

Facebook's Tectonic Filesystem: Efficiency from Exascale

Satadru Pan¹, Theano Stavrinos^{1,2}, Yunqiao Zhang¹, Atul Sikaria¹, Pavel Zakharov¹,
Abhinav Sharma¹, Shiva Shankar P¹, Mike Shuey¹, Richard Wareing¹, Monika
Gangapuram¹, Guanglei Cao¹, Christian Preseau¹, Pratap Singh¹, Kestutis
Patiejunas¹, JR Tipton¹, Ethan Katz-Bassett³, and Wyatt Lloyd²

¹Facebook, Inc., ²Princeton University, ³Columbia University

Exabyte-Scale Storage Use Cases at FB

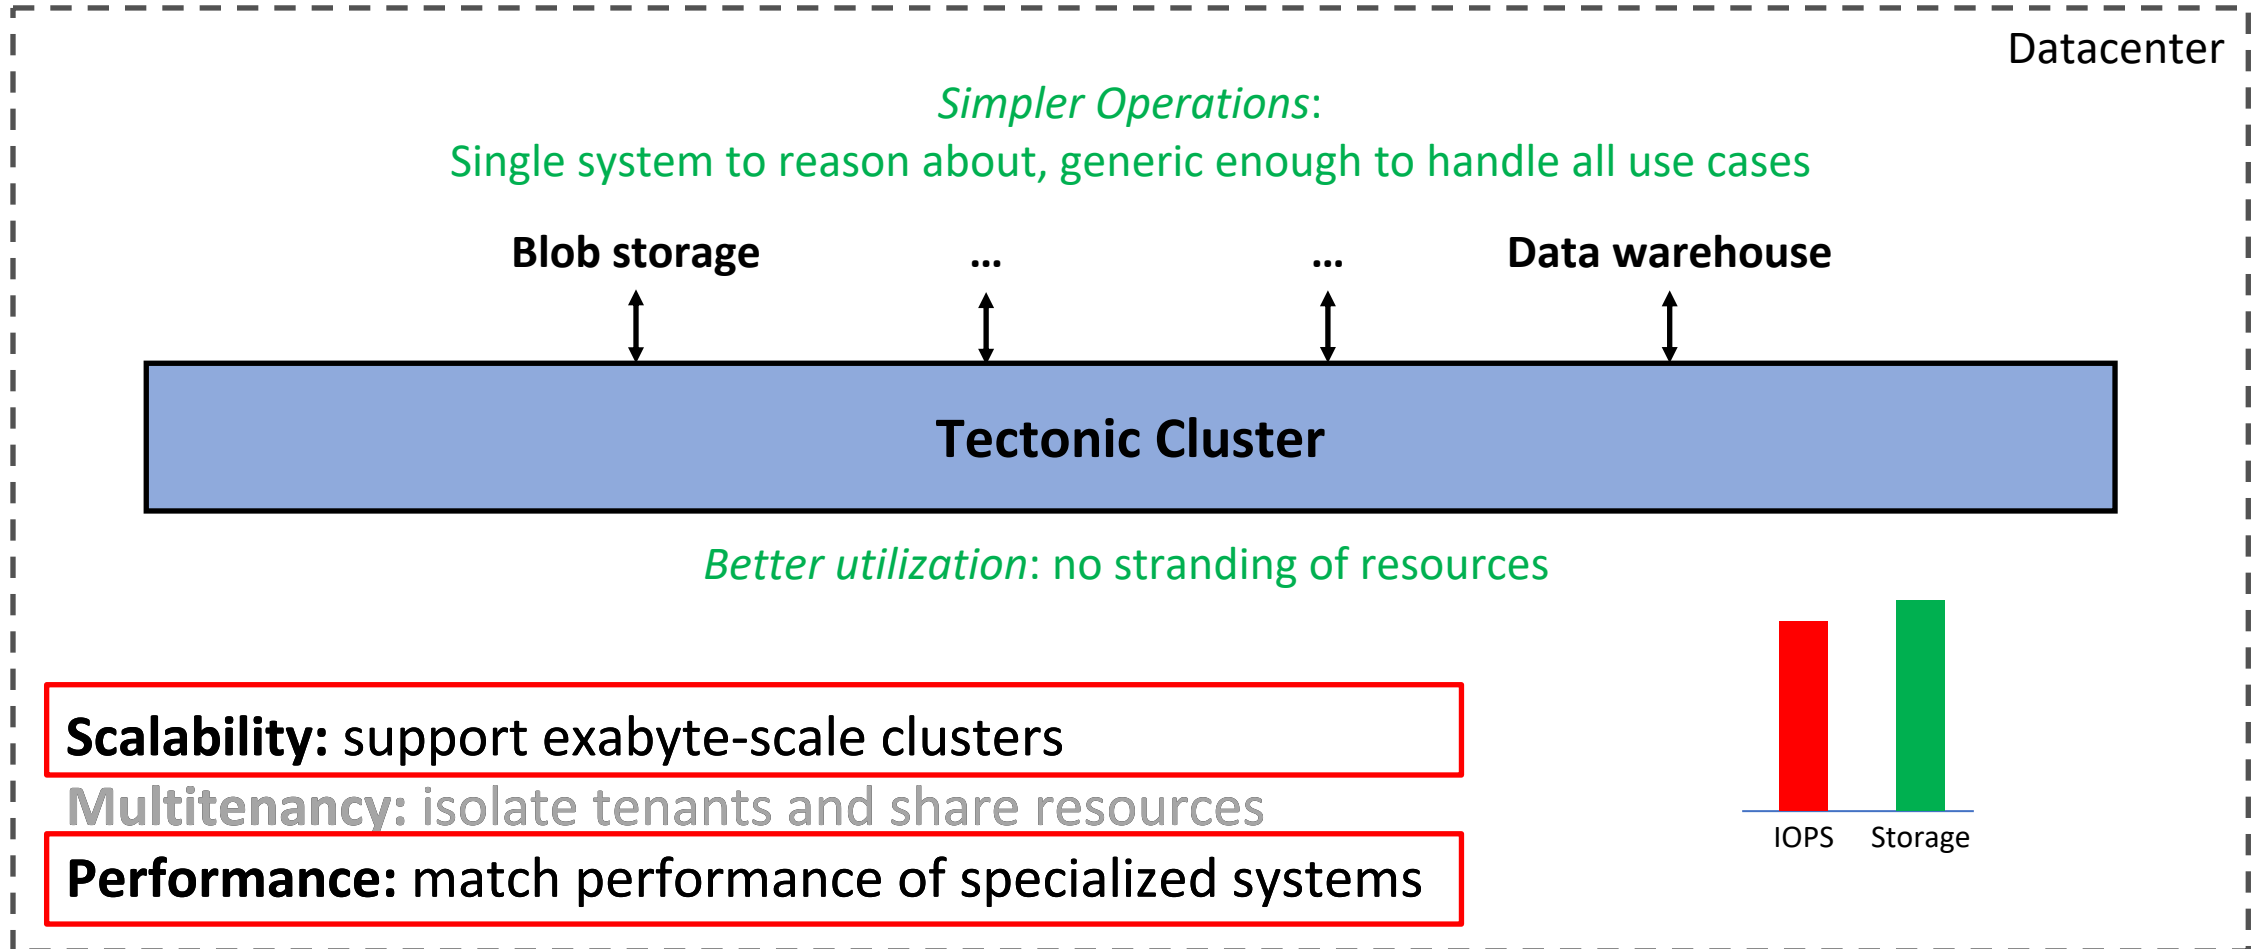
Blob storage

- Photos and videos in Facebook, Messenger attachments
- Exabytes of data
- Several KBs to several MBs in size
- Latency sensitive

