99 Deduplication Problems

Philip Shilane, Ravi Chitloor, and Uday Kiran Jonnala
EMC Corporation

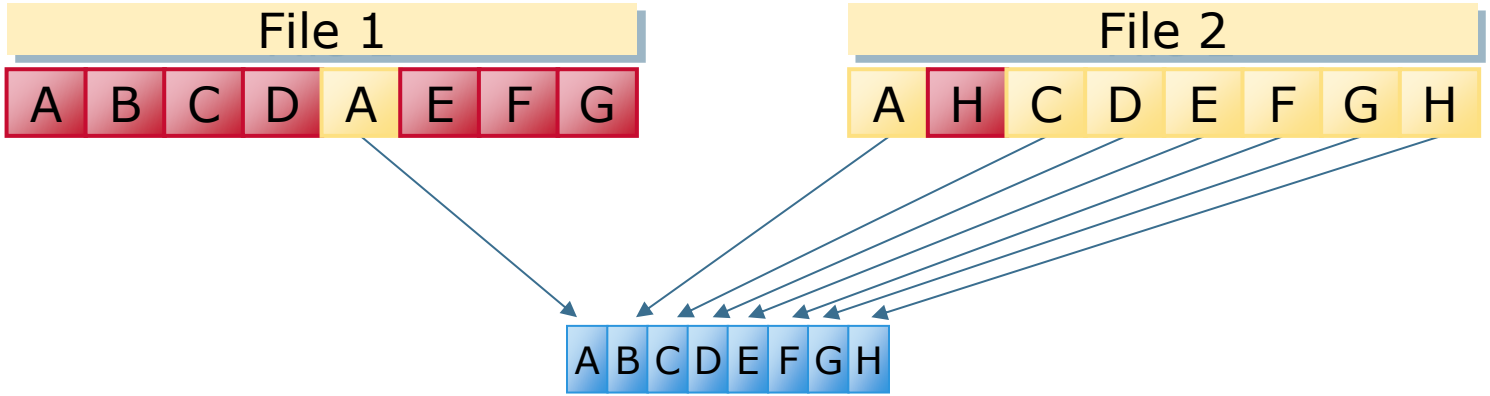

Deduplication Proliferation

- Numerous products from Dell, EMC, HPE, IBM, NetApp, Nimble, Pure, …
- Numerous publications

