DataDomain Cloud Tier: Backup here, backup there, deduplication everywhere!

Abhinav Duggal

Fani Jenkins

Philip Shilane

Ramprasad Chinthekindi

Ritesh Shah

Mahesh Kamat

Dell EMC

DELL EMC

Data Domain Appliance

- Purpose-built backup appliances
  - Backup and restore
  - Supports deduplication (typically 20x) & local compression (typically 2x)
  - 4TB to 1PB physical capacity appliances
- Data Domain filesystem
  - Log structured filesystem
  - Stream Informed Segment layout architecture [Zhu 08]
  - Mark and sweep based garbage collection [Douglis 17]
  - Supports traditional backup workloads (sequential) and modern (random access) backup workloads [Allu 18]

Data Domain Cloud Tier

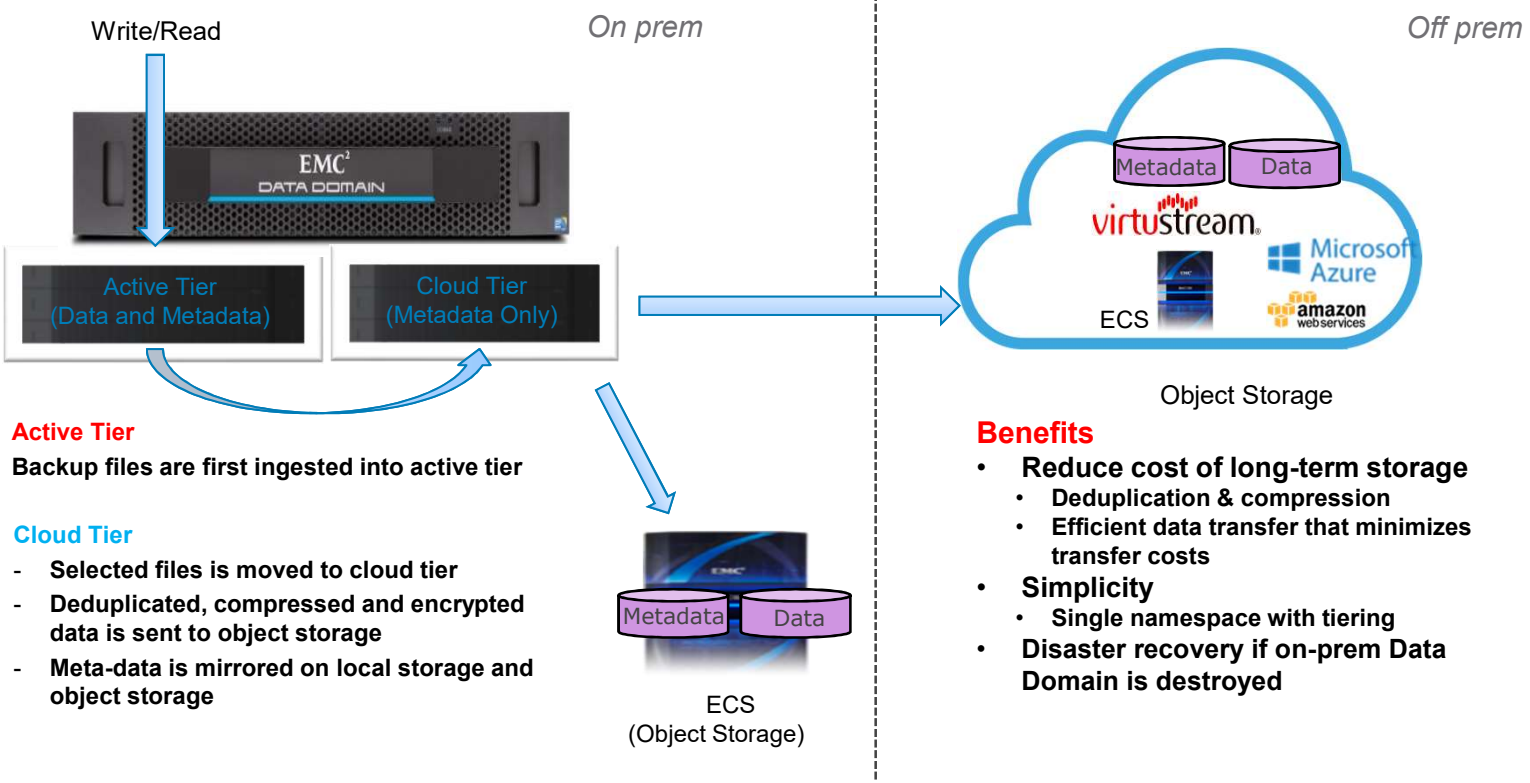

Active Tier

Backup files are first ingested into active tier

Cloud Tier

- **Selected files is moved to cloud tier**
- **Deduplicated, compressed and encrypted data is sent to object storage**
- **Meta-data is mirrored on local storage and object storage**

Benefits

- **Reduce cost of long-term storage**
  - **Deduplication & compression**
  - **Efficient data transfer that minimizes transfer costs**
- **Simplicity**
  - **Single namespace with tiering**
- **Disaster recovery if on-prem Data Domain is destroyed**