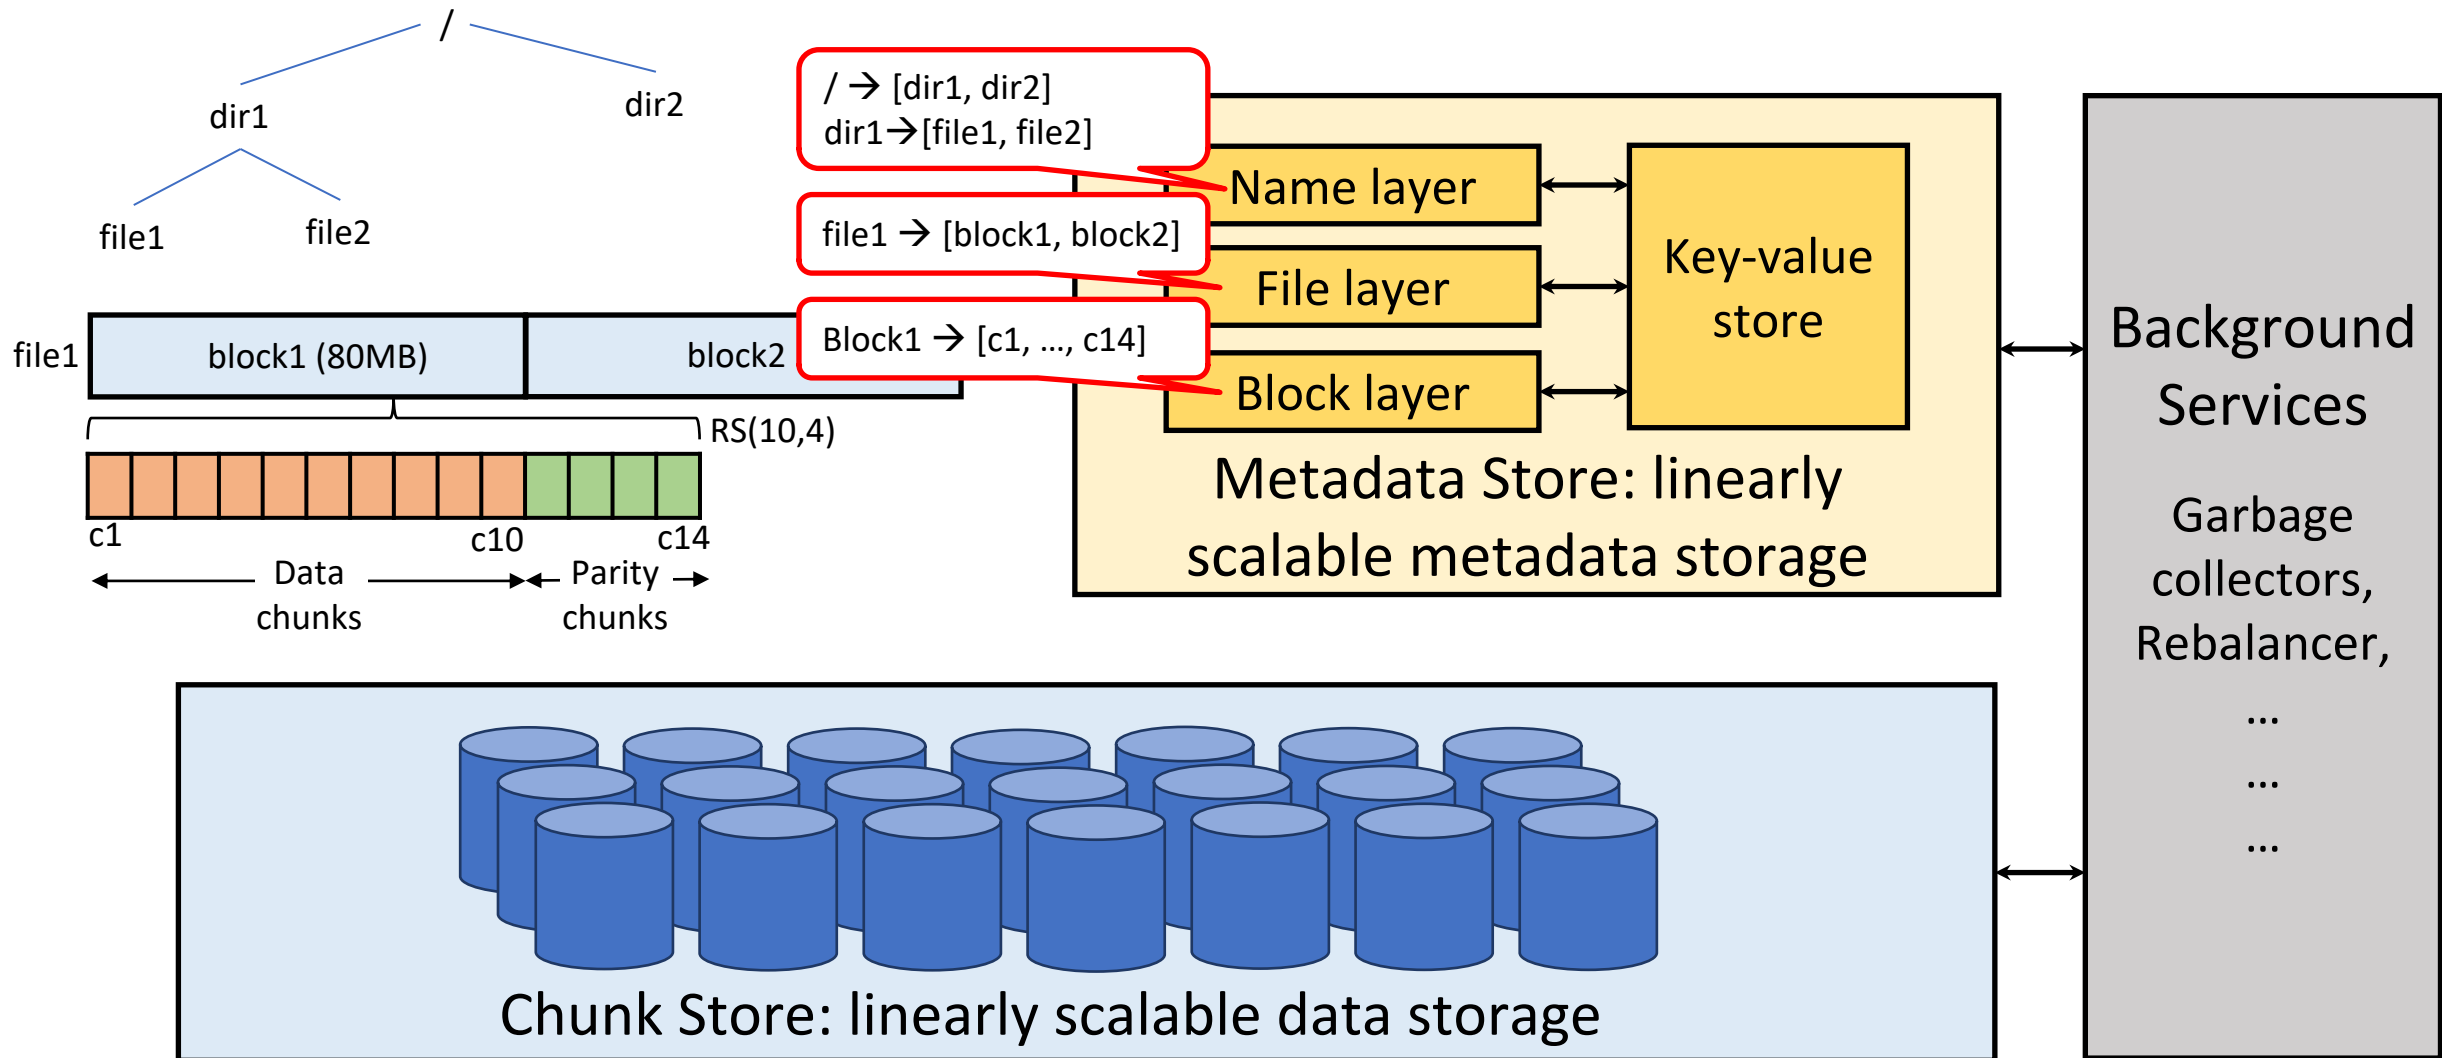
Data warehouse

- Hive tables for data analytics, machine learning
- Exabytes of data
- Reads are order of multiple MBs, writes are 10s of MBs
- Throughput sensitive