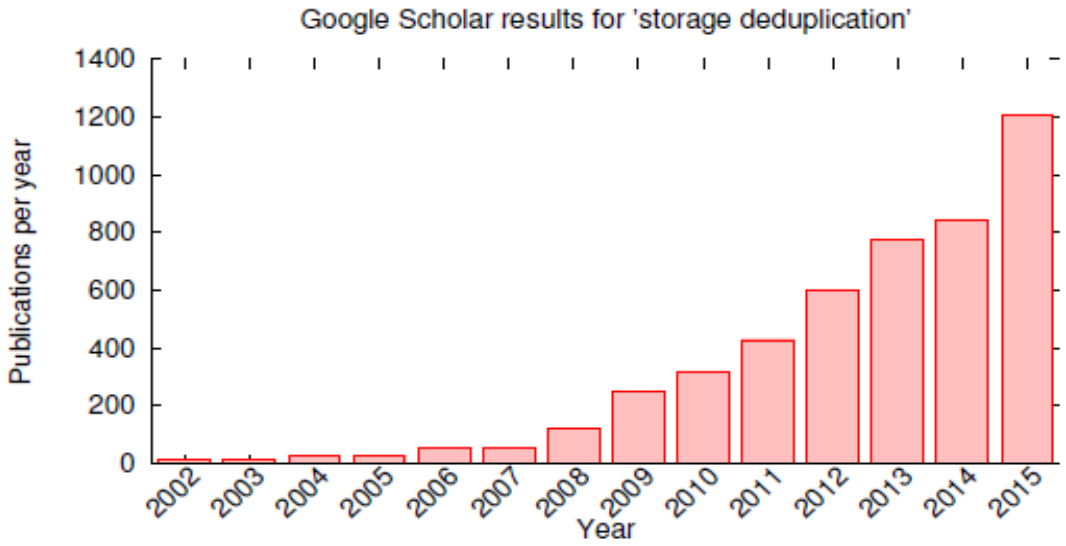

2 Non-problems and 99 Novel Problems

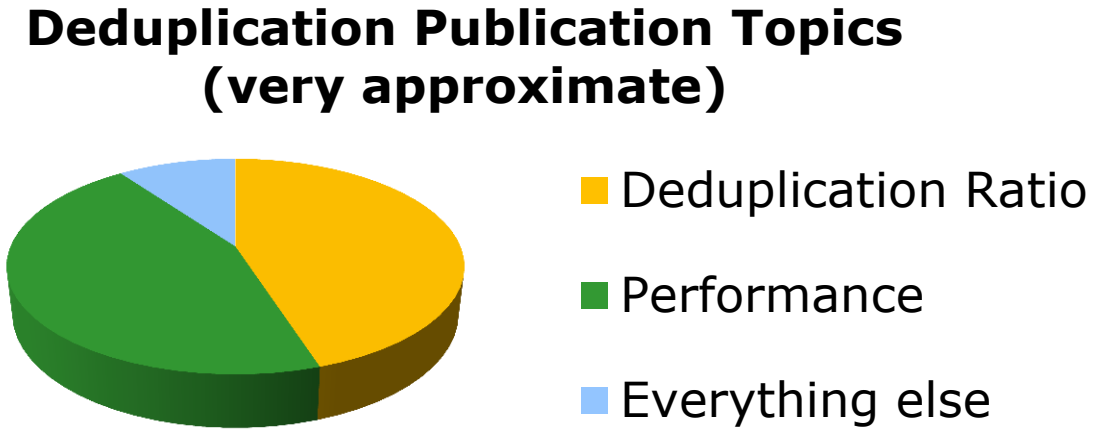

2 Non-problems and 99 Novel Problems

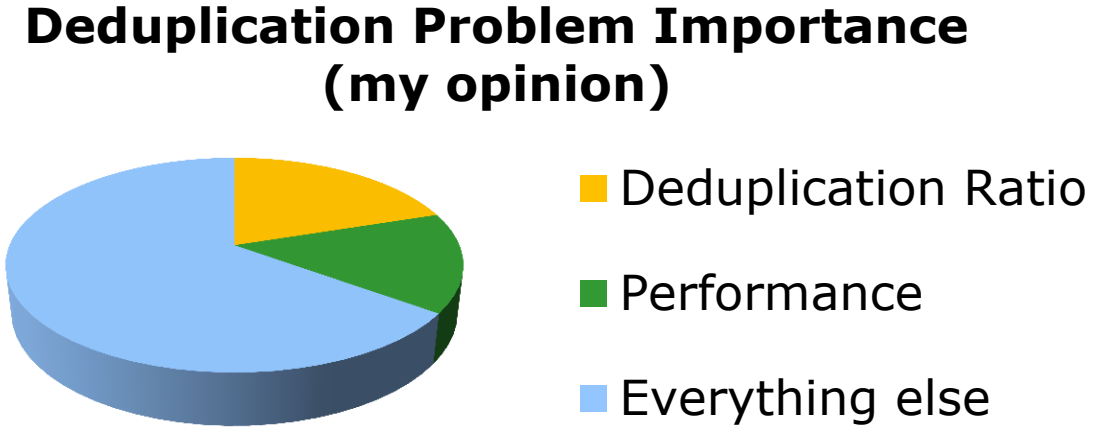

Please work on new problems!

- Customers expect fully featured storage products
- Novel problems are more fun to research

Outline

- Capacity
- Management
- Quality of Service
- Security and Reliability
- Chargeback for Service Providers
- Traces and Load Generators

Capacity: How Much Space Is Available?

How much more can be stored?