DELL EMC

File Representation in Data Domain Filesystem

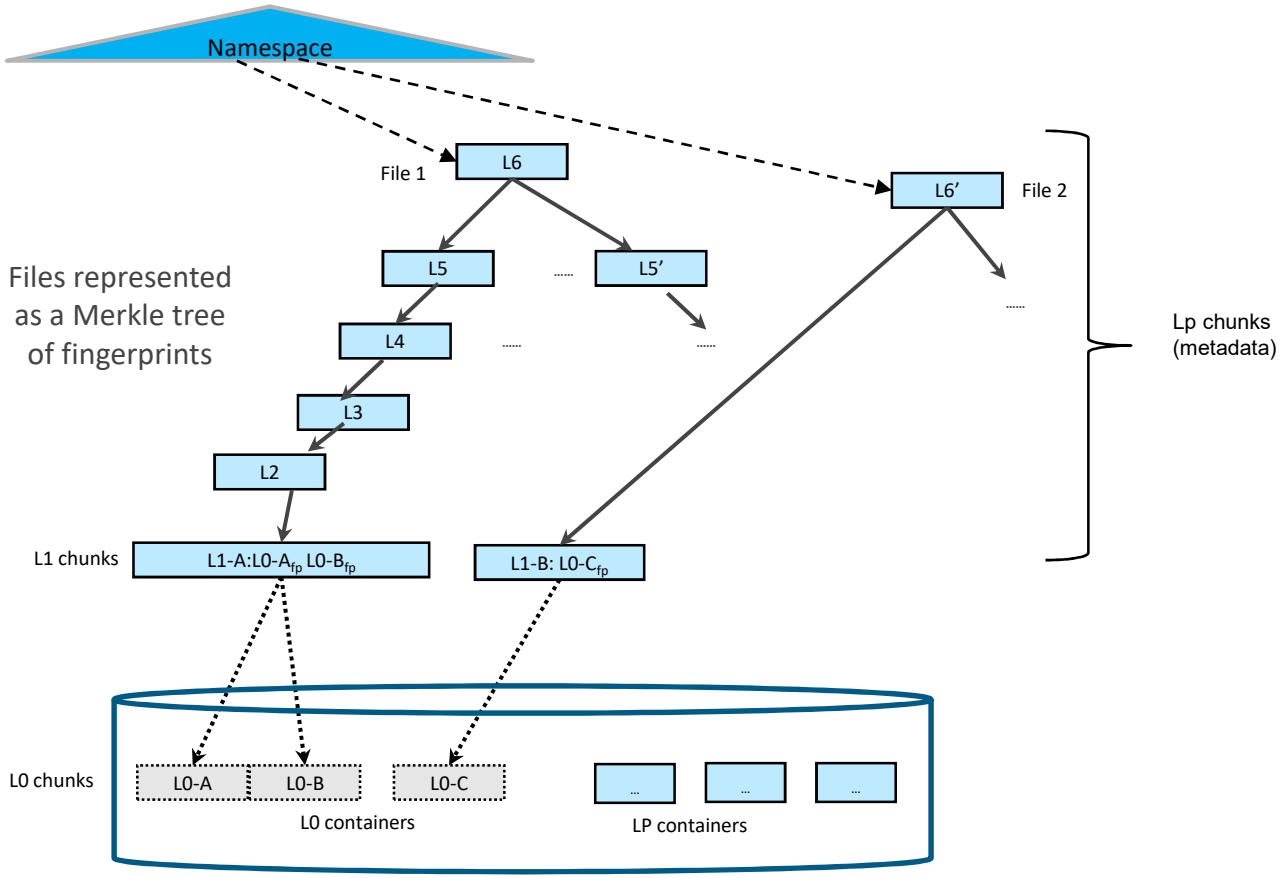

DELL EMC

Active Tier and Cloud Tier Architecture

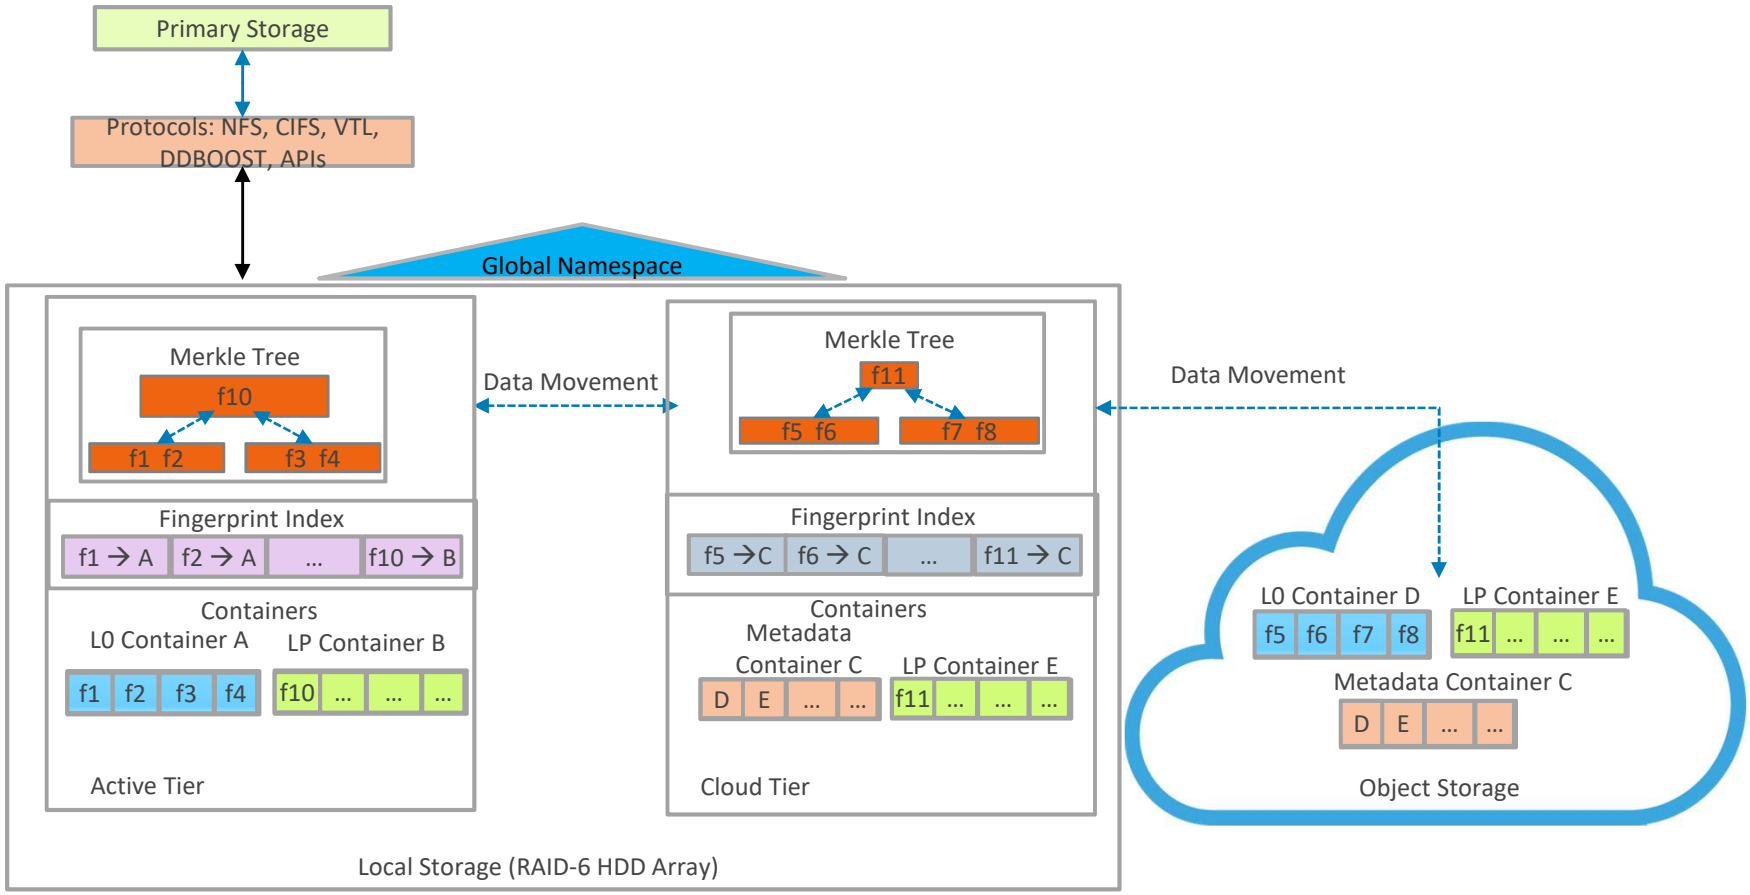

DELL EMC

Challenges in Deduplication System

- Depth-first enumeration of Merkle trees
  - Random IO to load every LP chunk
  - Repeated work due to deduplication
