

Understanding Primary Storage Optimization Options

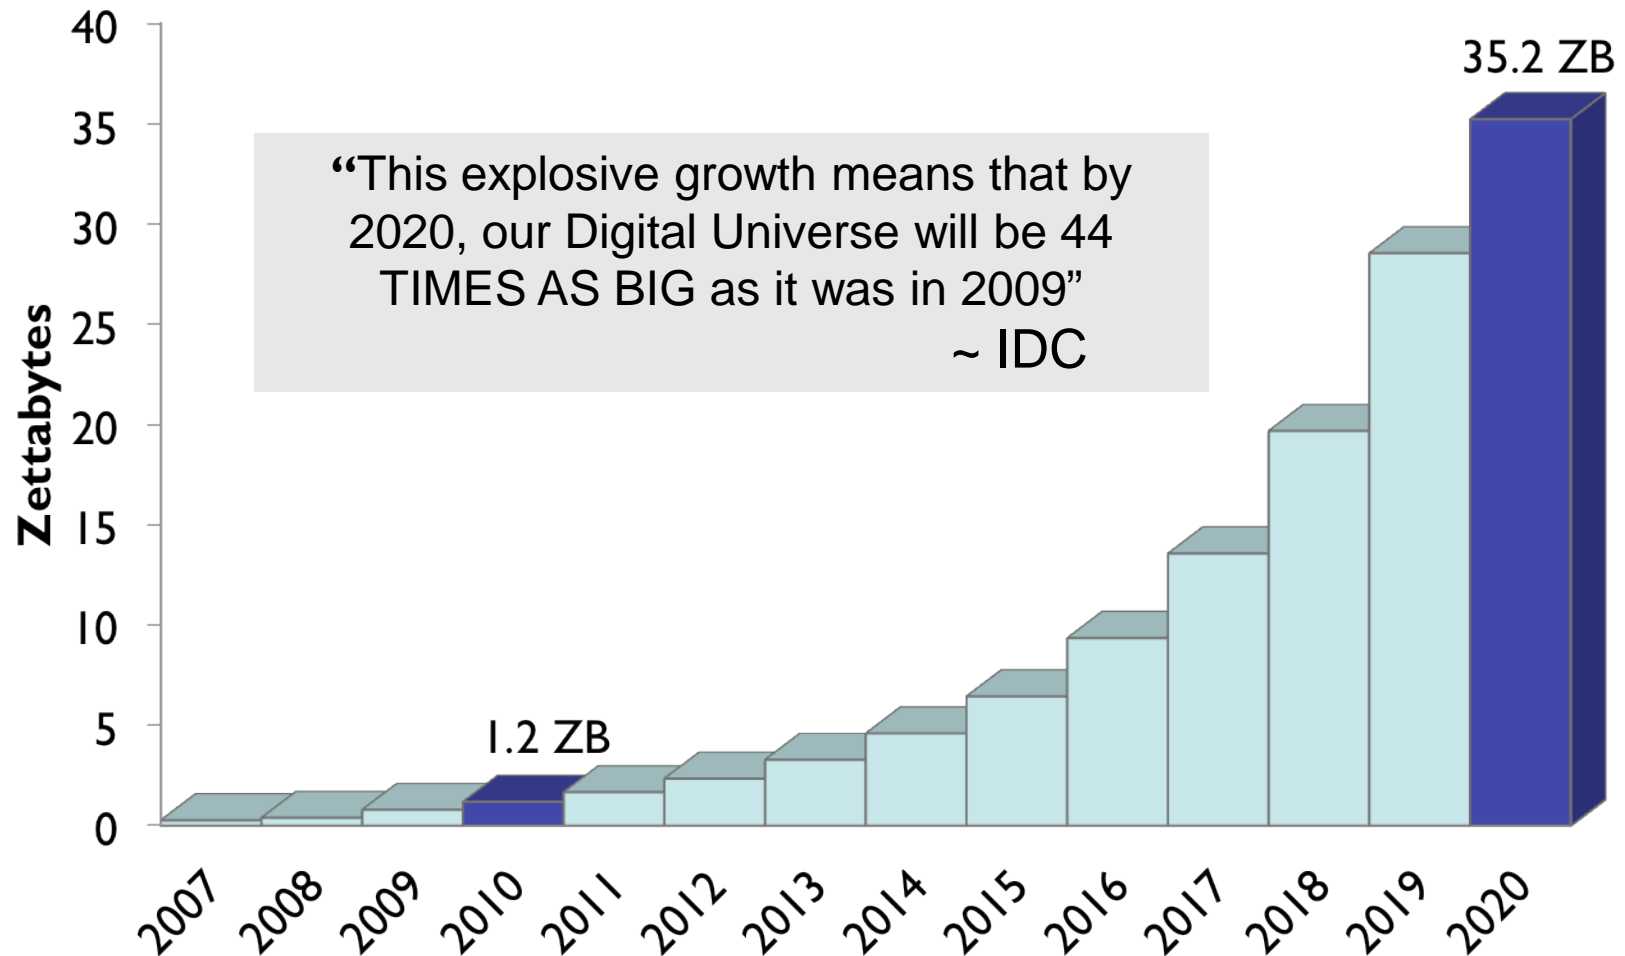
Jered Floyd

Permabit Technology Corp.

Primary Storage Optimization

- ❑ Technologies that let you store more data on the same storage
 - ❑ Thin provisioning
 - ❑ Copy-on-write snapshots
 - ❑ Compression
 - ❑ Deduplication

Data Growth is Accelerating



Source: IDC Digital Universe Study, May 2010

1 ZB = 1 Trillion Gigabytes

Requirements for Primary Storage Optimization

- ❑ To be broadly adopted, any technology must:
 - ❑ Support block, file and unified architectures
 - ❑ Have no impact on performance
 - ❑ Have no impact on reliability
 - ❑ Scale to support storage capacity deployed
 - ❑ Implement within existing architecture
- ❑ Older methods (e.g. thin provisioning) meet these requirements today
- ❑ Newer methods are beginning to emerge

Compression and Deduplication

- ❑ For that reason, we'll focus on compression and deduplication
- ❑ Both have been in backup “forever”
 - ❑ Compression – 40 years
 - ❑ Deduplication – 10 years
- ❑ Compression for primary storage, while available for many years, has never really been enabled
- ❑ Deduplication is relatively new to primary storage

2011 is The Year of Primary Storage Optimization

Musings on the future of data dedupe

One thing is clear: In 2011, the focus will shift from deduplication for nearline / secondary storage to deduplication for primary storage.

~Dave Simpson

INFOSTOR

Data storage trends 2011:

Predictions of hot data storage technologies