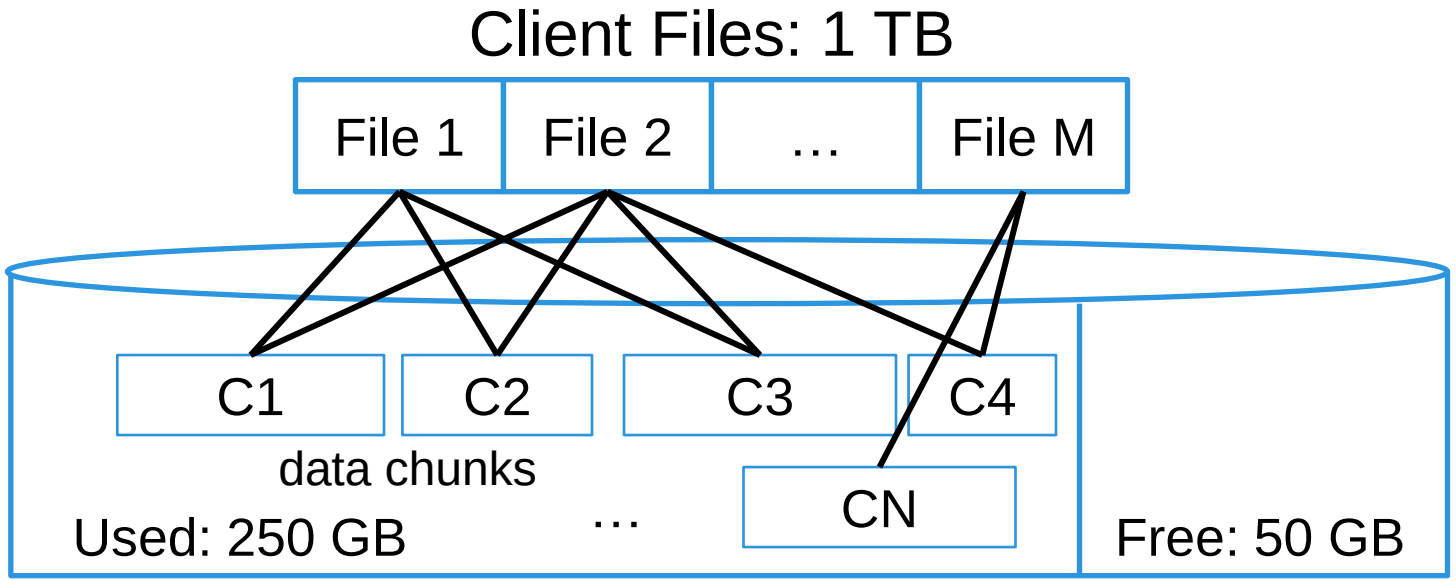

Capacity: What Should I Delete?

How much space will be freed by deleting a file?

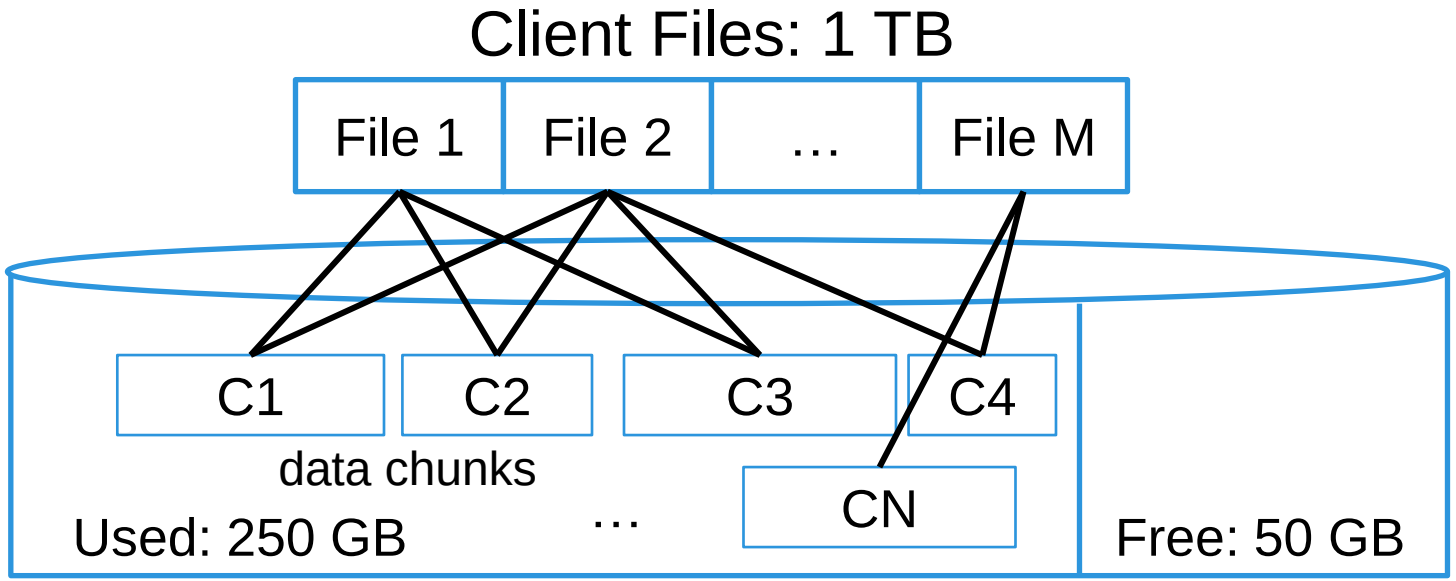

Capacity: What Should I Delete?

How much space will be freed by deleting a file?

