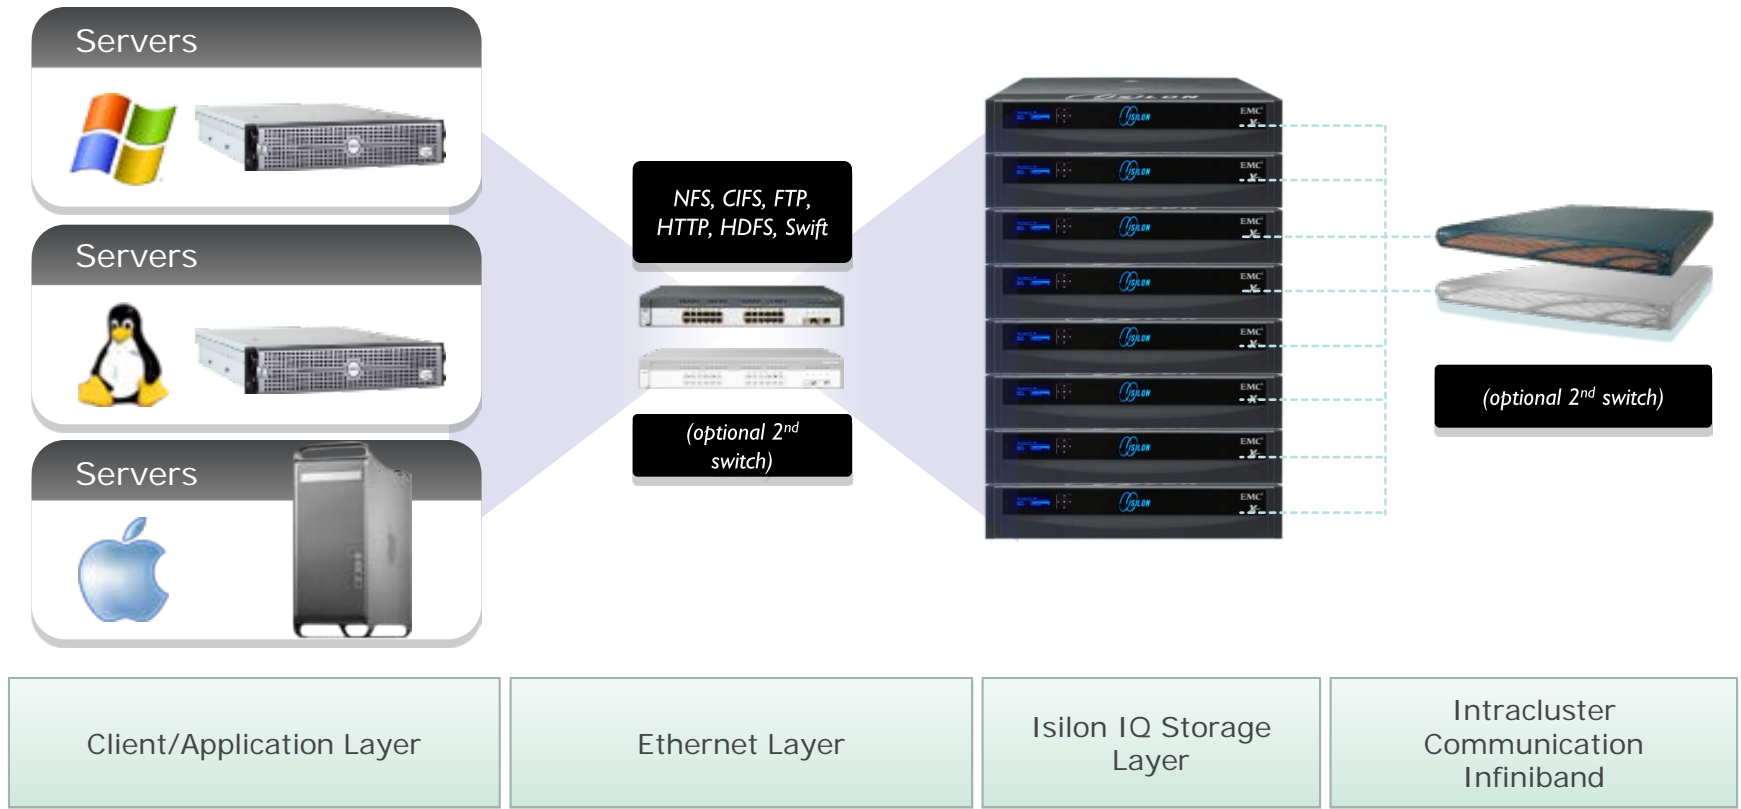
Outline

- ❑ Motivation
- ❑ OneFS Overview
- ❑ Architecture
- ❑ Features
- ❑ Integration
- ❑ Q&A