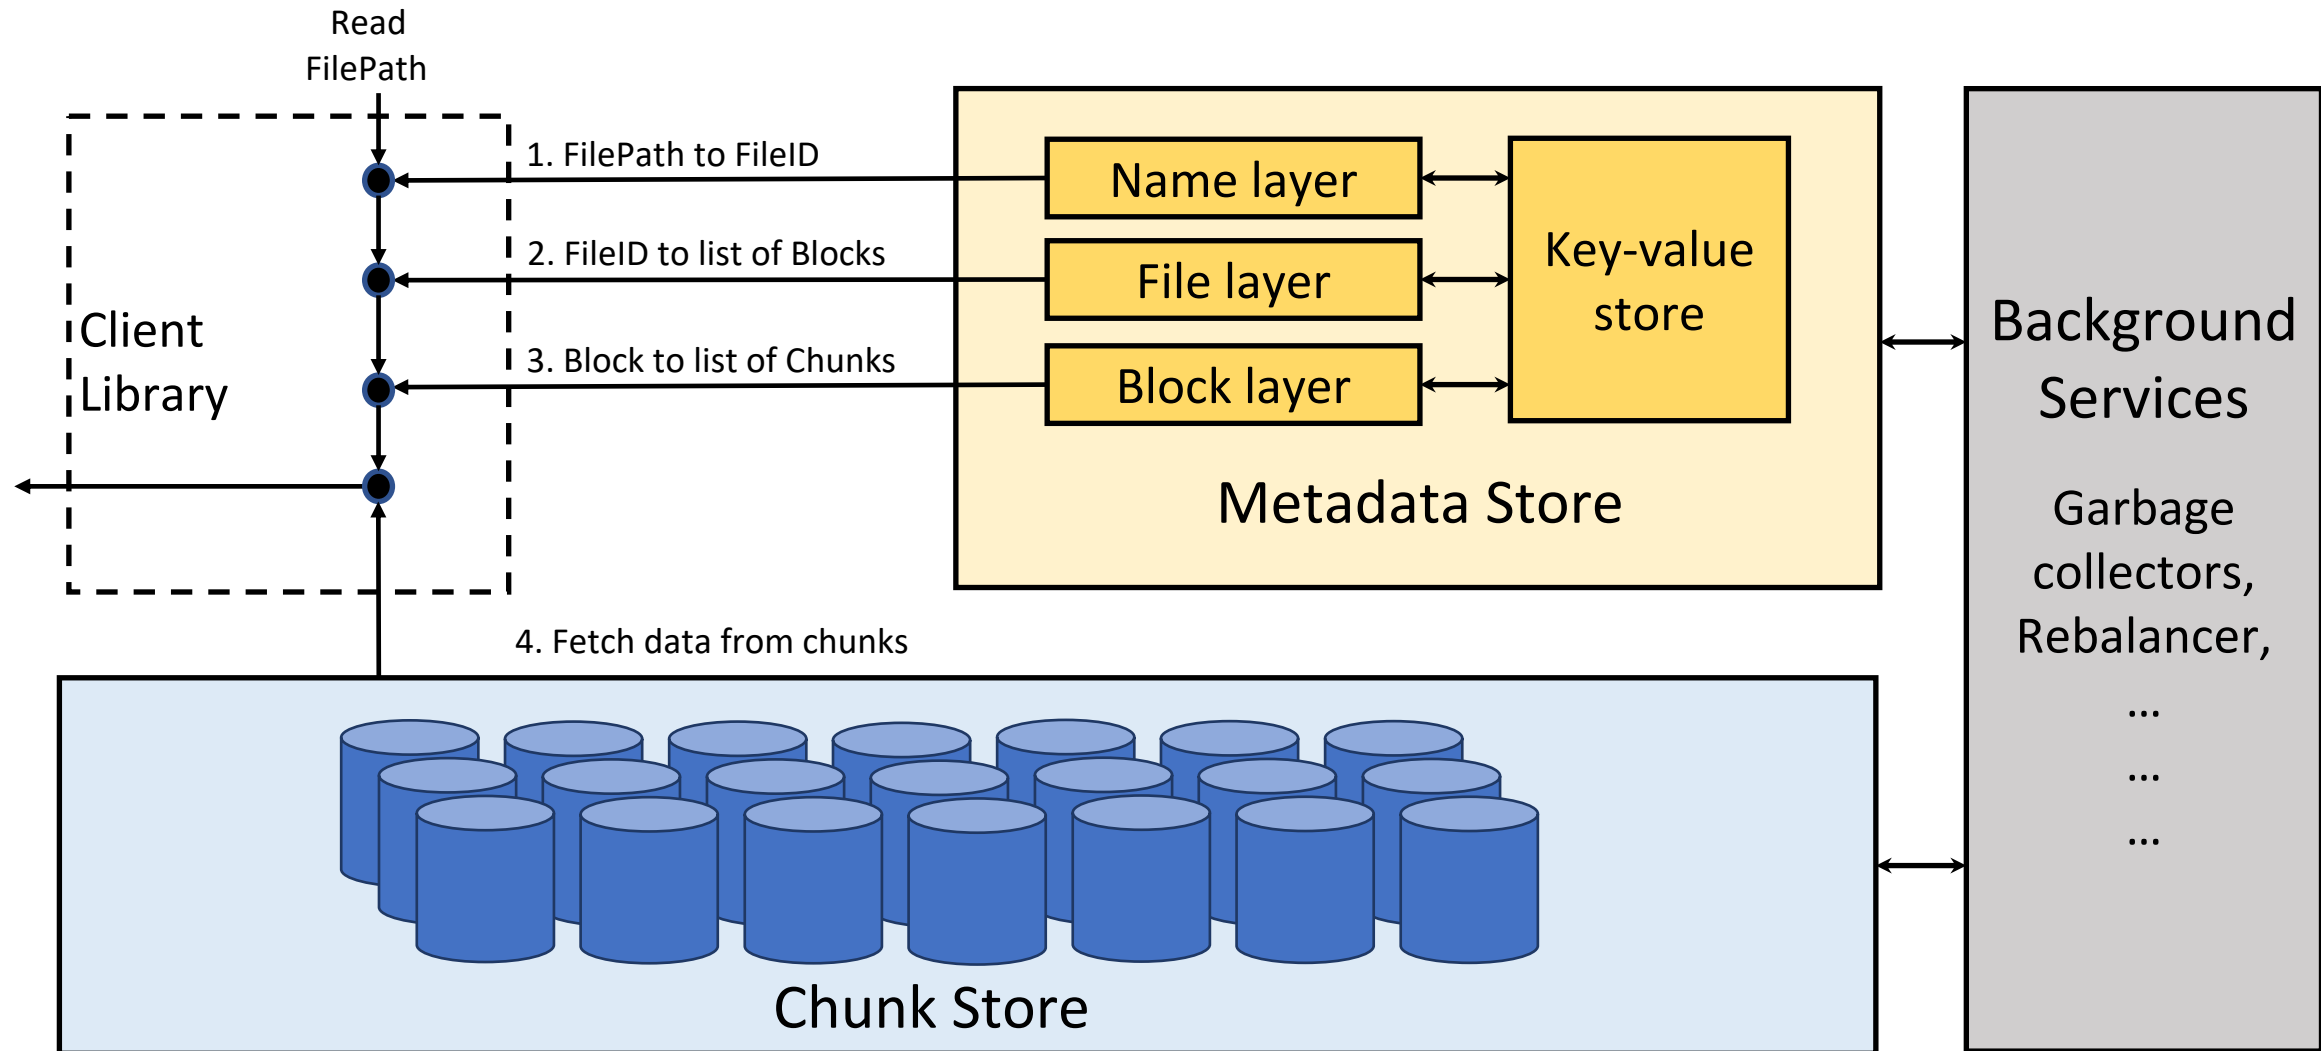
Tectonic Overview



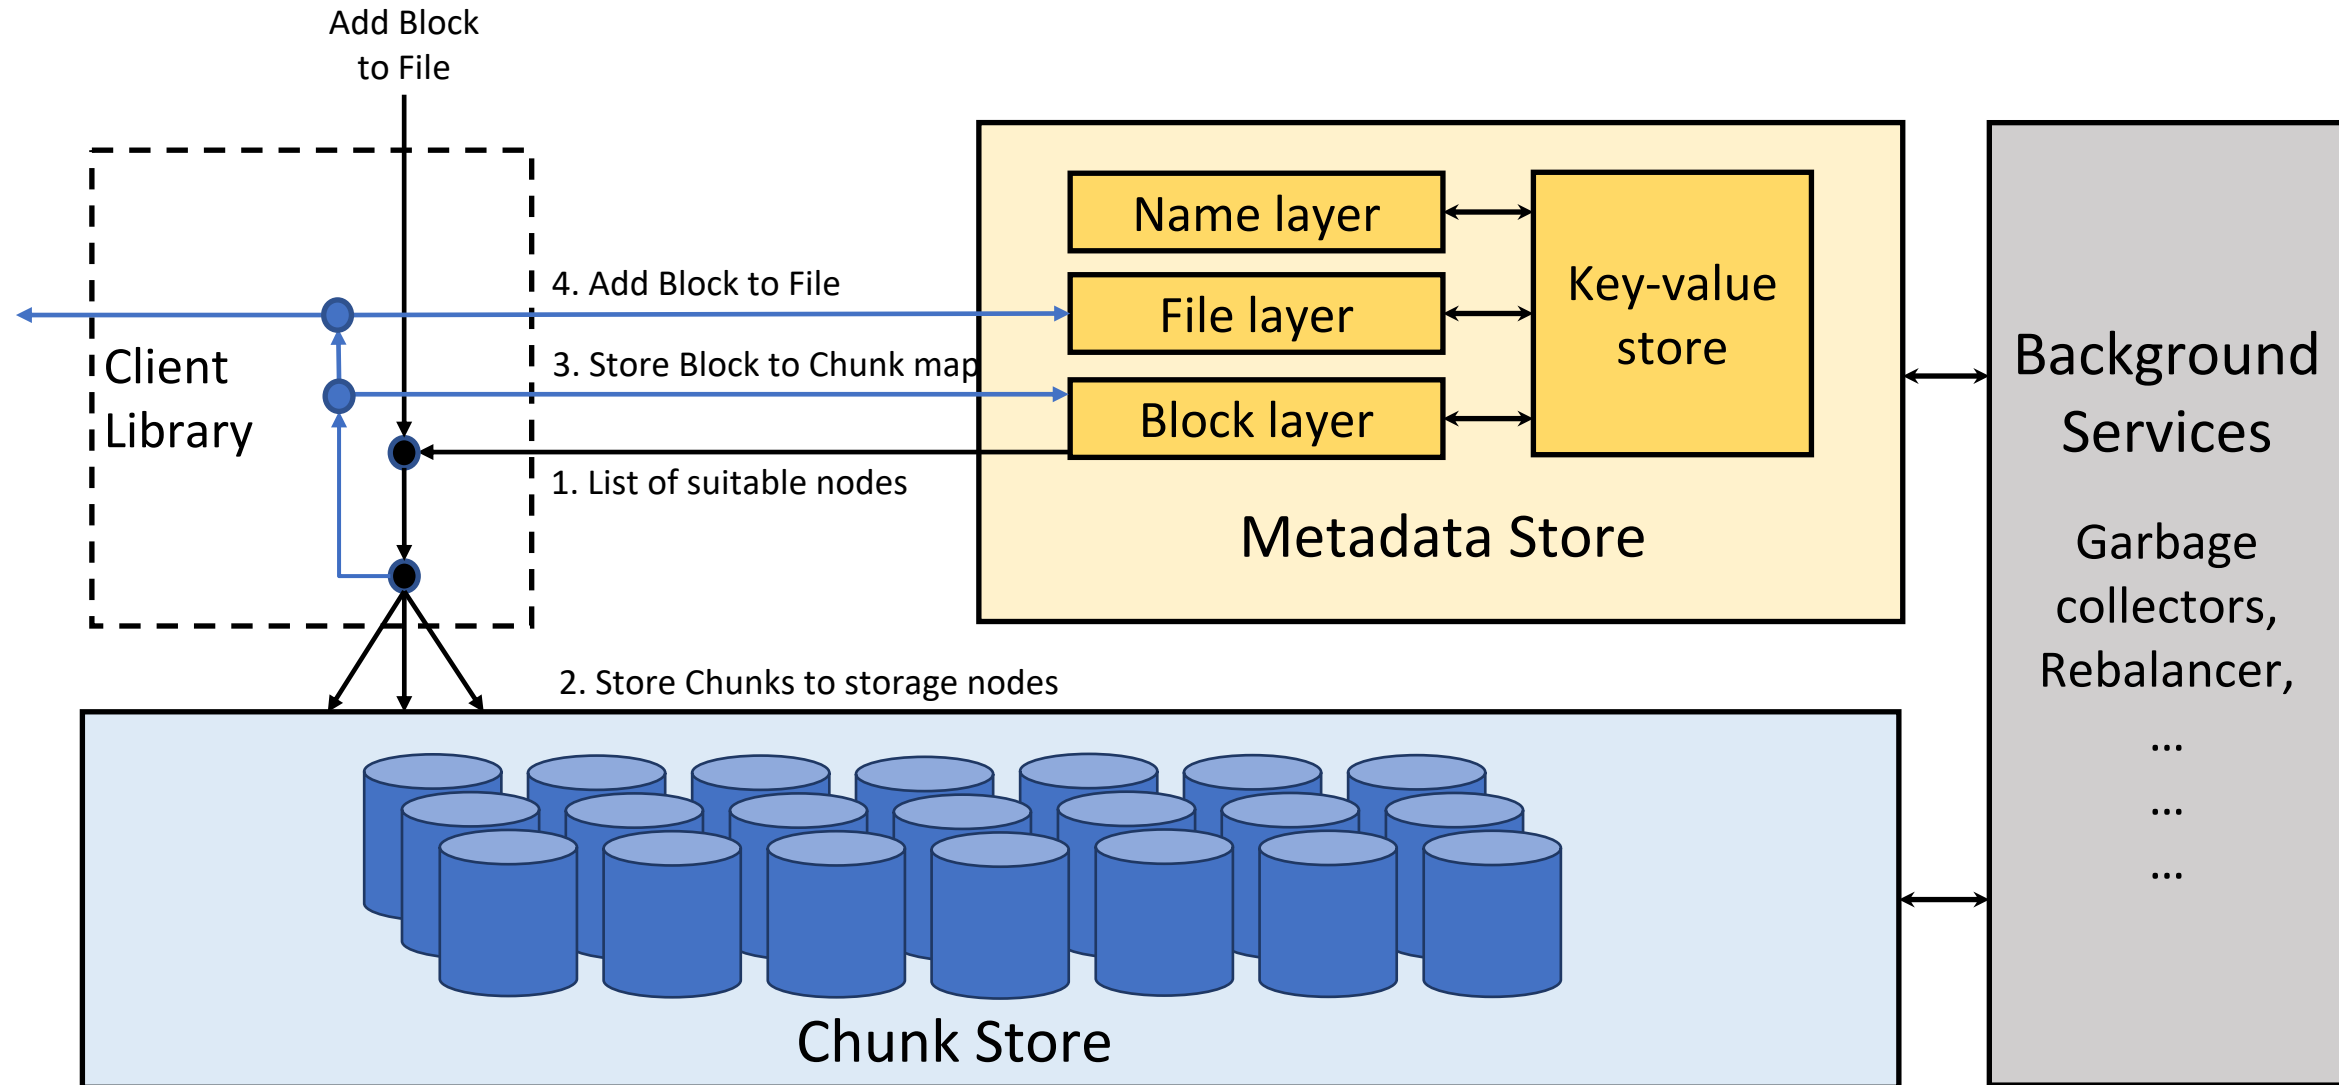
Scalability: Support Exabyte Scale Clusters



Scalability: Support Exabyte Scale Clusters



Scalability: Support Exabyte Scale Clusters



Performance: Match Specialized Systems