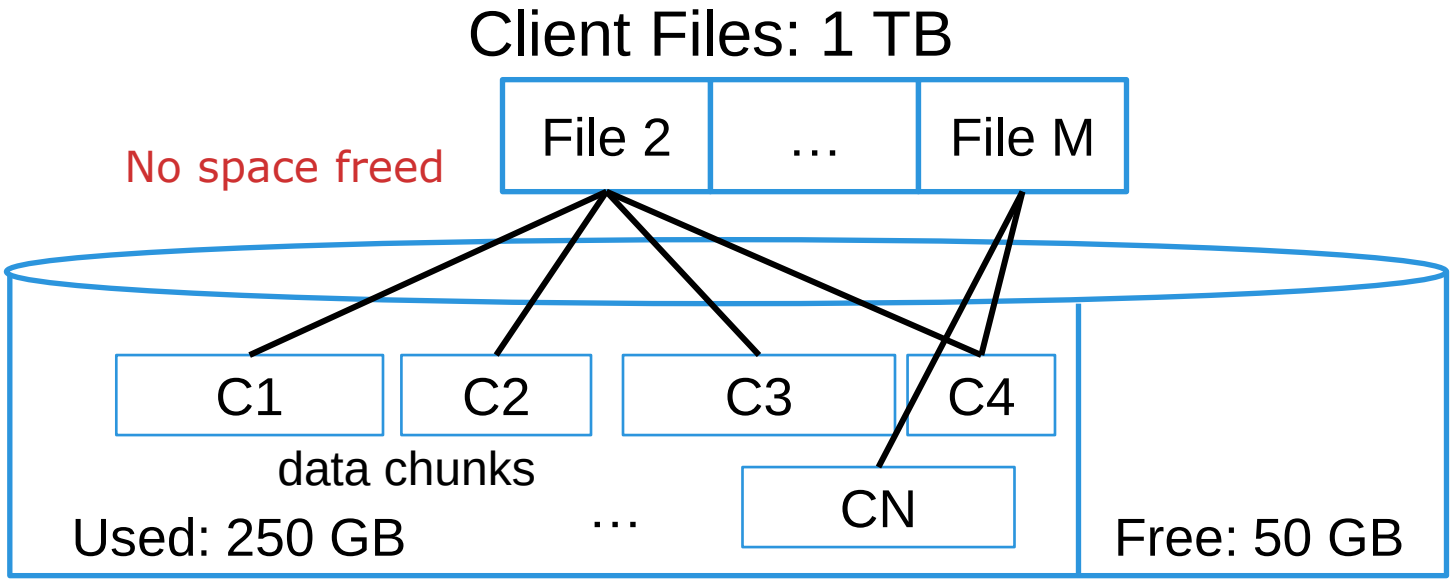

Capacity: What Should I Delete?

Freeable space changes dynamically

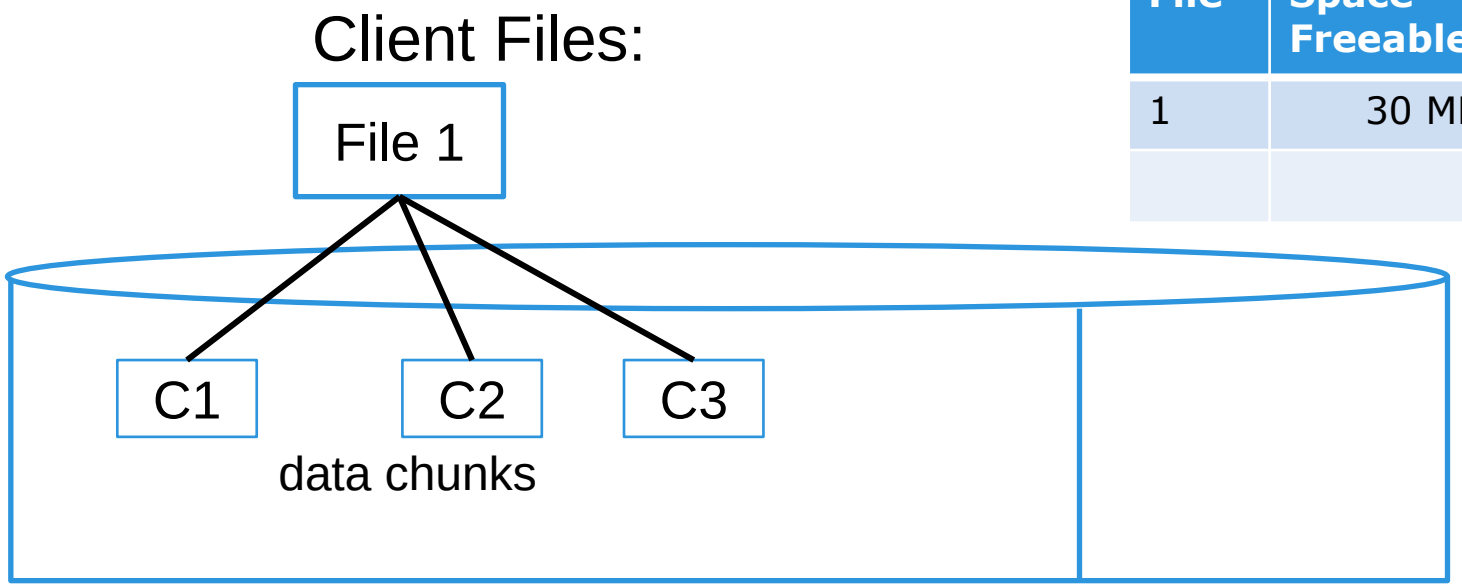

File	Space Freeable
1	30 MB

Capacity: What Should I Delete?