- Marking chunks live in a compact in-memory data structure
  - Build compact perfect hash vector (Fp PHV) which uses perfect hash functions to map fingerprints to unique position in bit vector (2.8 bits per fingerprint)

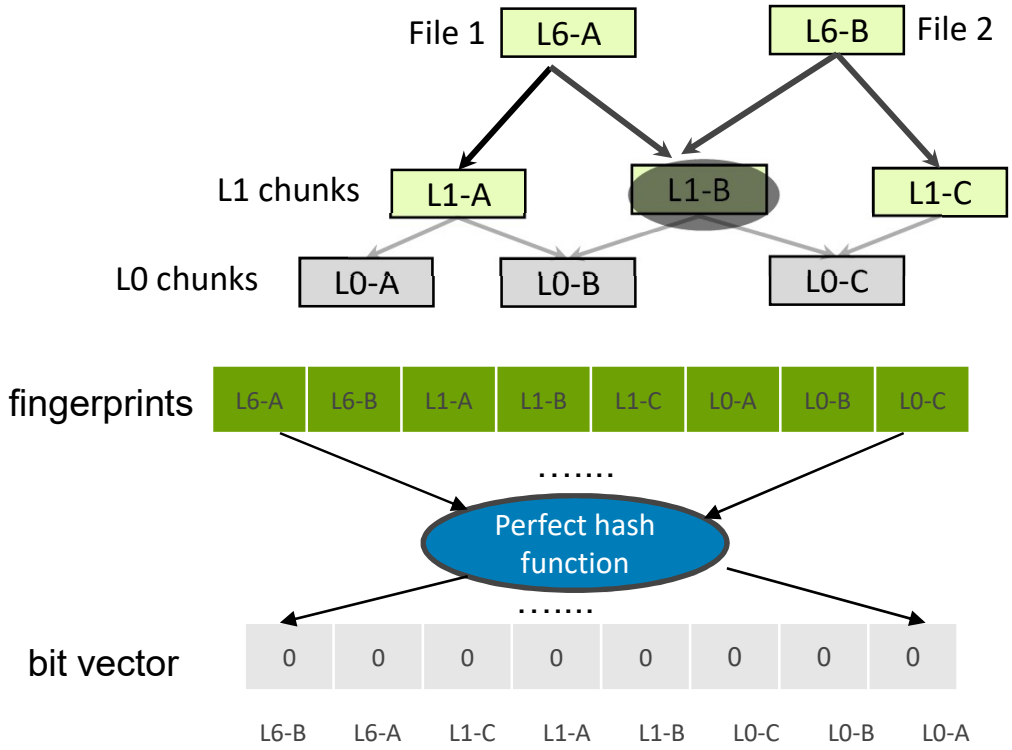

Requirements for Cloud Tier

In future, how much space will be freed if we move data to cloud tier?