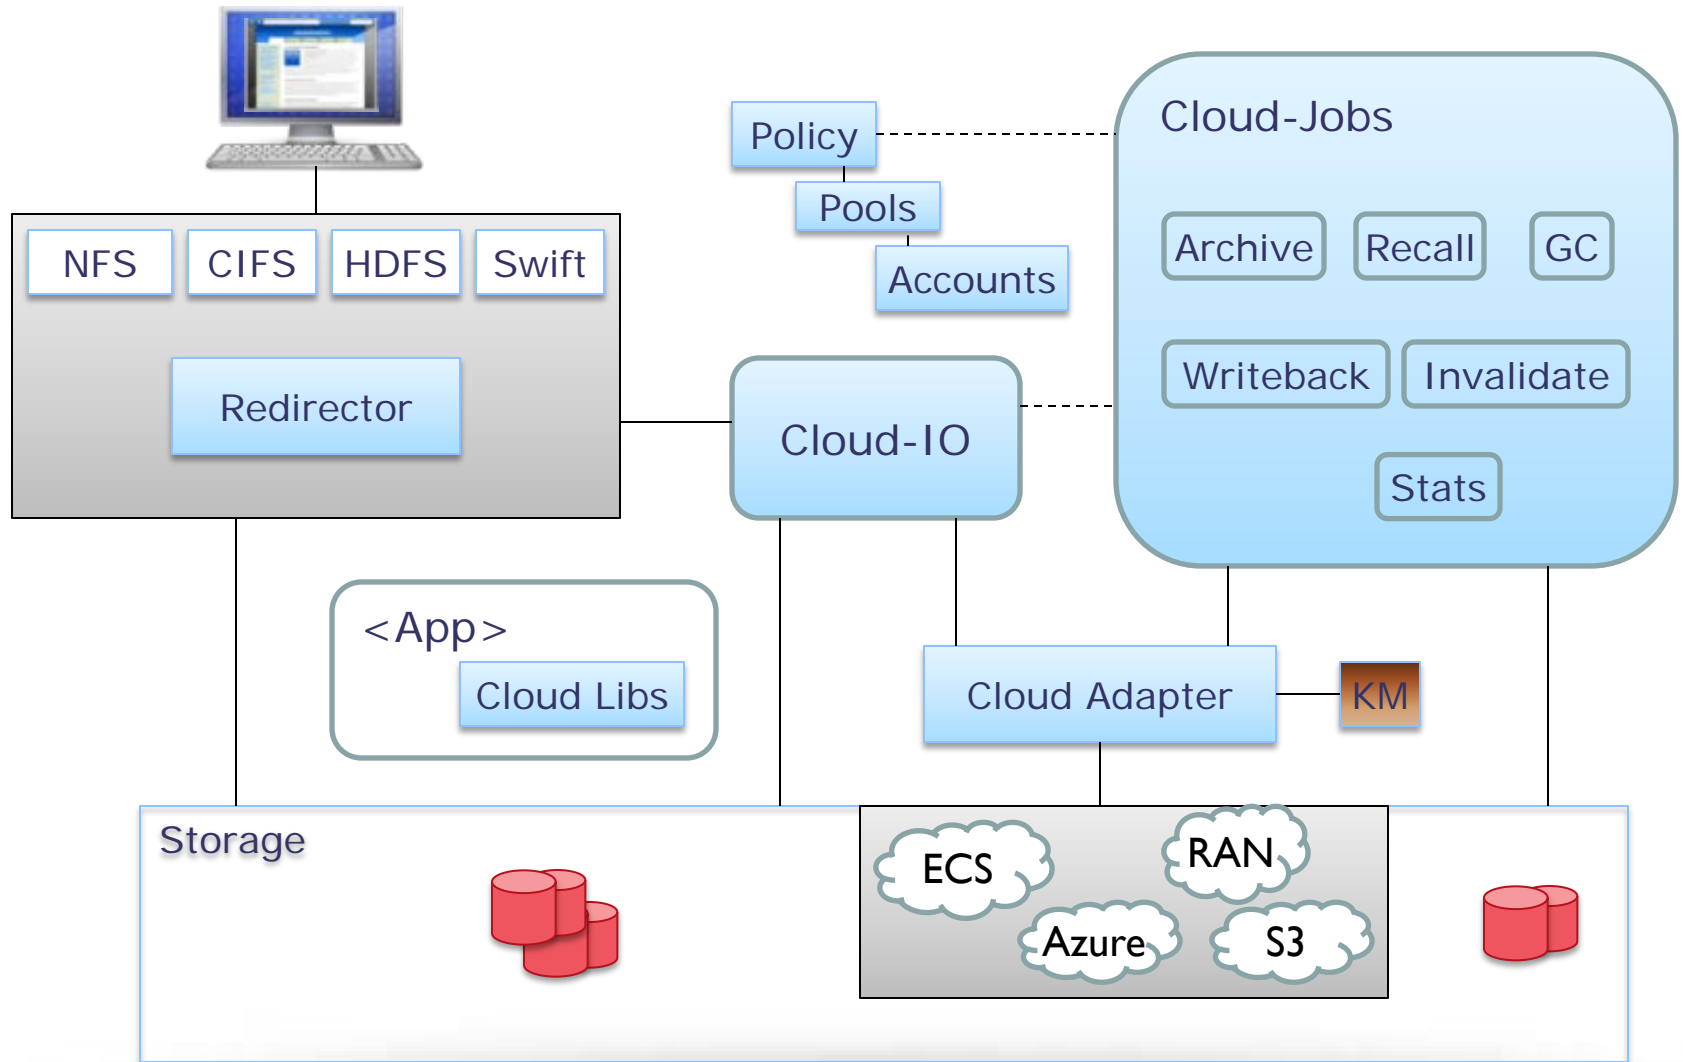
Motivation

- ❑ Cloud is always available, anywhere
- ❑ Use it as a low cost, virtually unlimited, storage tier for data not needed immediately