Freeable space changes dynamically

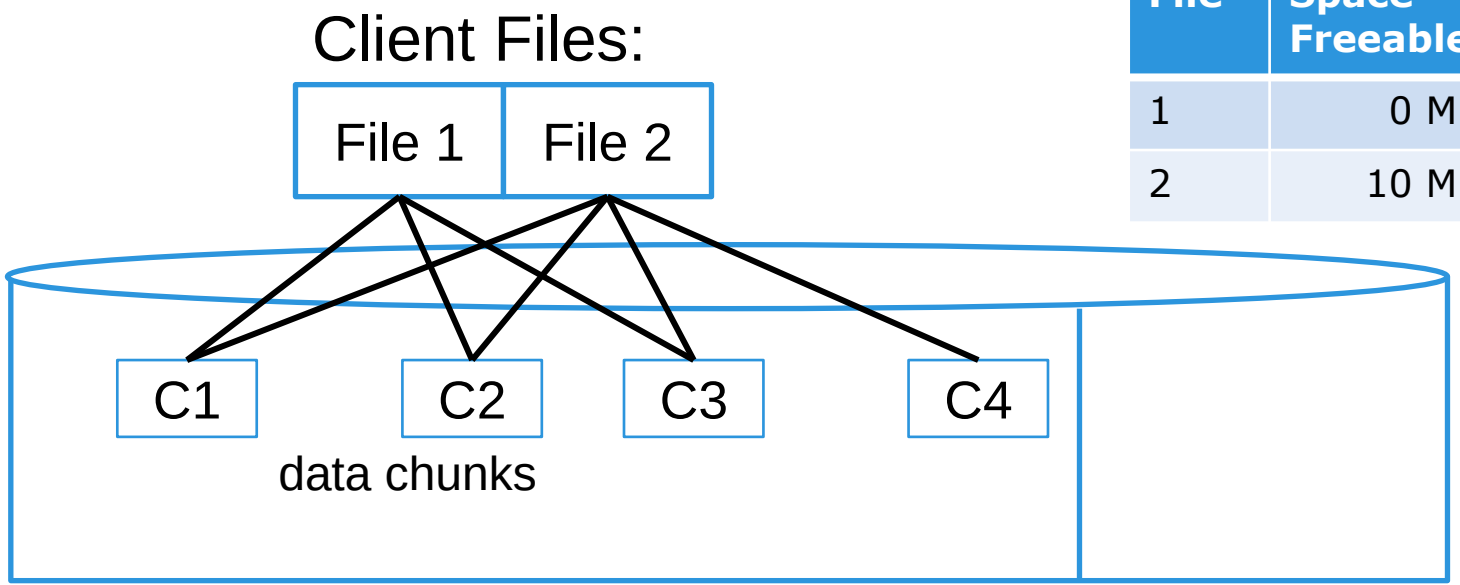

File	Space Freeable
1	0 MB
2	10 MB

Capacity: What Should I Delete?