Will our customers be able to seed large amount of data to cloud tier?

Will they be able to send incremental data efficiently after seeding?

How can we efficiently clean cloud tier when data dies?

Just “perfect hash” smart, “physical scan” fast,

and the cloud tier future will forever last.

Physical Scanning

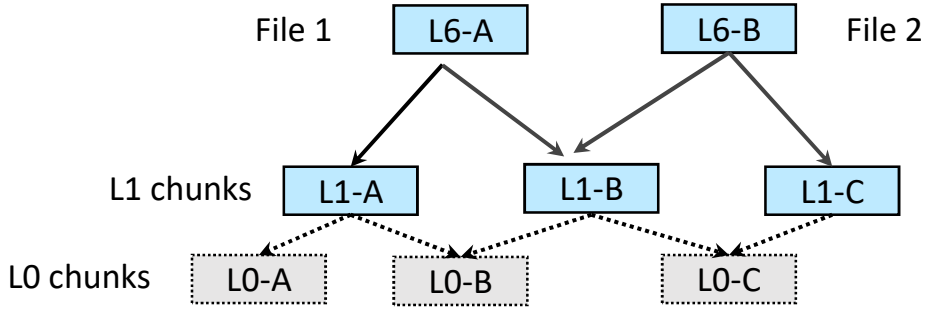

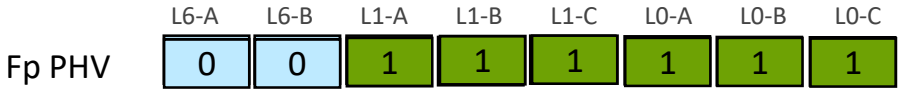

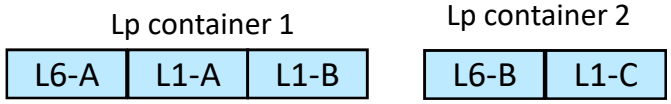

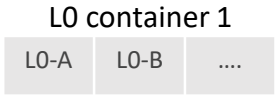

DELL EMC

How much space will be freed if we move data to cloud tier?

- Build perfect hash vector(PHV) by walking fingerprint index of active tier
- Physically scan candidate files meta-data chunks and mark them live (File 1 and File 2)
- Scan remaining files and unmark live chunks (File 3)
- Scan container set and accumulate chunk sizes of chunks still marked live

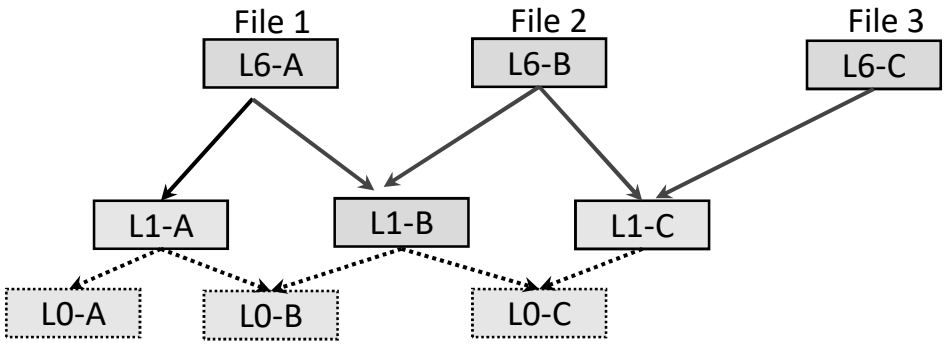

L6-A	L6-B	L6-C	L1-A	L1-B	L1-C	L0-A	L0-B	L0-C
1	1	0	1	1	0	1	1	0

Fp PHV

How can we seed data efficiently to the cloud tier?

