

Experiences with a Distributed Deduplication API

Fred Douglass

Andy Huber

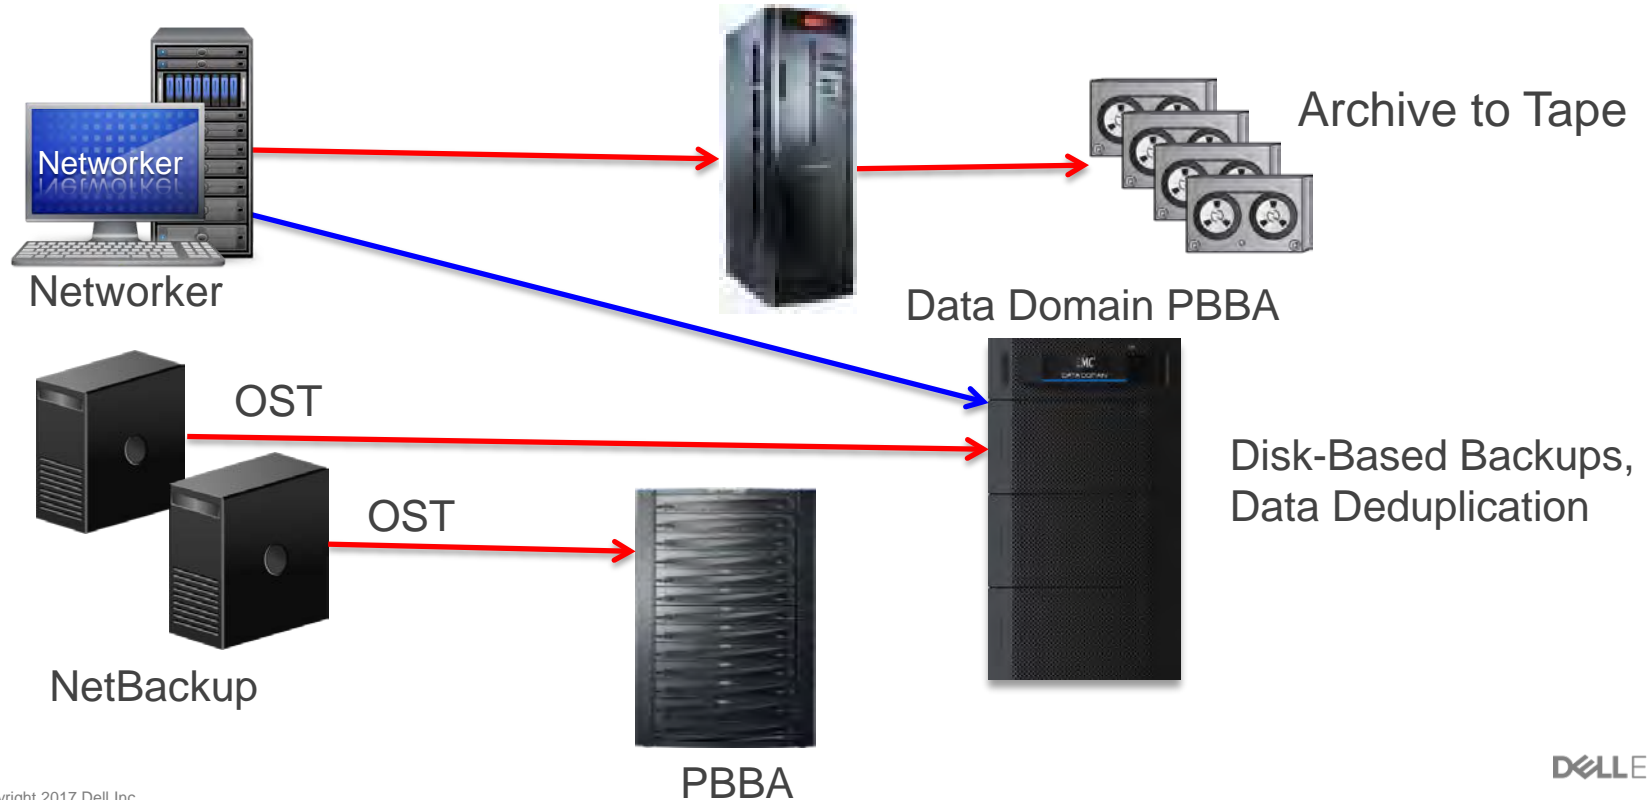
Donna Lewis

Rachel Traylor

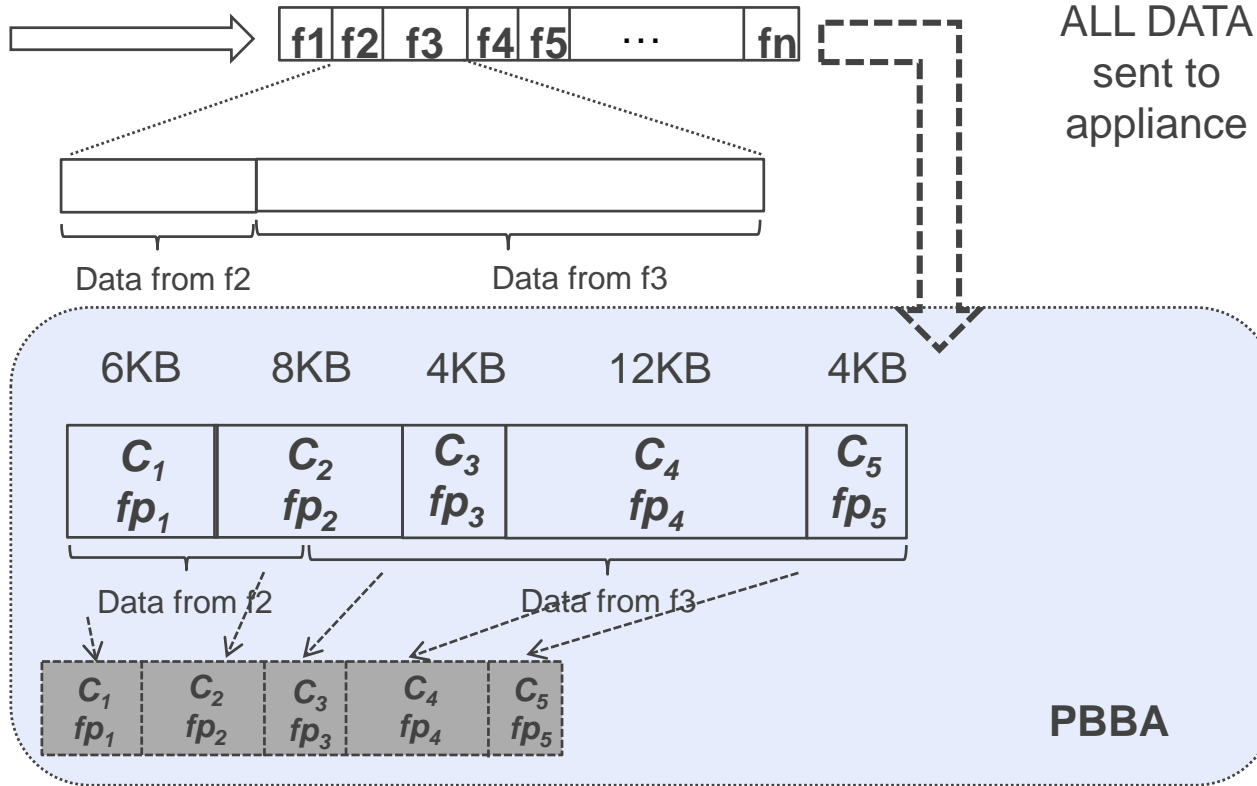



Background: Traditional Backup Model

Daily Full Backup, Periodic Full + Daily Incremental



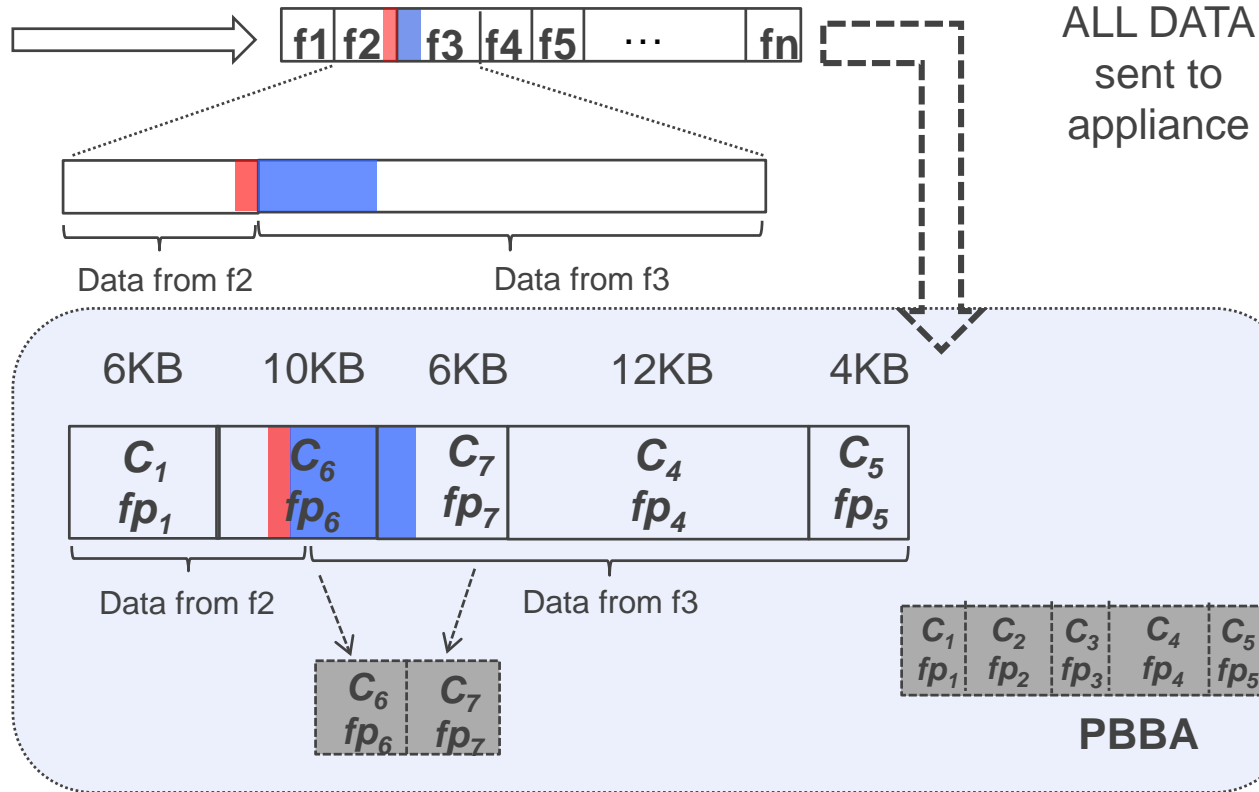