Freeable space changes dynamically

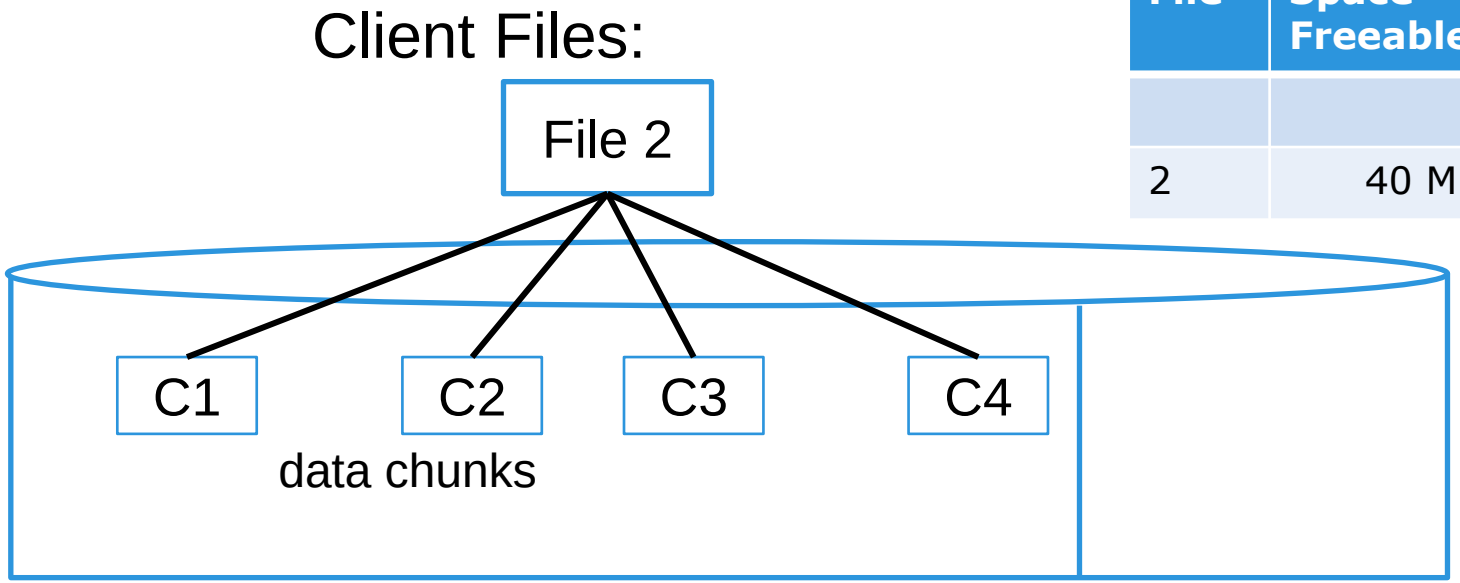

File	Space Freeable
2	40 MB

Management

- Initial Sizing
  - Estimation tools
  - Partitioning space for users
- Migration from legacy storage
  - Migrate data onto deduplicated storage
- Reporting
  - Capacity usage
  - Performance
  - Network usage
  - Per volume, system, cluster

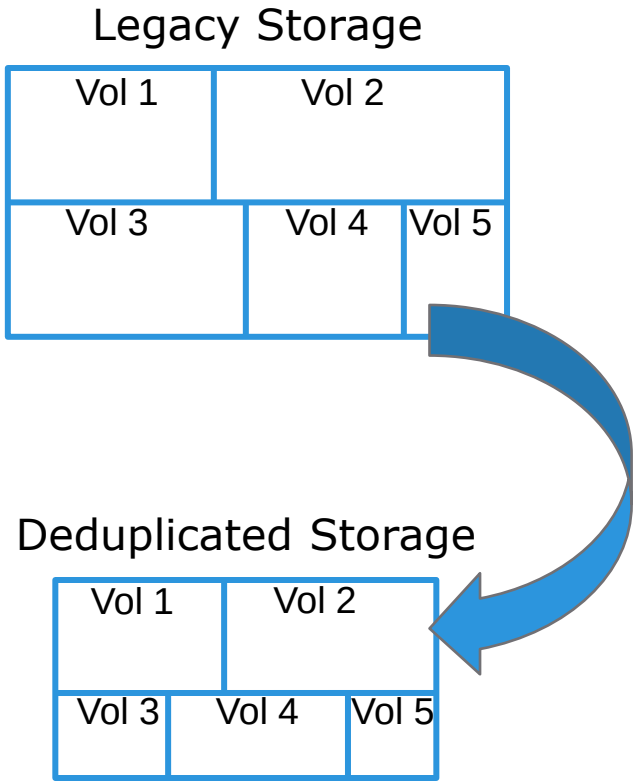

Quality of Service (QoS)

- Defined broadly: latency, throughput, and priority levels
- Customers may specify service level agreements
  - It is not good enough to simply avoid client timeouts
- Different from strictly performance improvements

QoS: Expected Performance

- Deduplication adds additional levels of indirection
- Shared content creates unpredictable performance
  - Chunks C and Y were written consecutively in File M
  - Due to deduplication, they are stored in different containers
  - Sequential read ahead may not help due to containers having chunks from multiple files