- Specialized storage systems optimize for the specific access pattern and performance requirements
- Tectonic uses *tenant-specific optimizations* to match the performance of specialized systems
- Optimizations are enabled by the Client Library, which runs in application binary
- Client library allows flexible and varying composition of Tectonic operations, which can be configured according to the needs of the tenant

Tenant-specific Optimizations: Appends


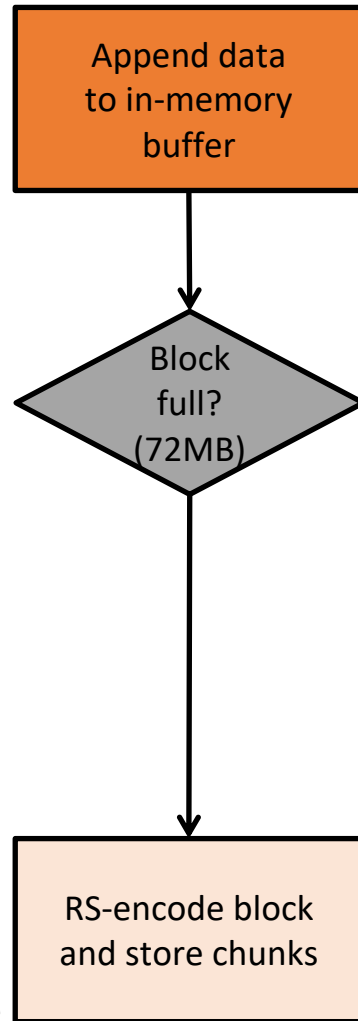
Data warehouse

Files are large
(100s of MBs)