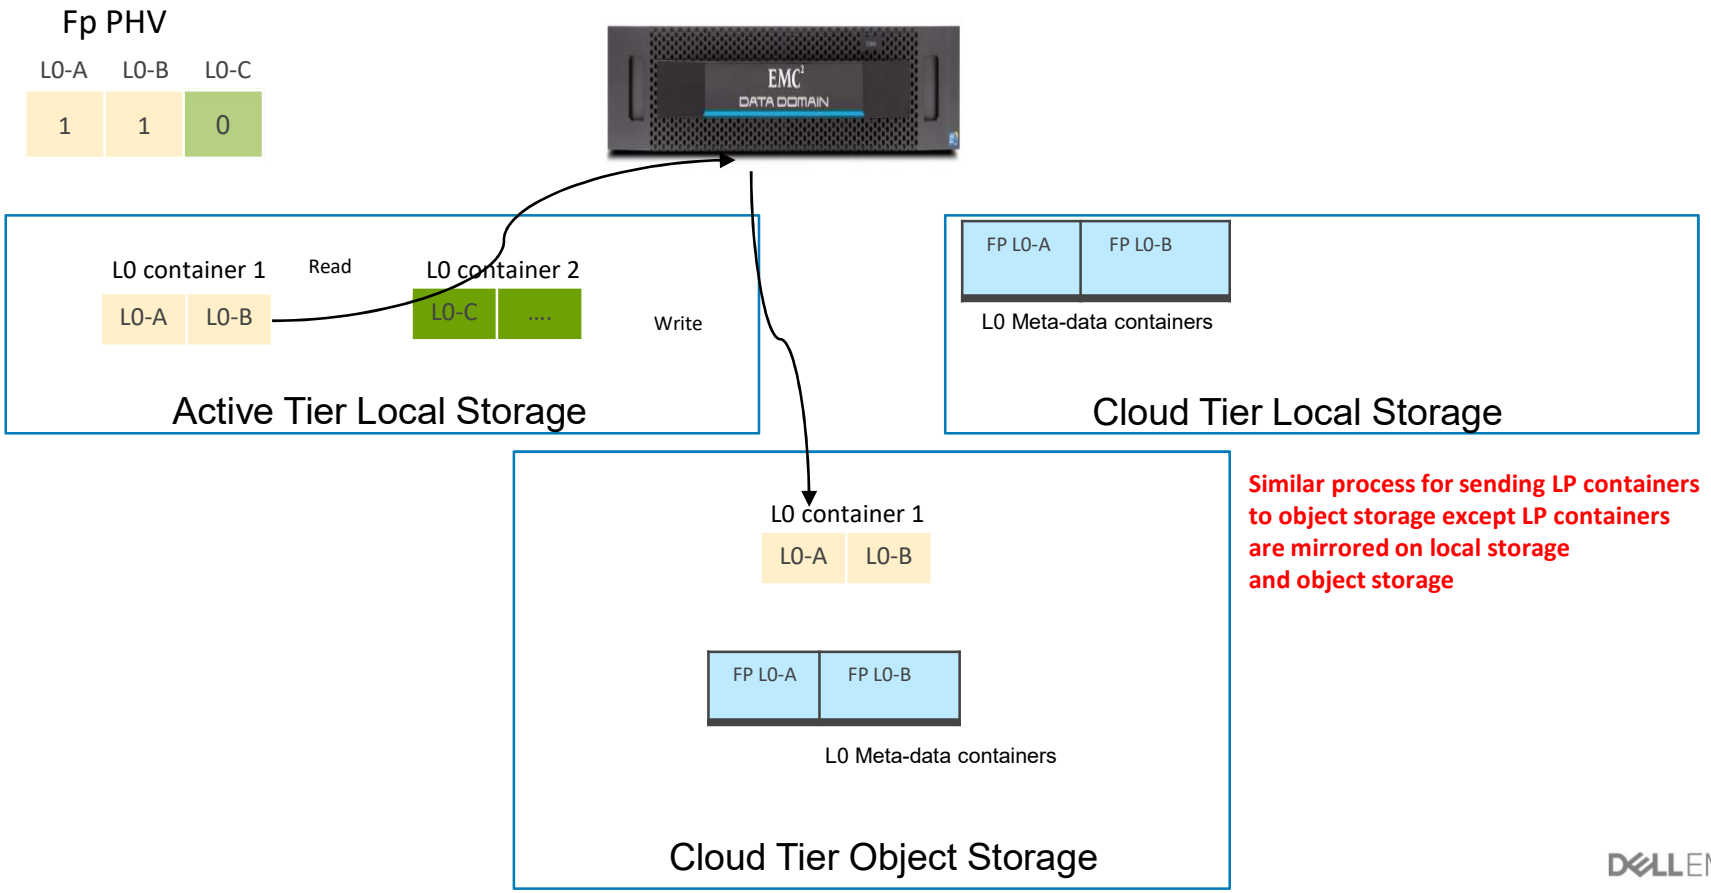

Similar process for sending LP containers to object storage except LP containers are mirrored on local storage and object storage

DELL EMC

How can we move incremental data to cloud tier?