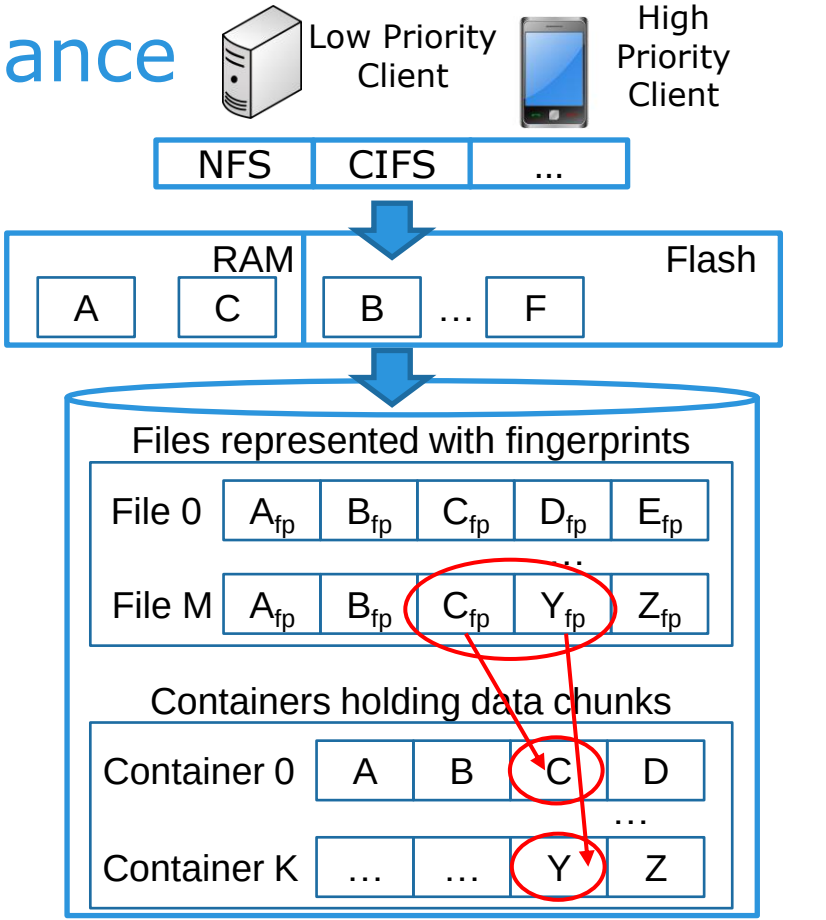

QoS: Caching

- Caching is a common approach to improve latency
- Deduplication creates a new form of content sharing
- The High priority client performs work that also benefits the Low priority client
- The High priority client could exhaust its quota, while the Low priority client can continue.

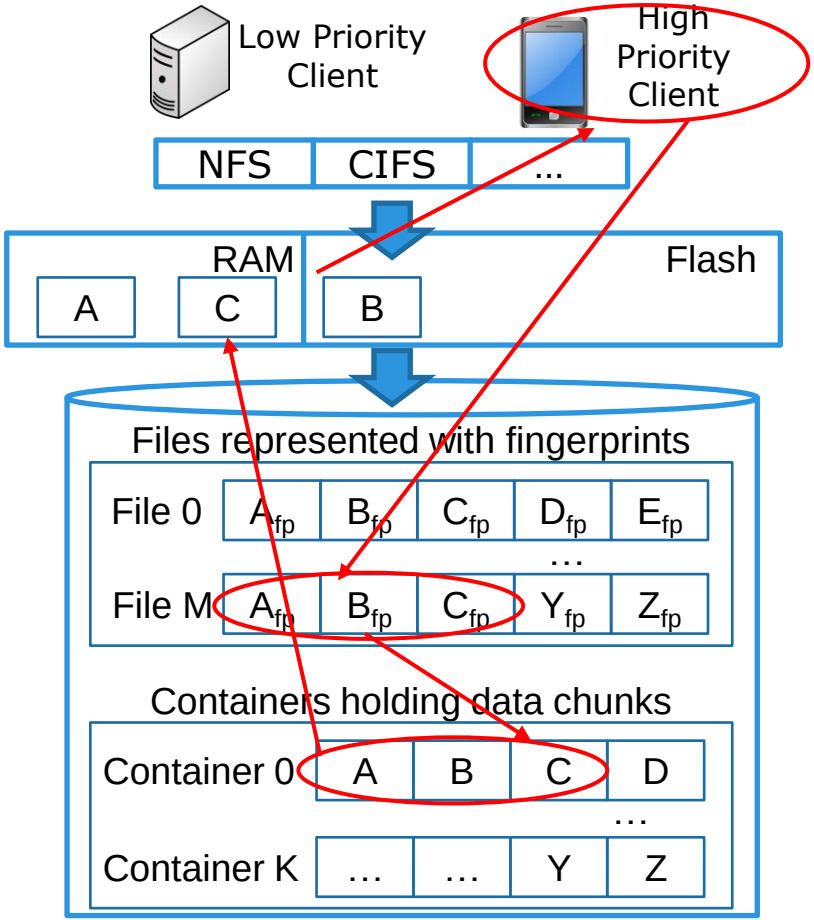

QoS: Caching

- Caching is a common approach to improve latency
- Deduplication creates a new form of content sharing
- The High priority client performs work that also benefits the Low priority client
- The High priority client could exhaust its quota, while the Low priority client can continue.