NFS Processing



Chunks,
fingerprints and
filters data

Compresses and
stores new data

NFS Processing - Subsequent Full



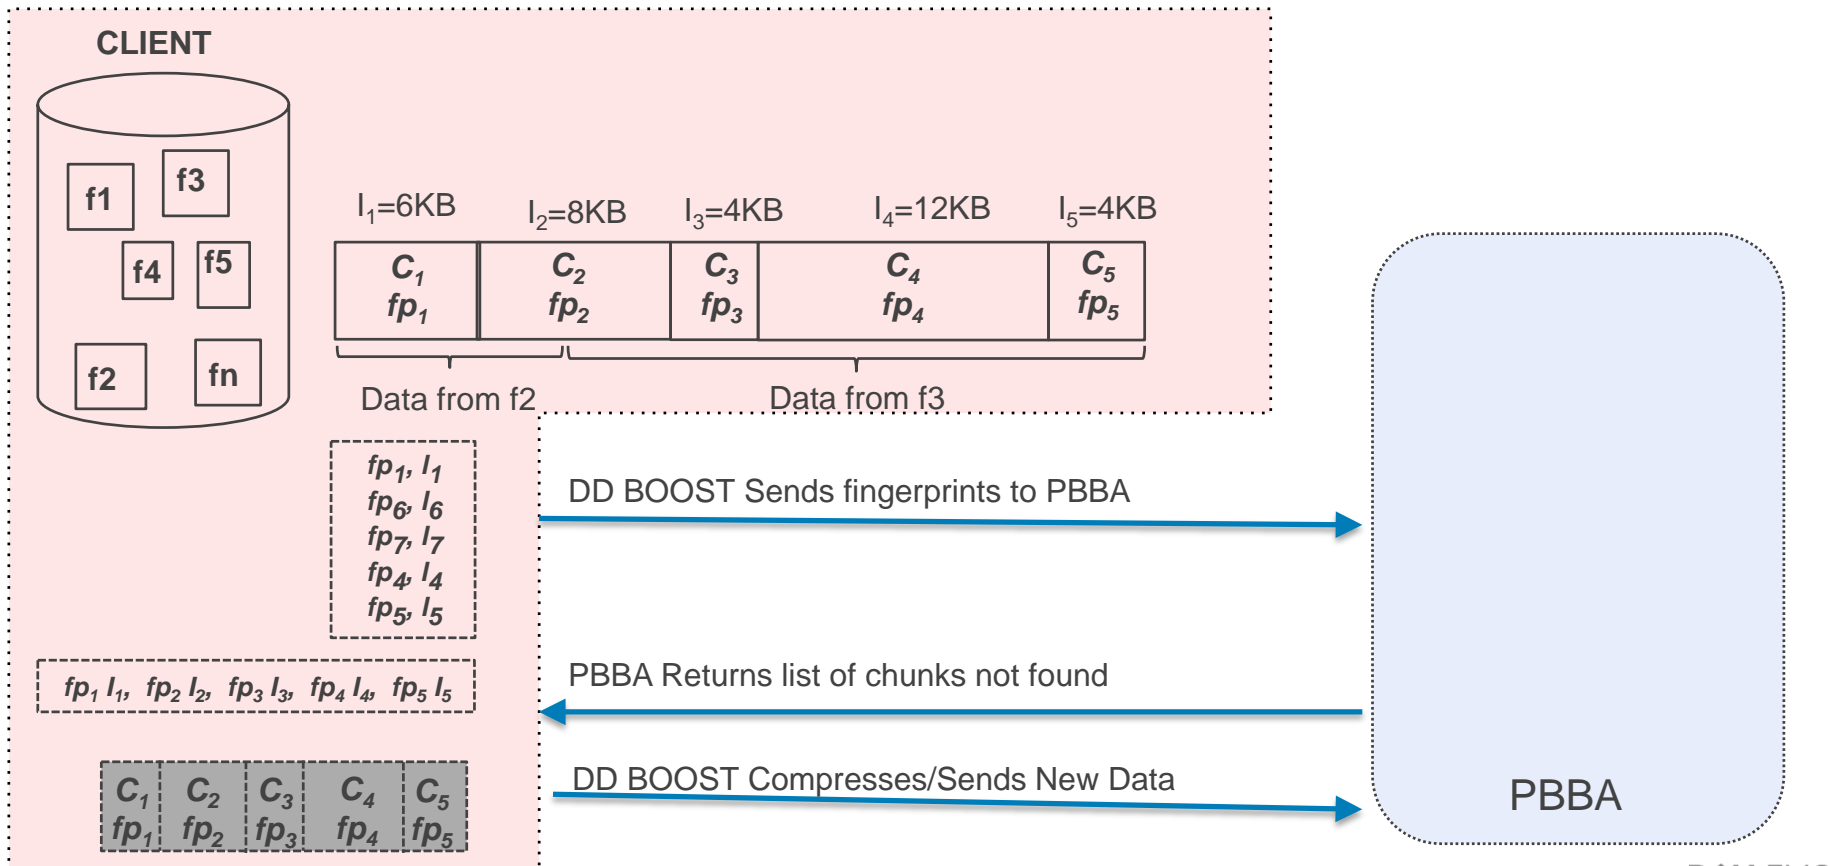
Chunks,
fingerprints and
filters data

Compresses and
stores only new
data

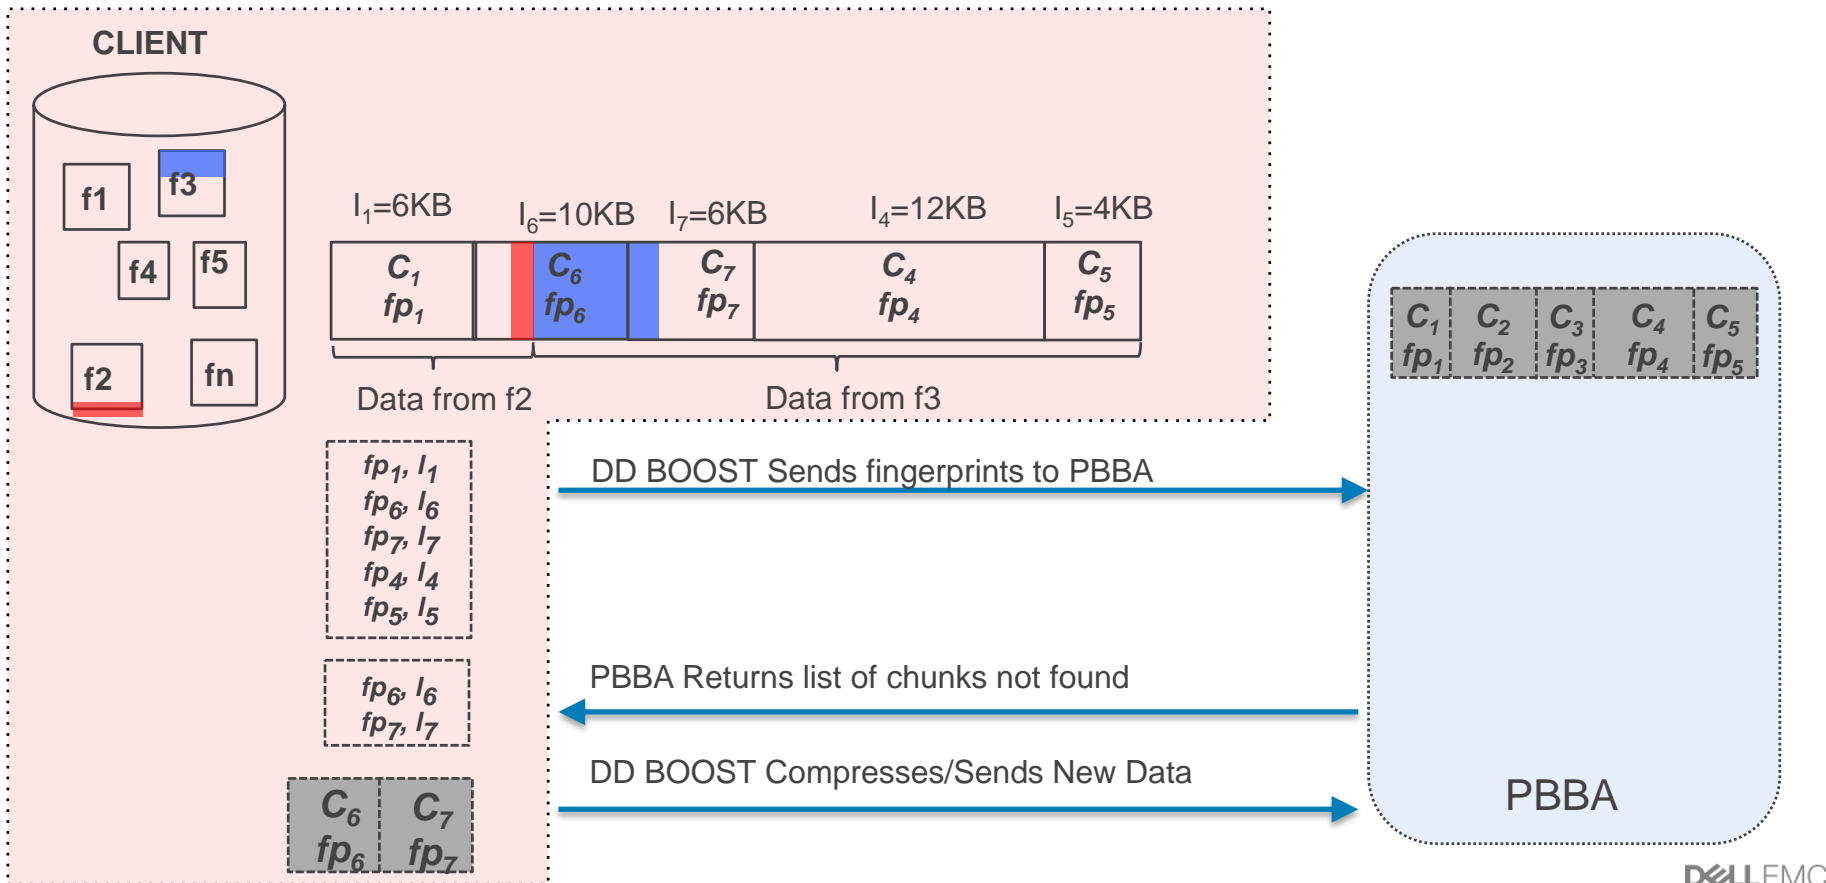
DD BOOST Distributed Deduplication

- Offloads Chunk and Fingerprint processing to the client
 - Reduces resource load on the PBBA significantly
- Filtering occurs between DD BOOST and the Data Domain PBBA
- Only new chunks and fingerprints are sent for storage, saving network bandwidth
 - Significantly reduces bandwidth