- ❑ Use higher performance, local OneFS NAS, tiers otherwise

OneFS Overview



Architecture: Cloudpools



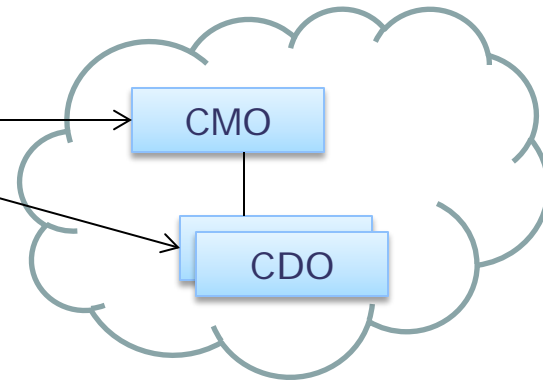
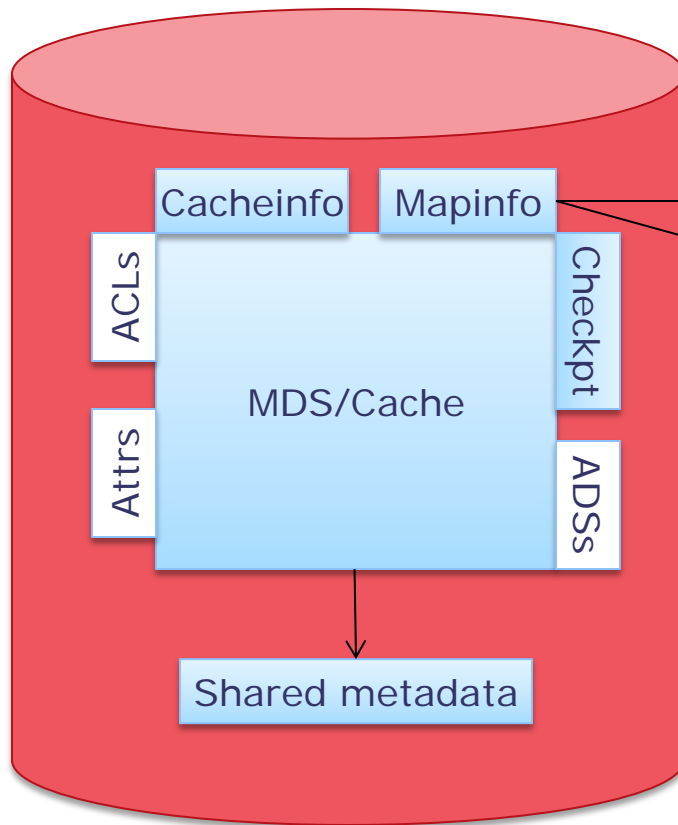
Features: Policy