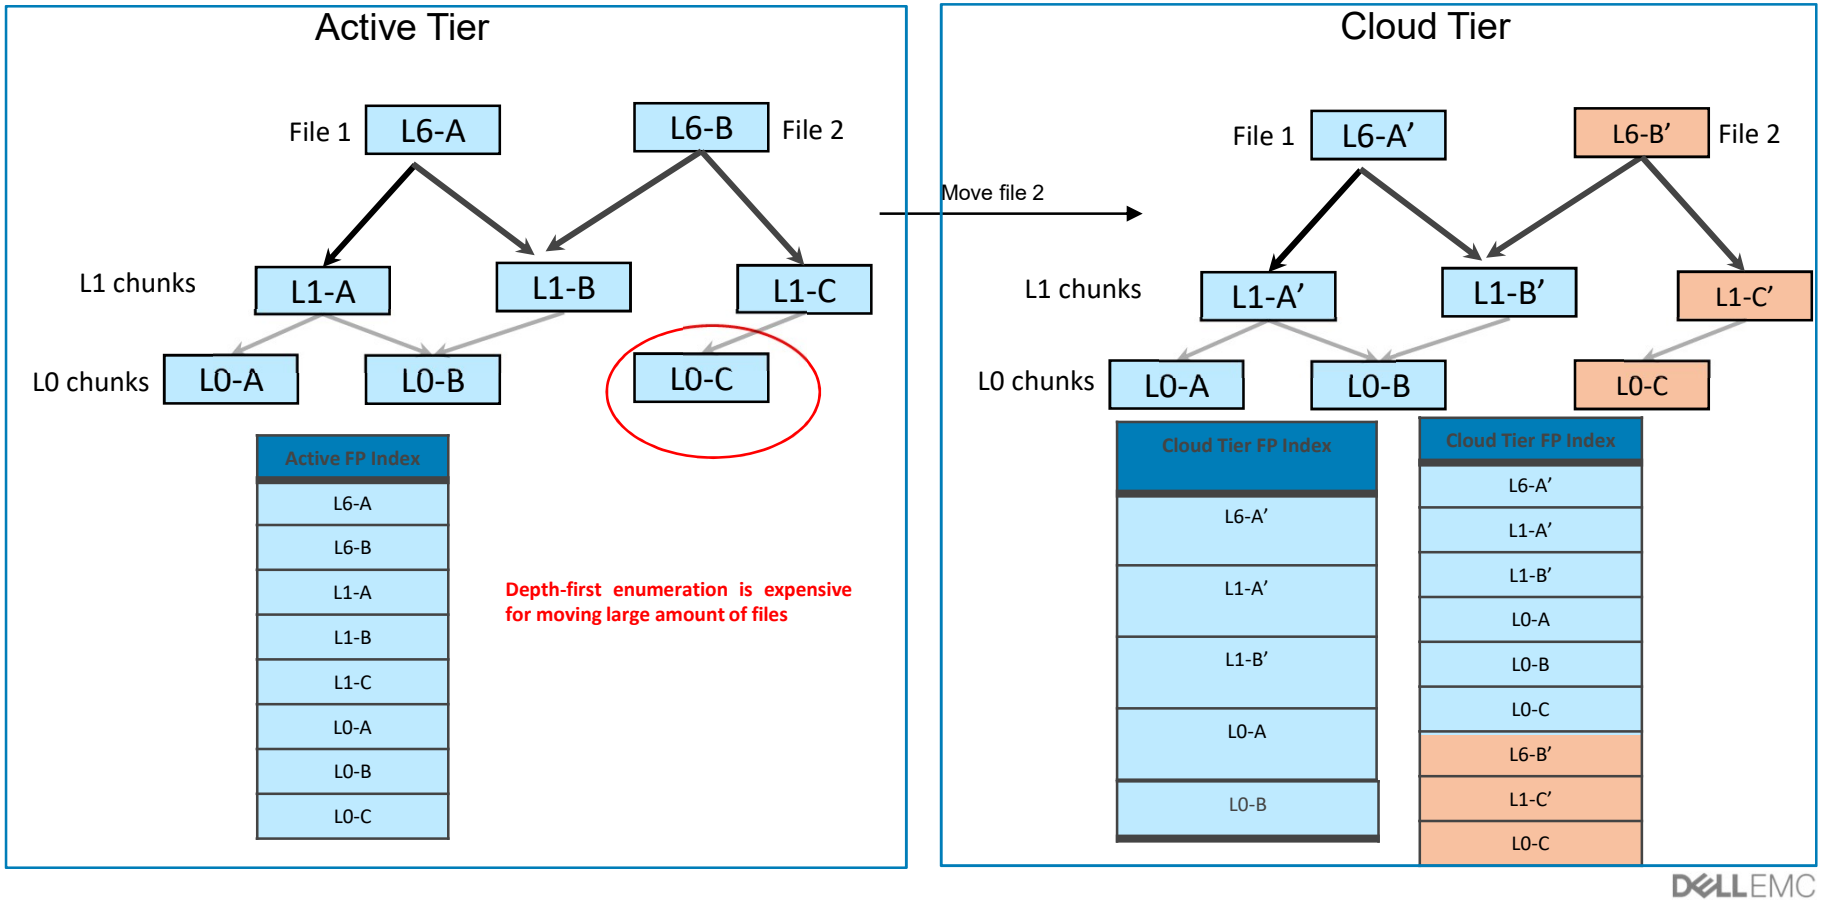

DELL EMC

How can we efficiently clean dead data on cloud tier?

- Build Fp PHV and physically scan candidate file meta-data chunks and mark chunks live
- Walk Metadata-Container set and find which L0 containers to clean
- Create and send recipe for copying offset ranges from old L0 containers into new L0 containers
- Microservice performs copy forward of L0 containers (Copy C1 & C2 into C3)
- Update meta-data by copy forwarding metadata containers (Copy C4 & C5 into C6)