DD BOOST Distributed Deduplication



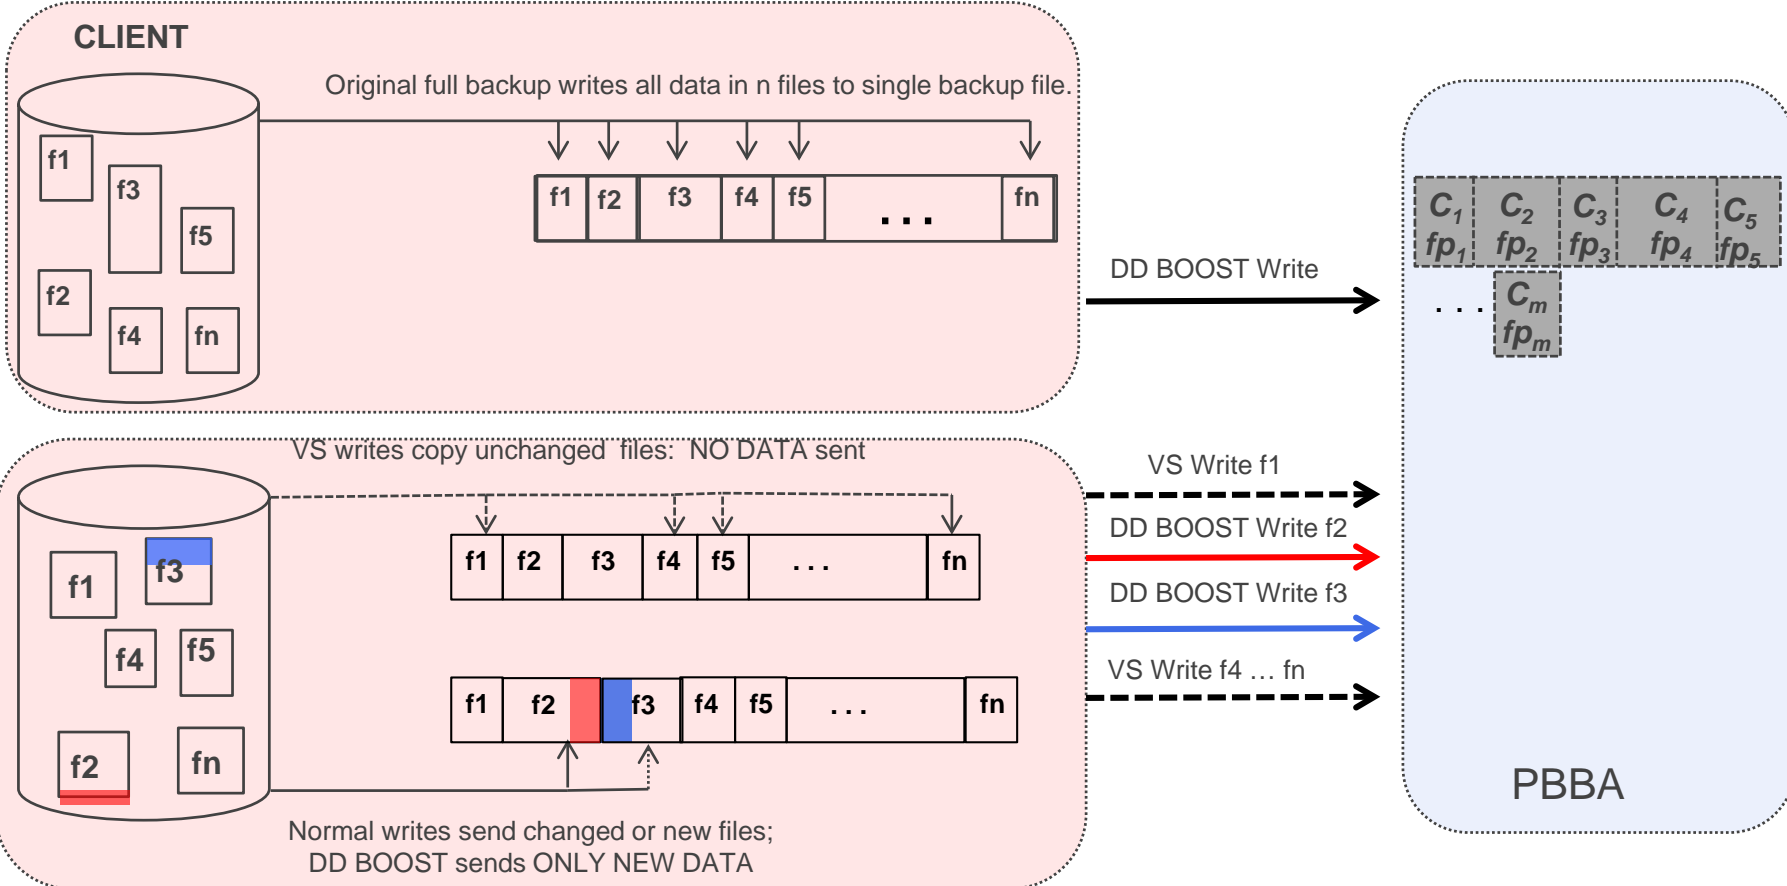
DD BOOST Distributed Deduplication



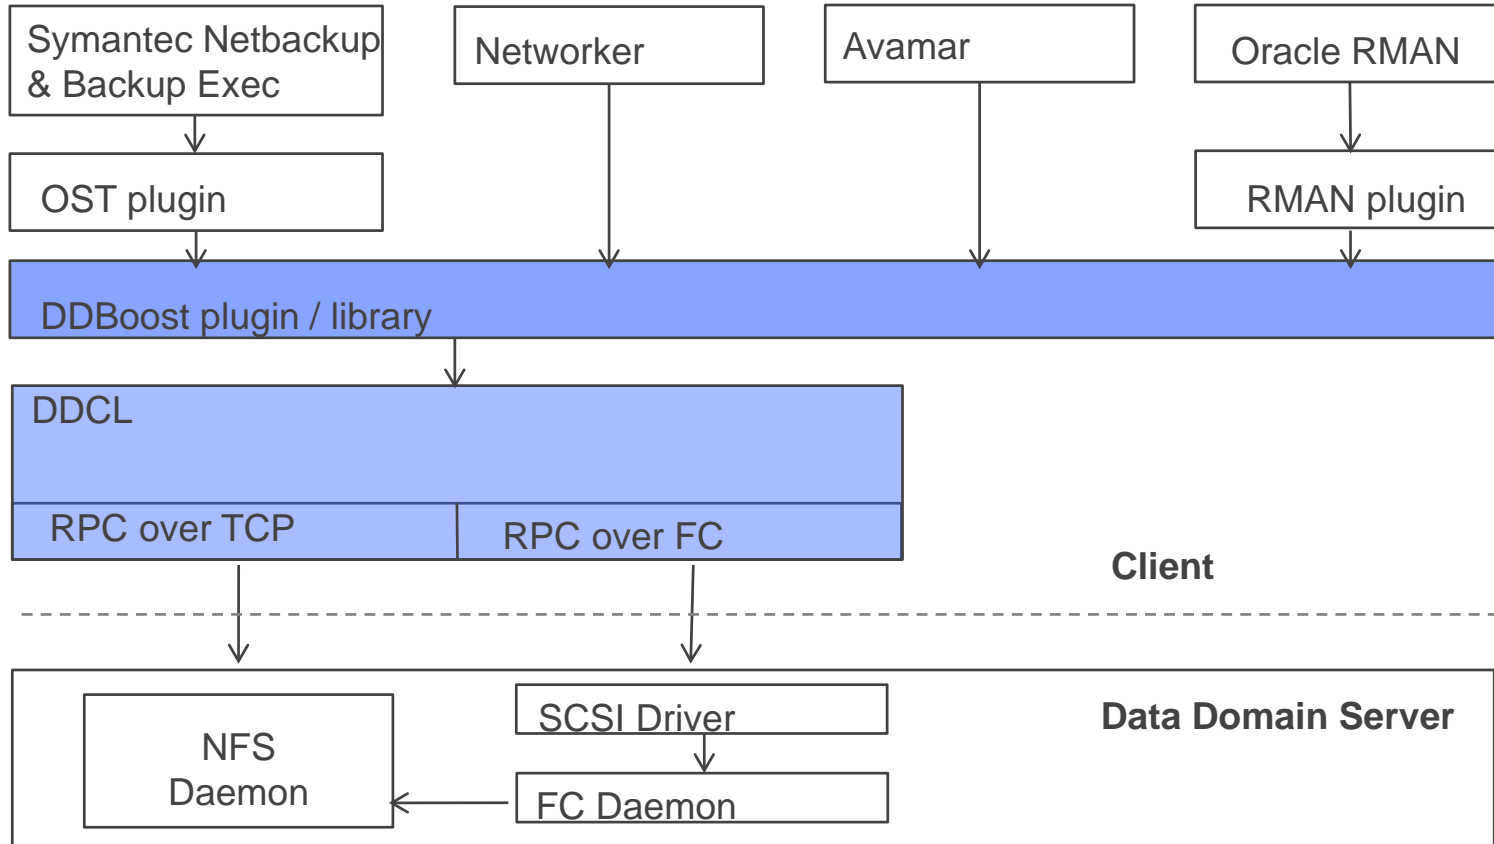
DD BOOST Virtual Synthetics

- Leverages change information known to the application
 - Identifies regions of unchanged data
- Method of directing PBBA to copy references from previous backups to the current one, data isn't sent
 - Significant bandwidth savings
 - Most effective with large regions
- New data processed using DD BOOST
 - Additional savings when interleaving