- ❑ Associated with a pool of cloud accounts
- ❑ Various criteria to select *which* content gets stored (archived) to the cloud:
 - ❑ Path, Changed-time, Size, Type, etc
- ❑ Various actions to select *how* content gets handled:
 - ❑ Compression, Encryption, Caching, Retention

Features: Archive

- ❑ File metadata stays on OneFS cluster
- ❑ File data chunked and stored on the cloud:
 - ❑ 1 Cloud Metadata Object (CMO)
 - ❑ N Cloud Data Object (CDO)
- ❑ Multiple concurrent file archiving
- ❑ File becomes 'Smartlink' after successful archive
- ❑ A new Smartlink has a tiny OneFS footprint

Features: What is a Smartlink?



- ❑ Special OneFS file
 - ❑ Pointer like
 - ❑ Data bearing
 - ❑ Versioned

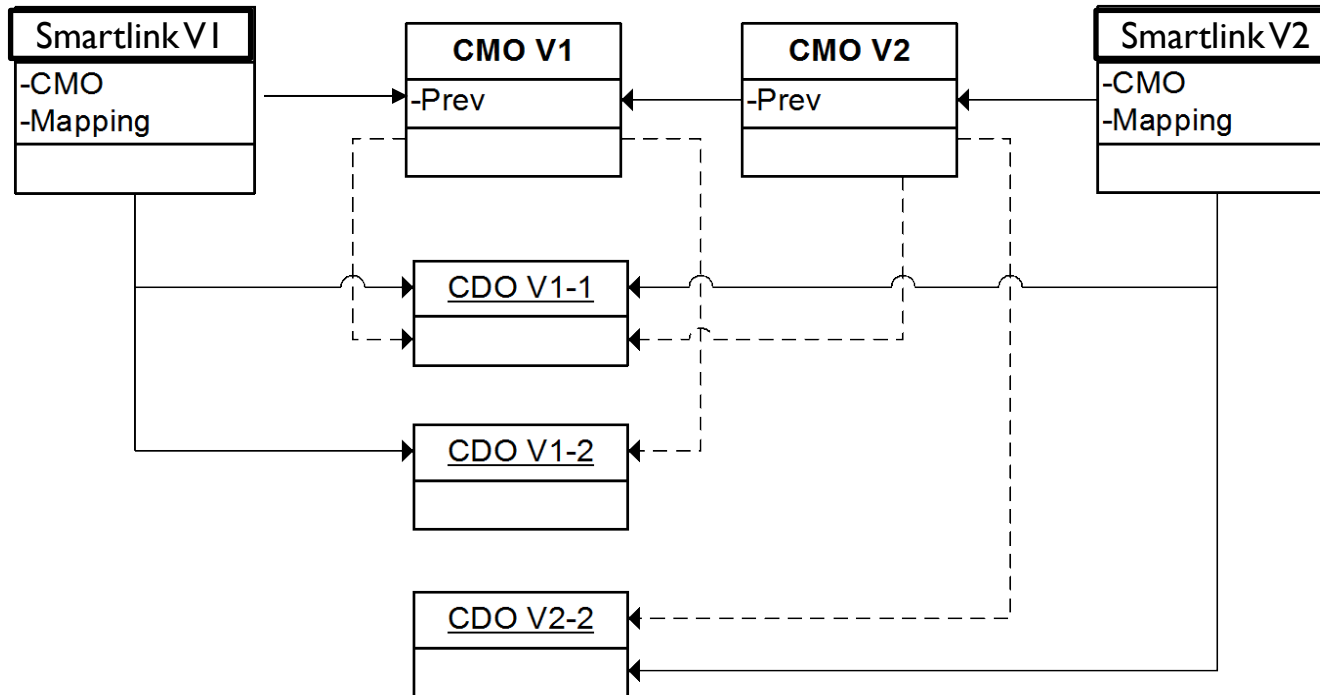
Features: Caching

- ❑ Smartlink access can trigger caching
- ❑ Cached data stored inside Smartlink file
- ❑ Caching metadata (Cacheinfo) in file ADS
- ❑ Cache maintains state per file region
 - ❑ Uncached, Cached, Dirty, Syncing
- ❑ Caching behavior controlled by policy
- ❑ Fully cached Smartlink != regular file

Features: [Total!] Recall

- ❑ Uncached data in Smartlink fetched from cloud
- ❑ Smartlink becomes regular file after successful recall
- ❑ Multiple concurrent file recall
- ❑ Unreferenced cloud objects deleted later

Features: Cloud Data Model



Features: Writeback

- ❑ Periodic job execution
- ❑ Dirty cache regions written to cloud
- ❑ Perform Smartlink cacheinfo and mapinfo update

Features: Invalidation

- ❑ Periodic job execution
- ❑ Addresses the need to ‘shrink’ Smartlink OneFS footprint
- ❑ Smartlinks with ‘Cached’ regions are cleared of local content
- ❑ LRU semantics

Features: Security

- ❑ Policy configurable
- ❑ AES 256 encryption
- ❑ Encrypt data before storing to cloud
- ❑ Encryption keys from OneFS Key Manager (KM)
- ❑ Block based encryption

Features: Garbage Collection (GC)

- ❑ Any CMO not referenced on OneFS is an orphan
- ❑ An orphan has a Date of Death (DOD) dictated by retention times
- ❑ [Cloud] GC orphan when its DOD expires
- ❑ Asynchronous, schedule driven, collection