The other one that I think is really important, and we're just beginning to see this come out now, is data deduplication and compression for primary storage.

~Tony Asaro



Hot technologies for 2011


Primary storage data reduction is back from our 2010 Hot Technologies list. In 2011, we'll see a lot more of primary data reduction in shipping products.



Storage Industry Consolidation And Deduplication

In short customers want a single deduplication method that works across platforms. It's the only logical way to leverage deduplication so its full advantages can be realized.

~George Crump



- ❑ Techniques to reduce the size of stored data
 - ❑ Lossy vs. lossless
 - ❑ Generic vs. content aware
 - ❑ Identifying repeated bytes
 - ❑ Identifying duplicate bytes
 - ❑ Identifying similar bytes or objects
 - ❑ Discarding irrelevant data
- ❑ Operate on a single file / object at a time

Data Compression Benefits

- ❑ Mature, generic algorithms
- ❑ Big savings on low entropy data (e.g. text)
- ❑ Big savings on rich media
- ❑ Broad hardware support for specific technologies
- ❑ Low memory requirements

Data Compression Challenges

- ❑ High processor requirements
- ❑ No cross-object savings
- ❑ Complex licensing on media formats
- ❑ Always in the data read path
- ❑ Modifies byte stream on storage
- ❑ Savings constant regardless of data scale

Data Deduplication

- ❑ Conceptually a sort of compression...
- ❑ Techniques to eliminate data being stored
 - ❑ Single-instance vs. sub-file
 - ❑ Fixed block vs. variable block
 - ❑ Generic vs. content aware
 - ❑ Always lossless
- ❑ Operates across a file system, LUN, or entire storage pool