DD BOOST: Virtual Synthetics



DD BOOST Architecture



Evaluation

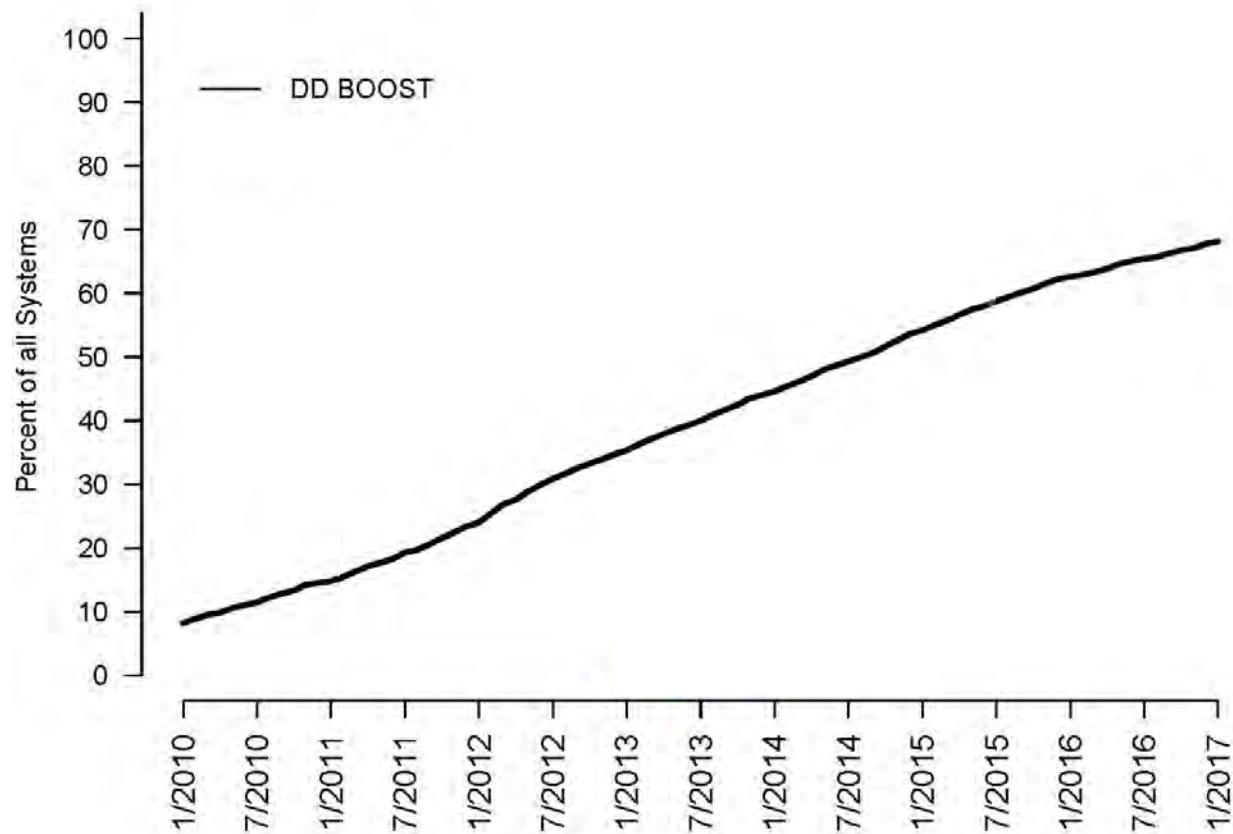
Analyzed Customer Telemetry

- Adoption Rate
- Bandwidth Savings

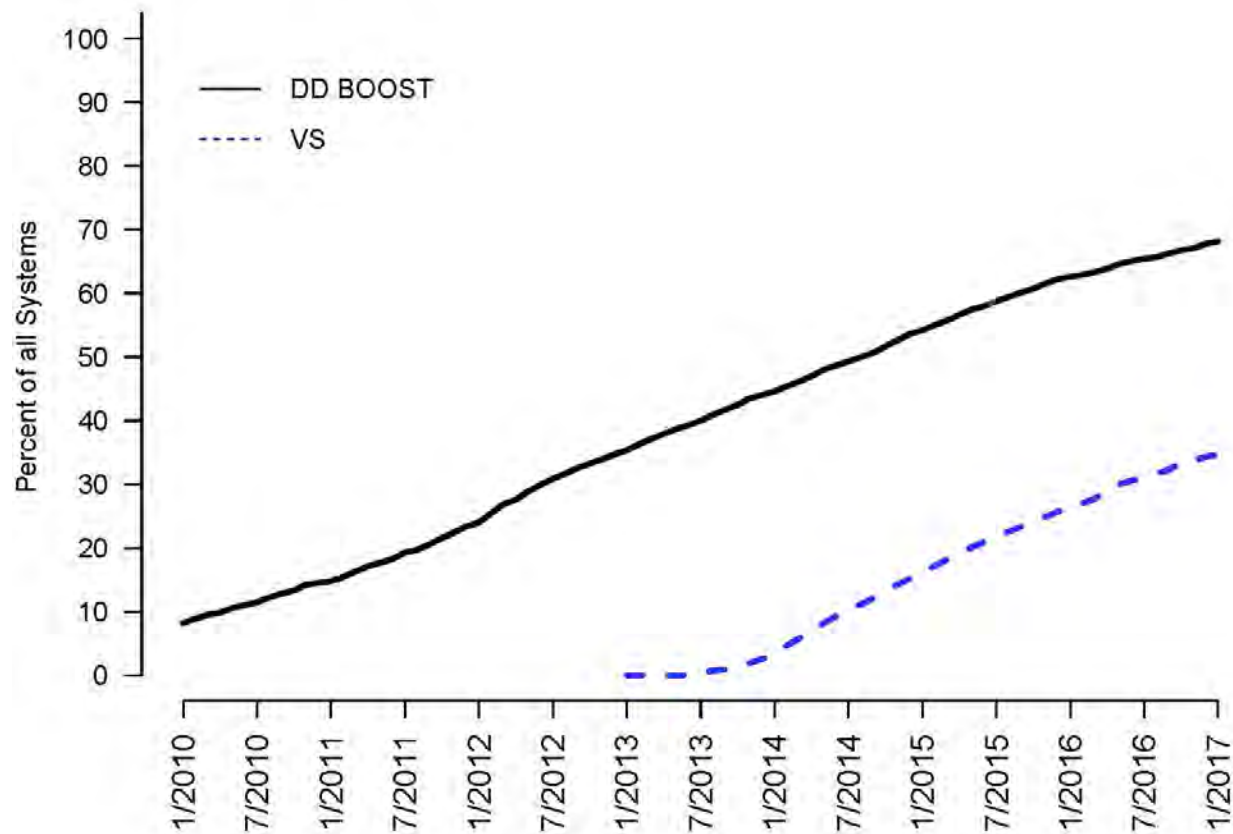
Performance Benchmarks

- DD BOOST and NFS
- DD BOOST and DD BOOST Virtual Synthetics

DD BOOST Adoption




DD BOOST Adoption




DD BOOST Bandwidth Savings

No Virtual Synthetics



Systems using Virtual Synthetics



Performance Evaluation Method

Setup:

- Internal Load Generator – DD BOOST, NFS, Virtual Synthetics
- Parallel Stream Processing (1, 8, 16, 32, 64)
- Streams divided evenly across four systems
- 10 Gb/sec Connections

Generations