Files are read
after the creator
closes the file

Minimize bytes
written to store file
to improve overall
throughput

Read-after-write
consistency
only after file close



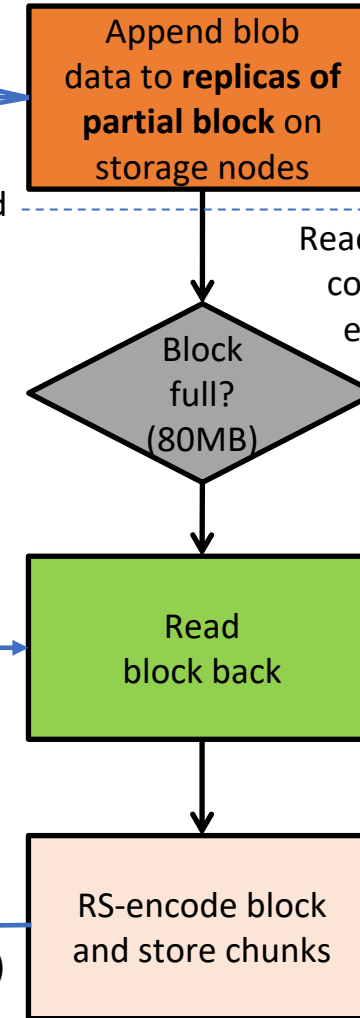
Blob storage

Blob sizes are small
(100s of KBs)

Blobs appended to log
structured file

Blobs need to be persisted
before acknowledging
upload

Minimize latency
for blob uploads,
Later optimize
storage



3-way replicated

Read-after-write
consistency on
every append


RS(9,6)

RS(10,4)