Features: Cloud Libs

- ❑ Provide APIs to handle Smartlinks
 - ❑ Low level: detect, read, write, truncate, delete, etc
 - ❑ High Level: app specific
- ❑ Integrate with OneFS apps (snapshot, backup, replication, etc)

Features: Stats

- ❑ Collect statistics on cloud access from the cluster
- ❑ Inform user or throttle operations when thresholds are exceeded
- ❑ Support per minute, hour, day, month granularity of presentation

Features: Metering

- ❑ Intent to understand how much cloud storage is being used
- ❑ Usage data fetched for cloud accounts used by a cluster
- ❑ Reported to Isilon support central

FS Integration

A look at Smartlink and OneFS integration

FS Integration: Protocol support

- ❑ Smartlink accessible over NFS, CIFS, HDFS, and Swift just like a regular file
- ❑ Redirect handling of Smartlink to cloudpools
- ❑ Return accessed blocks from local storage on cache-hit
- ❑ Fetch accessed blocks from cloud on cache-miss

FS Integration: Snapshots

- ❑ Integrated into OneFS snapshots
- ❑ Smartlink snap represents point-in-time data from OneFS *and* Cloud storage
- ❑ Support cloud COW by using provider specific operations
- ❑ Snap-diff support for Smartlink

FS Integration: Backup

- ❑ Integrate into OneFS NDMP backup – no change to backup application
- ❑ Smartlink shared metadata (accounts, pool, policy, keys) backed (and restored) as well
- ❑ Configurable backup retention time
- ❑ Option to backup Smartlink as regular file

FS Integration: Replication

- ❑ Integrated into OneFS SyncIQ – bidirectional content synchronization across clusters
- ❑ Smartlink replicated to target cluster in a fraction of time on average
- ❑ Smartlink replication typically takes up miniscule network bandwidth – cluster to cluster (or cloud)
- ❑ Smartlink can be replicated as regular file and any changes to either synced across clusters

FS Integration: Changelists

- ❑ OneFS changelists allow customer identification of changed files across two snapshots
- ❑ Prevent false positives for Smartlinks
 - ❑ Cache read
 - ❑ Cache invalidation

Q & A