Data Deduplication Benefits

- ❑ Big savings on redundant data (e.g. VM, database)
- ❑ Lower CPU requirements than compression
- ❑ No impact on data read
- ❑ Underlying data isn't modified
- ❑ Savings scale with more data stored

Data Duplication Challenges

- ❑ Higher memory requirements
- ❑ Limited applicability for media files
- ❑ No standardization for software implementations
- ❑ Limited scale in most solutions

Deduplication: Backup vs. Primary

- ❑ Primary workflows require massive scale and high performance
- ❑ Backup data model allows for simplifications not applicable to primary storage
 - Bloom filters are unfeasibly computationally expensive for random-access deletion
 - Differencing methods require larger blocks than primary dedupe allows
 - Buffering for large look-back window
 - Locality knowledge to individual sources
 - Large block similarities
- ❑ Backup dedupe doesn't adapt to primary storage use case

	Backup	Primary
Data Flow	Stream-Oriented	Random Access
Latency Critical	No	Yes
Typical Chunk Size	1 MB and up	4 KB to 64 KB
Index Lookups	Thousands/sec	Millions/sec
# Objects	100s Millions	100s Billions
Unique Data	1 to 60 TB	1 TB to PBs

Gartner Priority Matrix for Storage Technologies

benefit	years to mainstream adoption			
	less than 2 years	2 to 5 years	5 to 10 years	more than 10 years
transformational		"Big Data" and Extreme Information Processing and Management Data Deduplication Enterprise-Grade Solid-State Drives Thin Provisioning		
high	iSCSI Serial Attached SCSI: HDD and SSD Interface Wide-Area SANs	Continuous Data Protection Data Encryption Technologies, HDDs and SSDs Enterprise Information Archiving Metro-Area SANs Quality of Storage Service Storage Resource Management TCP/IP Offload Engines and iSCSI	Cloud-Oriented Object-Based Storage Cross-Platform Structured Data Archiving Massive Array of Idle Disks Online Data Compression Storage Cluster File Systems Virtual Machine Recovery	pNFS
moderate	Distributed Tape for Backup and Archiving	Automatic Storage Tiering Bare-Metal Restore Content-Addressed Storage Controller-Based I/O Prioritization Distributed Virtual Tape Fibre Channel Over Ethernet File Virtualization Technologies Heterogeneous External Storage Virtualization SAN Management Server-Based Replication Storage Multitenancy	Appliance-Based Replication Emerging Data Protection Schemes Enterprise Endpoint Backup Hybrid Storage Gateway Appliance Information Dispersal Algorithms Public Cloud Storage Storage Bridge Bay	
low			Switch-Based Storage Virtualization	