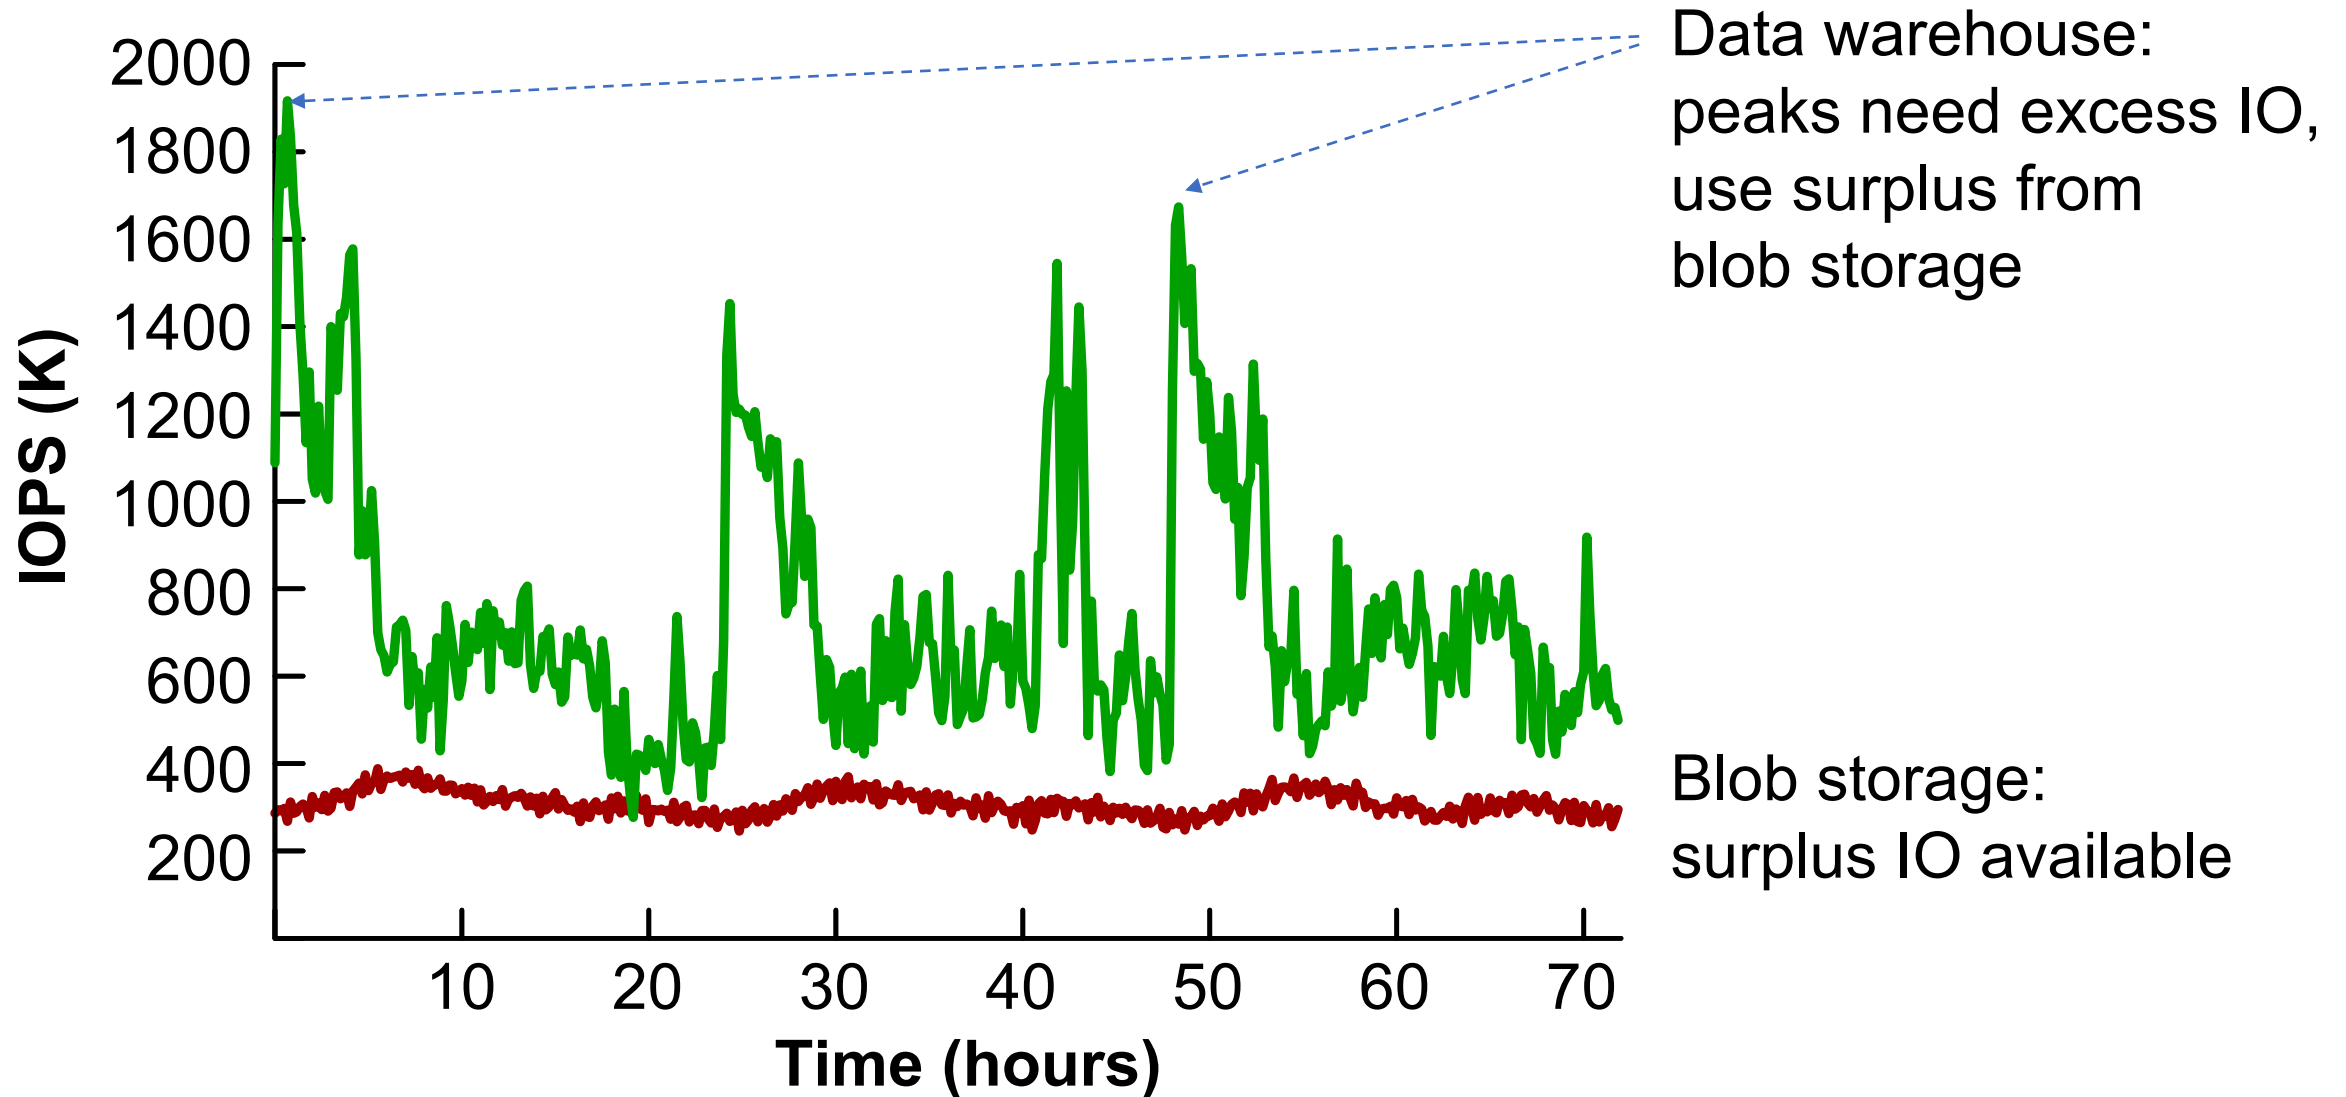
Results

- Tectonic clusters are ~10x the size of HDFS clusters, which simplifies production operations
- Blob storage latency in Tectonic comparable to Haystack
- In a multitenant cluster, data warehouse uses surplus IO from blob storage to serve its peaks

Efficiency From Storage Consolidation



Efficiency From Storage Consolidation



Tectonic Provides Datacenter-Scale Storage

- Replaced previous constellation of specialized storage systems
 - Simpler operations
 - Better resource utilization
- Tectonic's design addresses the key challenges:
 - Scalability: disaggregated linearly scalable components
 - Performance: tenant-specific optimizations via client library
 - ...

Thank You