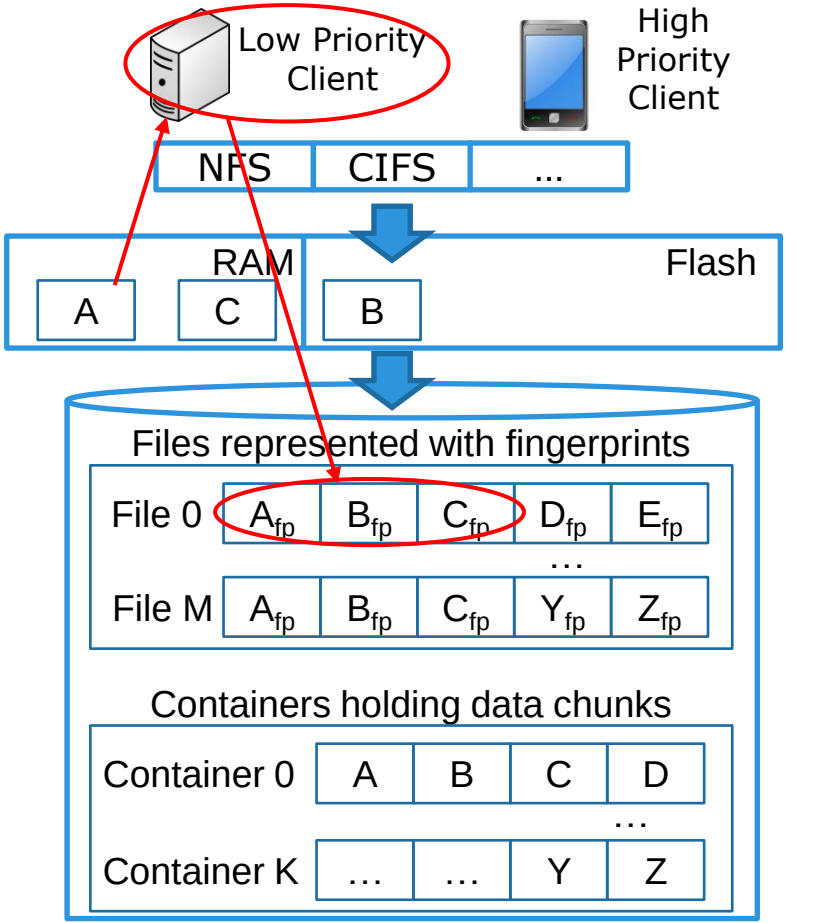

QoS: Resource-intensive Background Tasks

- Garbage collection
- Replication
- Verification (online fsck)
- Disk reconstruction

Security and Reliability

- Maintain advantages of deduplication while securely storing data, preventing
  - Unauthorized access
  - Knowledge of content
  - Data tampering
- There is already research on these topics, but more is needed
  - Converged encryption
  - Timing issues to reveal content

Security and Reliability

Likelihood that needed data will be available

Deduplicated systems remove redundancy but add:

- RAID, Erasure Encoding, Versioning, and Remote Replication

How do we analytically compare the reliability of these approaches?

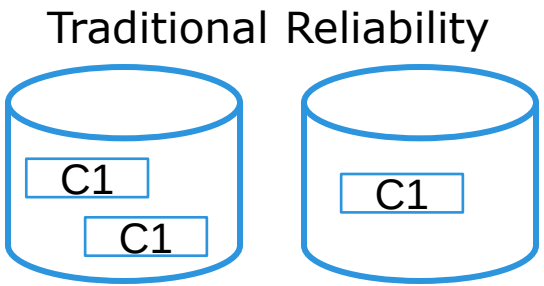

Security and Reliability

Likelihood that needed data will be available

Deduplicated systems remove redundancy but add:

- RAID, Erasure Encoding, Versioning, and Remote Replication

How do we analytically compare the reliability of these approaches?

Traditional Reliability

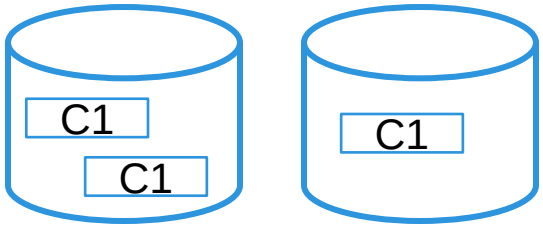

Deduplication Reliability

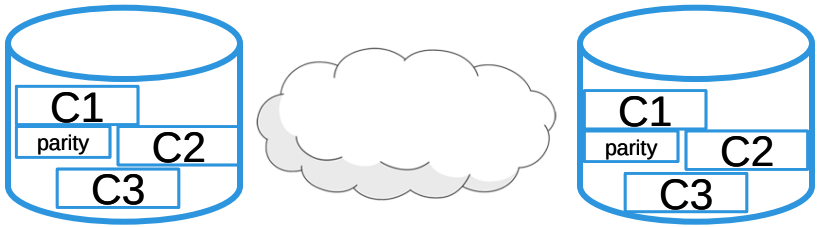

Chargeback for Service Providers

- QoS across tenants sharing content
- Service provider must charge appropriately
  - Too high and a customer can purchase storage itself
  - Too low and the service provider loses money
- Deduplication complicates billing
  - Capacity, CPU, I/O, network, other services
- Billing timeliness is important

Traces and Load Generators

- Unlike standard storage traces, deduplication needs the content or at least content hashes
- Only a small handful of such traces exist
  - Block traces from Florida International University
  - Static snapshots from Microsoft and Stony Brook University
- For both engineering and research experiments, we need realistic content
  - Anonymizing the data is critical

Time	R/W	LBA	Size	Hash
1	R	0	8	0xabc
2	W	8	8	0x1ab
3	R	16	16	0xffa
…				

Conclusion: Please work on new problems!

- Customers expect fully featured storage products
- Novel problems are more fun to research
- Numerous, novel deduplication problems:
  - Capacity
  - Management
  - Quality of Service
  - Security and Reliability
  - Management
  - Chargeback for Service Providers
  - Traces and Load Generators
  - Many more…

Questions?


EMC²