As of July 2011

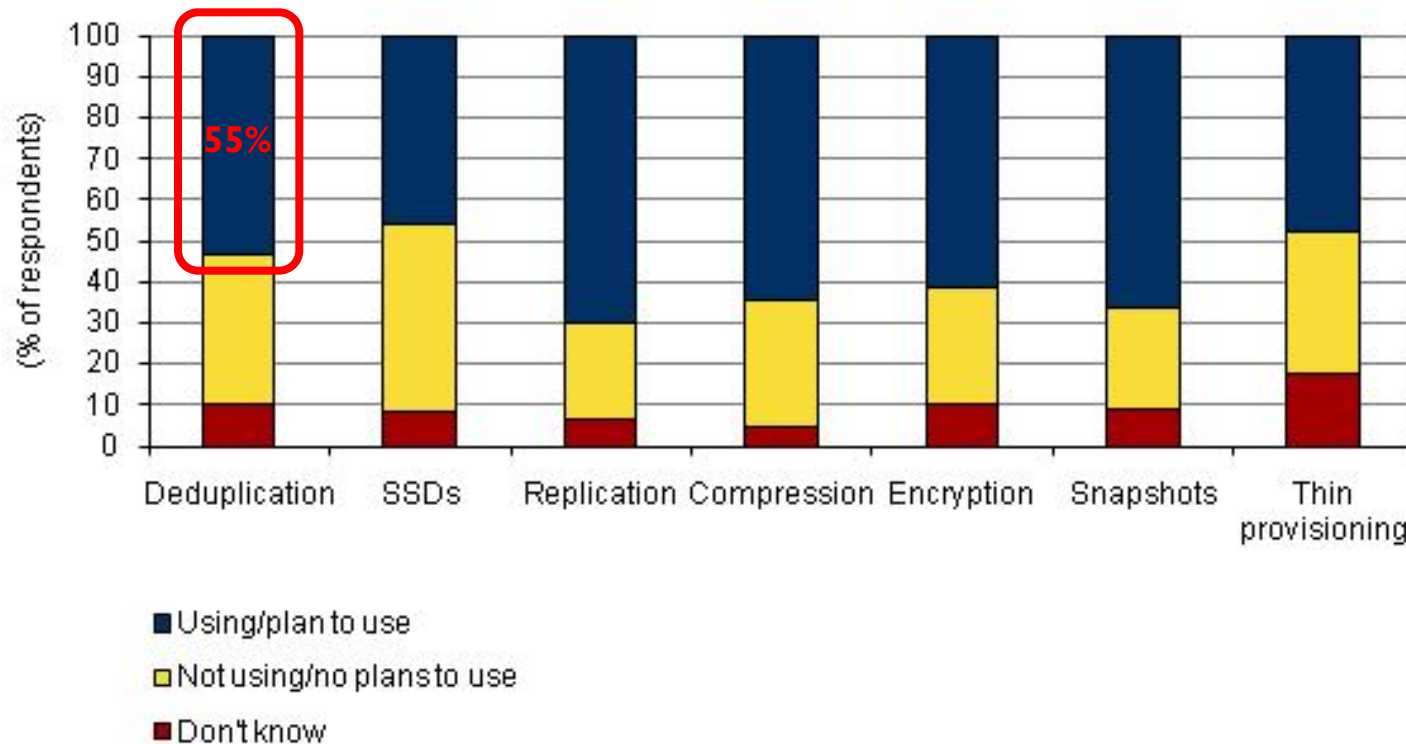
"Big Data" and Extreme Information Processing and Management
Data Deduplication
Enterprise-Grade Solid-State Drives
Thin Provisioning

Deduplication identified as a "transformational" storage technology

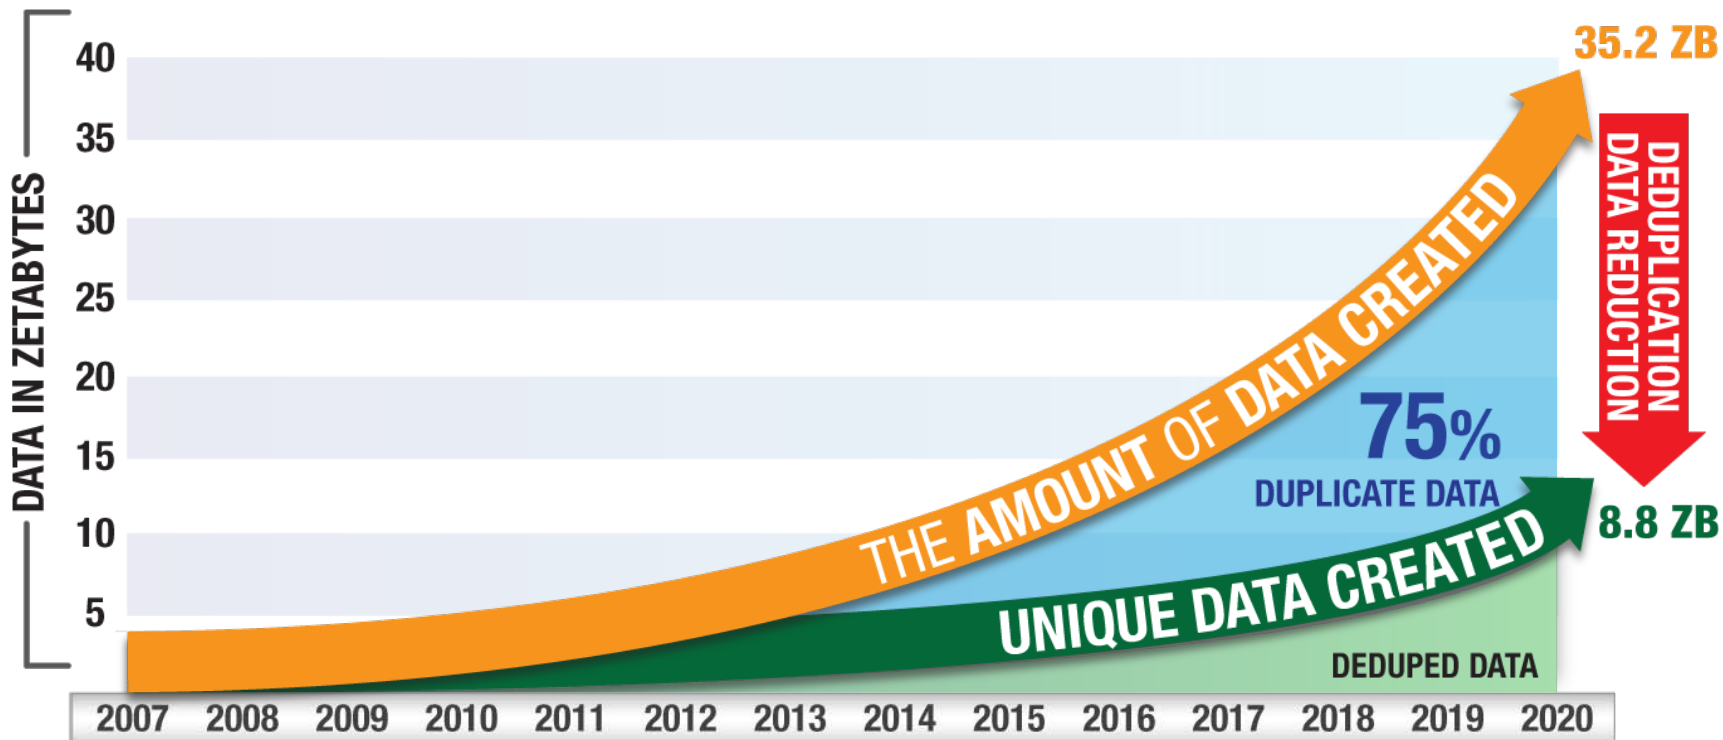
Source: Gartner Hype Cycle for Storage Technologies, 2011
 26 July 2011 | ID:G00214638