Gen 0 – Initial Backup “Generation 0”

Low-gen – First Changed Backup, Generation 1

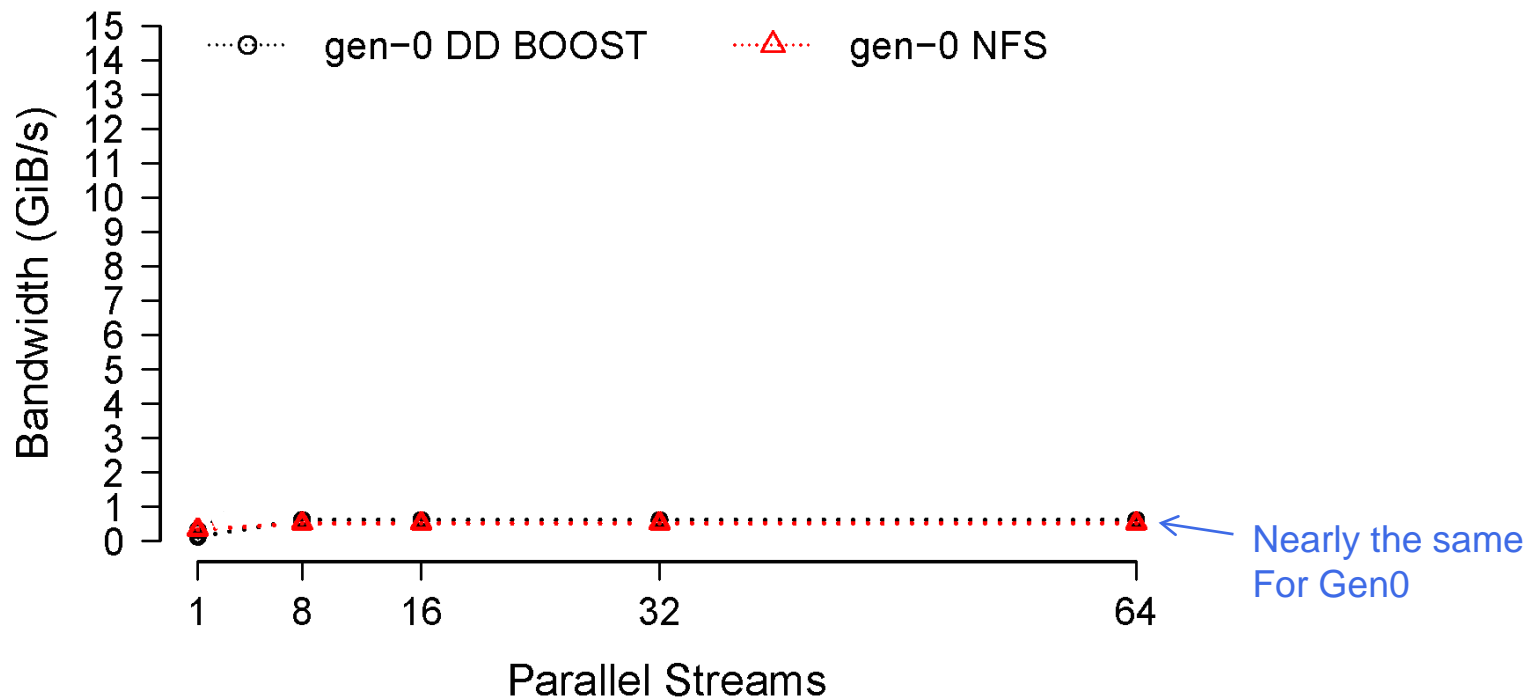
High-gen – Generation 41

Performance Evaluation Method – Change Rate

- Normal Distribution
 - Concentrates updates more frequently in a smaller range of data
 - Validated against customer datasets
 - One cluster of added data, one cluster of deleted, one cluster of modified per 1GB of data written.
 - 1% added, 1% deleted, 3% modified
- Uniform Distribution
 - Used in previous studies
 - Changes are distributed uniformly throughout the file
 - Useful for performance benchmarking release over release
 - Ineffective measurement for Virtual Synthetics