Similar process for cleaning LP Containers

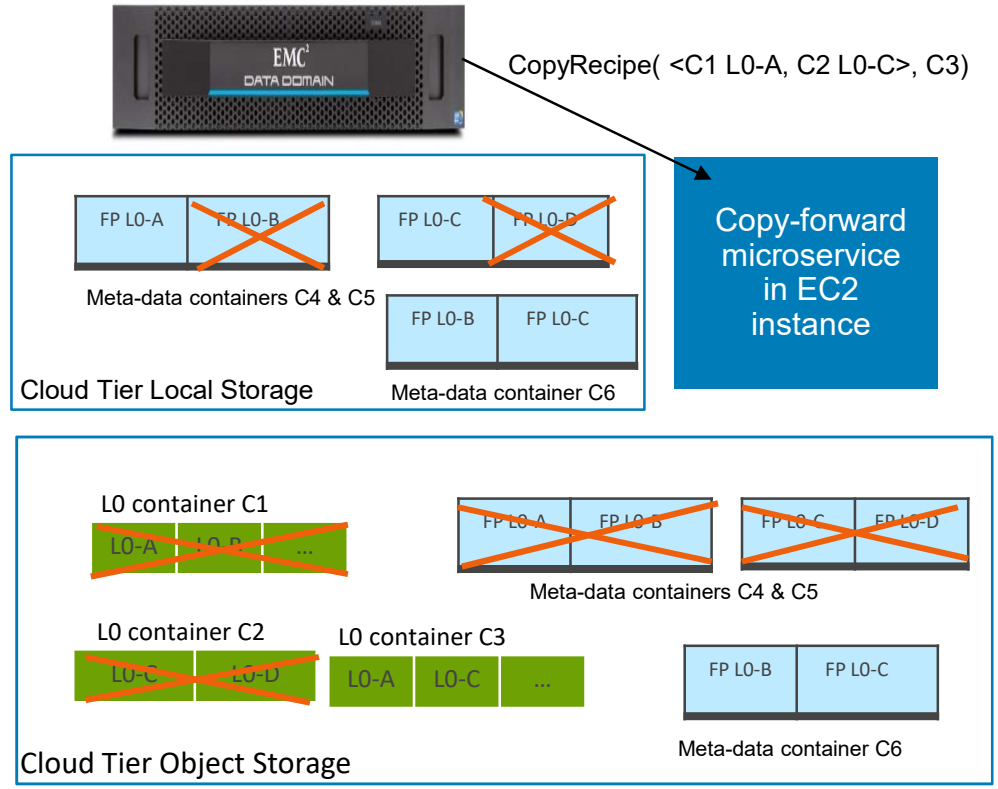

In private cloud like ECS, we have an API which performs similar copy service

DELL EMC

Evaluation Methodology

- Deployed systems:
  - Filtered out systems with less than 1TB in physical capacity on cloud tier
  - Randomly selected 200 systems out of hundreds of field systems
  - For GC, we looked at these 200 systems who ran cleaning at-least once(around 40%)
- Experimental evaluation
  - Synthetic Load generator to generate data
    - First generation is randomly generated, followed by generations with random deletes, shuffles and addition
    - 5% change rate between generations

Field Analysis

So far our customers have saved hundreds of thousands to millions of dollars due to the benefits of deduplication and local compression

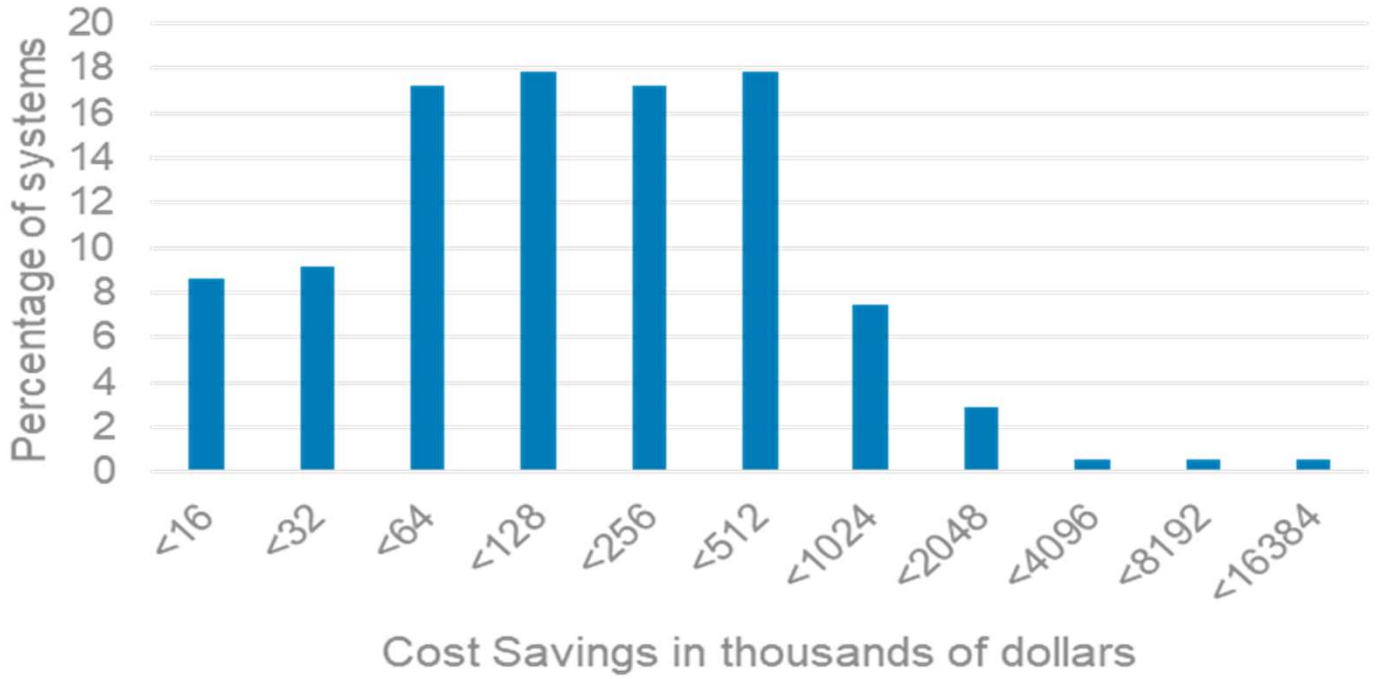

DELL EMC

Summary of Field Evaluation Results

- Some customers are writing less than 100TB (logical) per month and others more than 500TB (logical) per month
- Total compression ranges from 4X to 100X
- In order to reduce the active tier on-prem cost, some of our customers have moved more data to their cloud tiers as compared to their active tiers
- Churn (data deleted per period) on cloud tier is 0-5% of physical capacity per month whereas on active it is 10% per week
- Most customers are running cloud cleaning infrequently