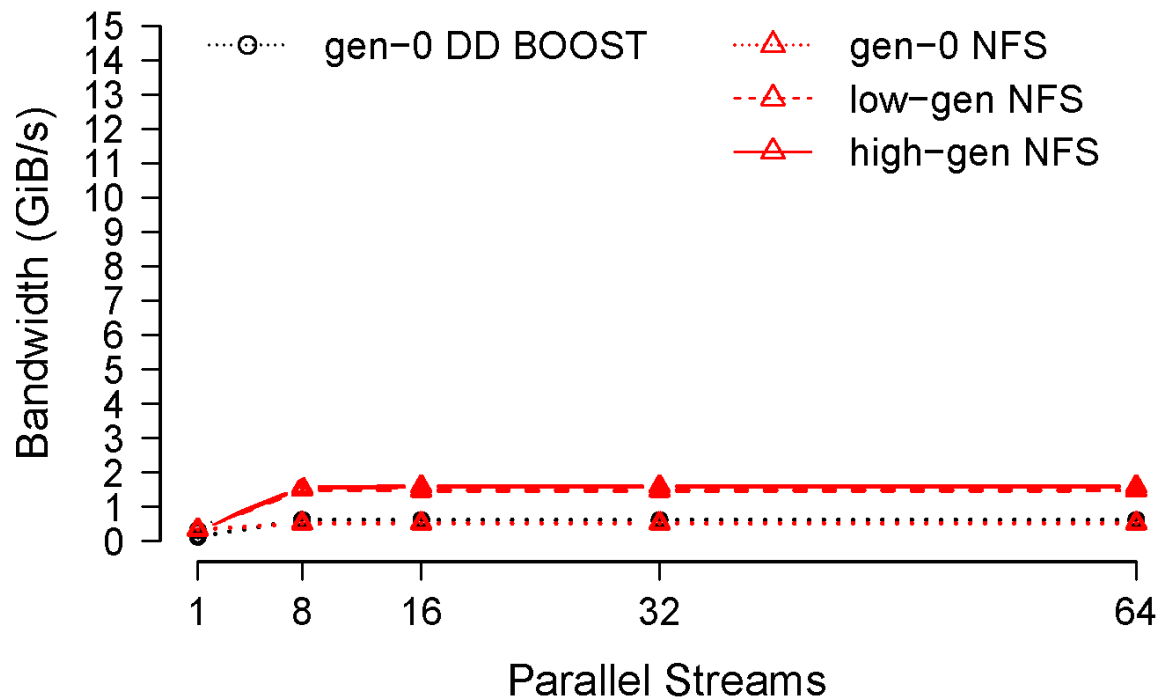
Performance: DD BOOST vs NFS

Backup Measurements



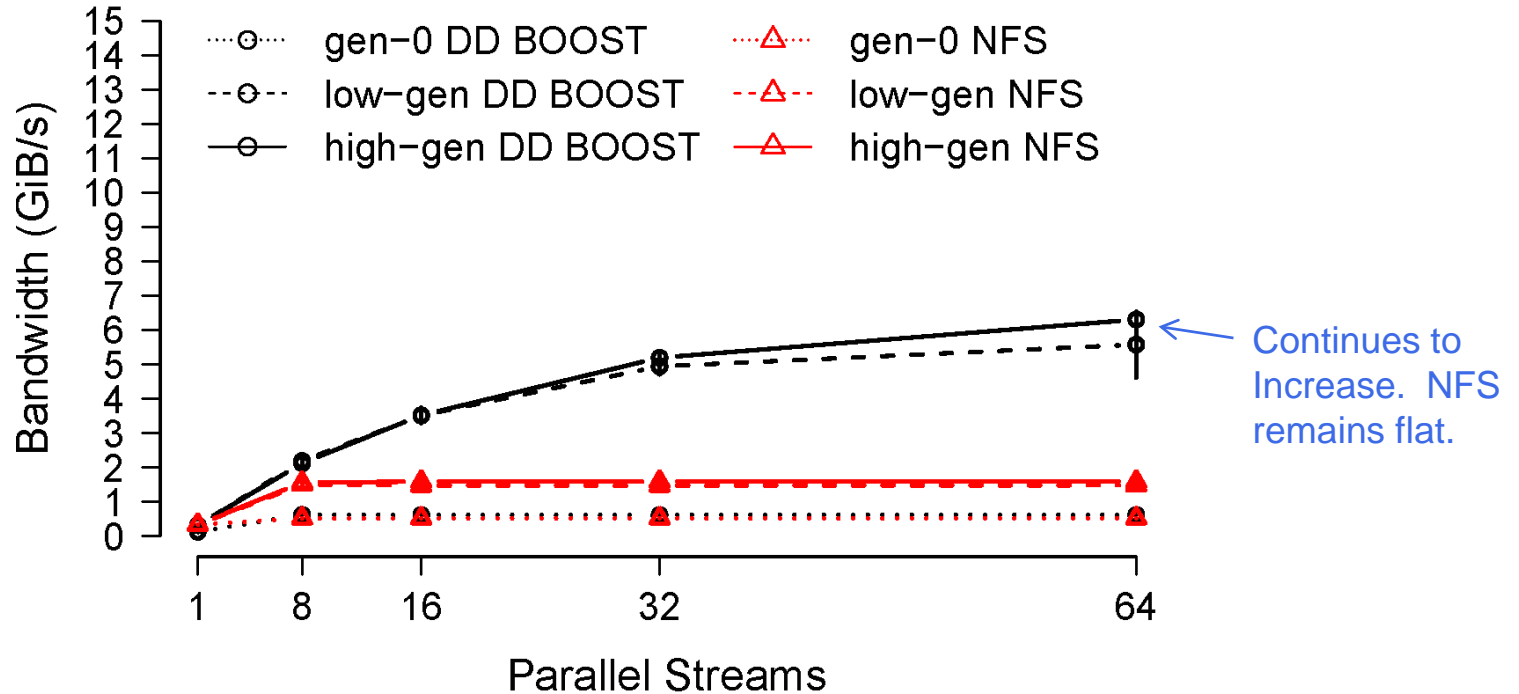
Performance: DD BOOST vs NFS

Backup Measurements



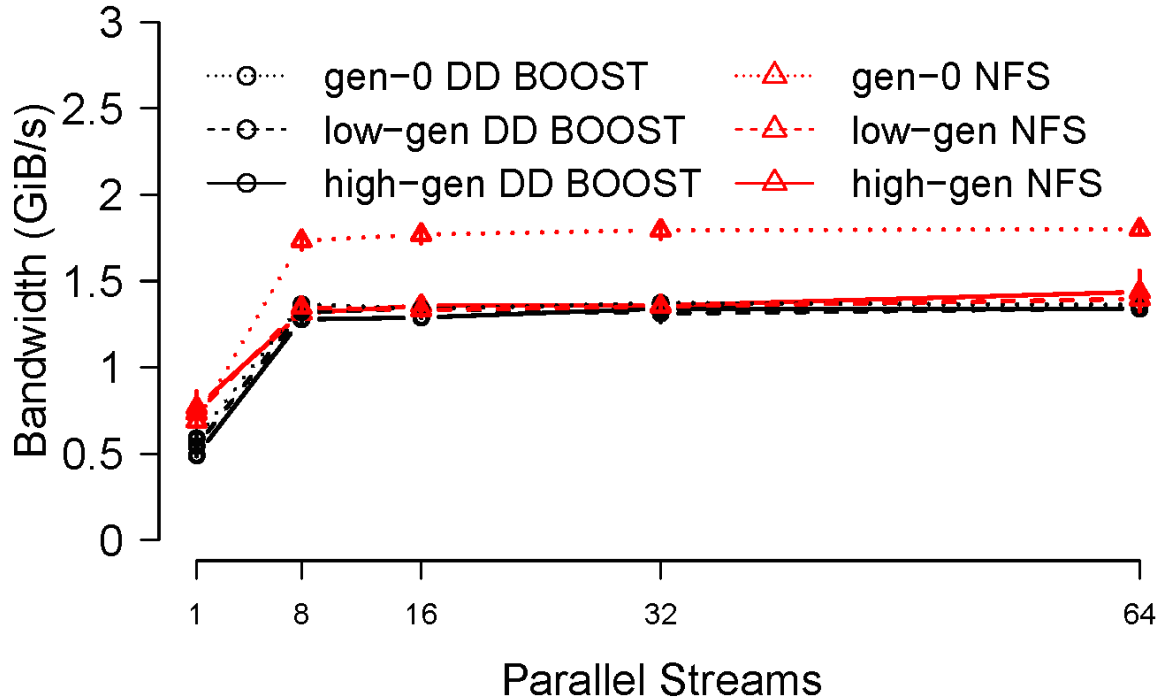
Performance: DD BOOST vs NFS

Backup Measurements



Performance: DD BOOST vs NFS


Restore Measurements



DD BOOST has a limited read size, but applies read-ahead logic


Performance: DD BOOST and Virtual Synthetics

Normal Distribution



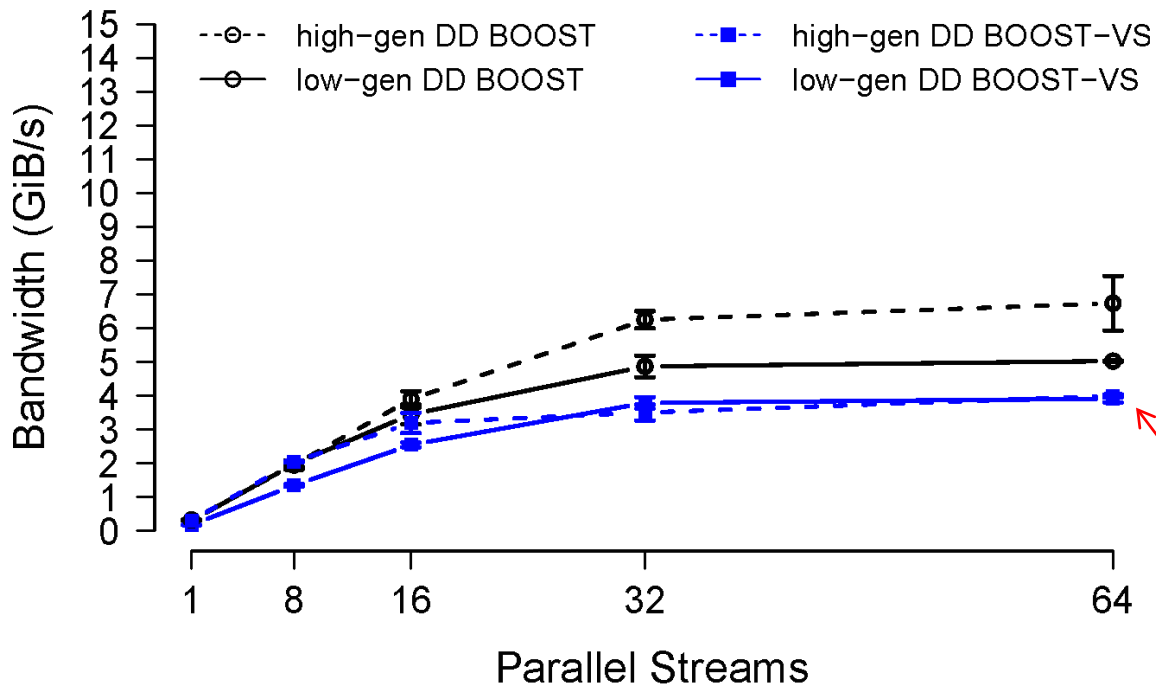
Performance: DD BOOST and Virtual Synthetics

Normal Distribution



Performance: DD BOOST and Virtual Synthetics


Uniform Distribution




Greater overhead if synthesizing smaller more distributed regions

Performance: DD BOOST and Virtual Synthetics

Normal Distribution



Uniform Distribution



DD BOOST : Lessons Learned

Deduplication Obstacles

- Compression and Encryption
- In-line Transformations

File Usage

- Smaller files and File Management Operations

Virtualization

- DD BOOST's limited resource requirements are an ideal fit

Related Work

- Low-bandwidth Network File System (LBFS)
 - Similar chunking and fingerprinting strategy
 - Targeted at reducing network traffic for file operations
- Rsync
 - Reduces network bandwidth when syncing a directory
 - Embedded deduplication with network transport
 - Similar approaches: DOT, czip, and Jumbo Store
- Deduplication optimizations and tradeoffs

Conclusion and Ongoing Work

- Customer Telemetry shows significant bandwidth savings
- Virtual Synthetics further reduces bandwidth, allowing customers to manage full backups while only writing changes
- Steady increase in adoption

Ongoing Work:

BoostFS – FUSE-based offering to eliminate integration effort

First Linux version released in Fall 2016

D~~E~~LL EMC

Additional DD BOOST Functionality

- Lightweight load-balance/failover mechanism
- Fastcopy or Clone
- Per File Replication
- Compare, Return Differences
- Snapshot Storage Unit
- Control Token Based Authentication
- Set Extended Attributes
- Retrieve Data Movement Statistics