Internal Evaluation I

- FMIG & Restore from private cloud ECS is faster than public cloud AWS by nearly 2X
- High-end platform FMIG & Restore performance is up-to 30% better than mid-range platform
- Gen42 FMIG & Restore performance 5X-6X better than Gen0 FMIG & Restore
- FMIG performance is mostly 50% better than restore performance

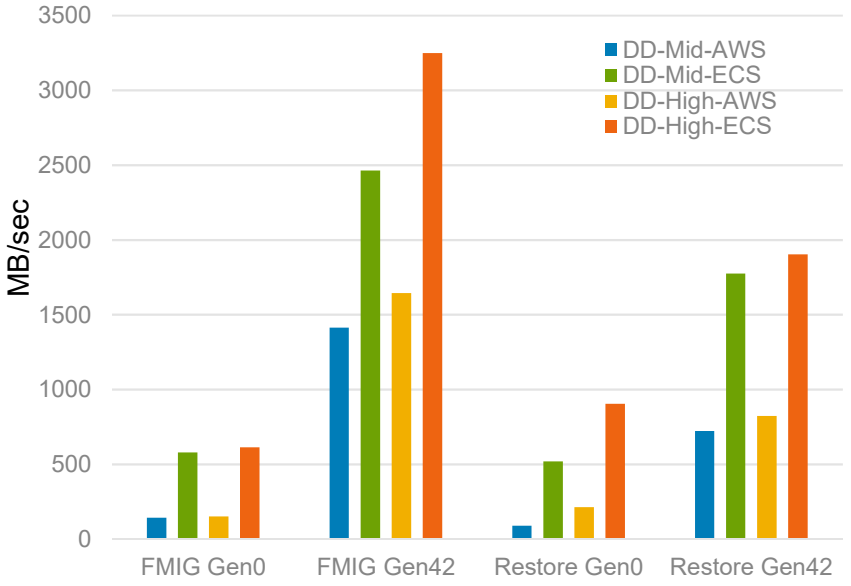

Internal Evaluation II

- FMIG Gen 0 is similar to Seeding Gen 0
- As generations increase, seeding is 2X faster than FMIG

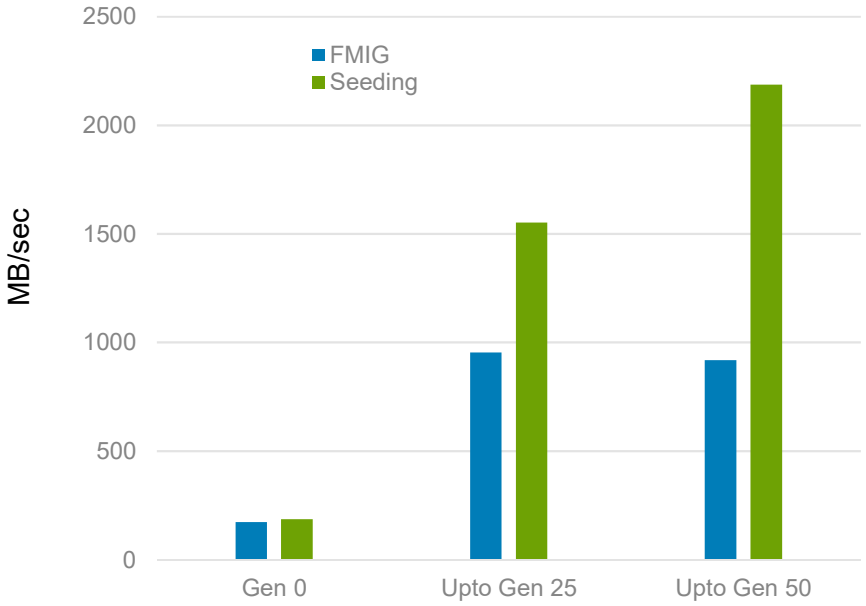

DELL EMC

Internal Evaluation III

- Copy forward using private cloud provider API is 500 MB/s which is good enough for the low-churn we see on the cloud tier
- Copy forward speed using microservices increases linearly with the increase in number of instances

Conclusion

- Data Protection continues to be a priority as our customers move parts of active tier data to the cloud
- Adding cloud tier to our active tier architecture involved adding several techniques
  - Mirroring meta-data helps support efficient deduplication & GC
  - Using perfect hashing and physical scanning to solve problems around free space estimation, seeding & cloud GC
  - Cloud GC algorithm using cloud provider API/Microservices to perform copy forward in the object storage
- Experiences from our initial customers show that customers are using cloud tier to either reduce cost of on-prem active tier storage and/or retaining the data for long term archival storage


DELL EMC