IDC 2010 Trends in Storage and Virtualized Environments

Advanced Storage Technologies Deployment



Deduplication Impact



Source: IDC Digital Universe Study, May 2010

1 ZB = 1 Trillion Gigabytes

Compression vs. Deduplication

	Compression	Deduplication
Impact on Read	High	None
Impact on Write	High	Moderate
Savings on VM	Low	High
Savings on Media	High (some formats)	Low
CPU Requirements	High	Moderate
Memory Requirements	Low	High
Scalability	Unlimited	Varies by Implementation
Impact on Reliability	Moderate	Low


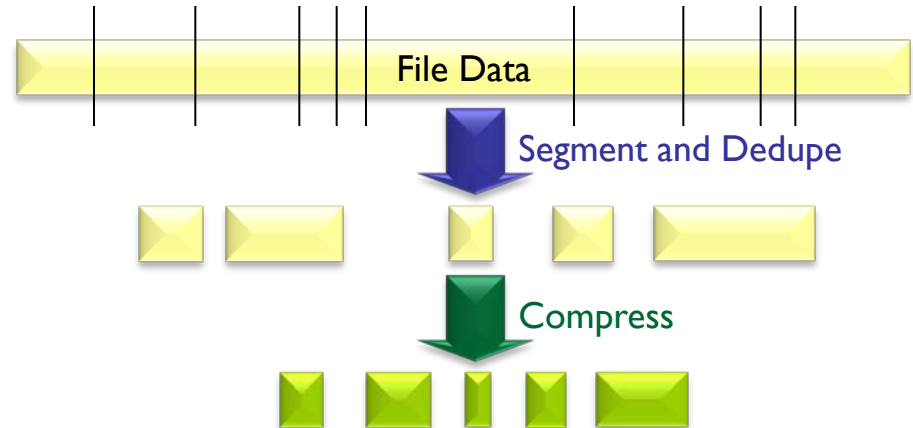
Compression and Deduplication

- ❑ Compression identifies "micro" duplicates

- ❑ Dedupe identifies "macro" duplicates

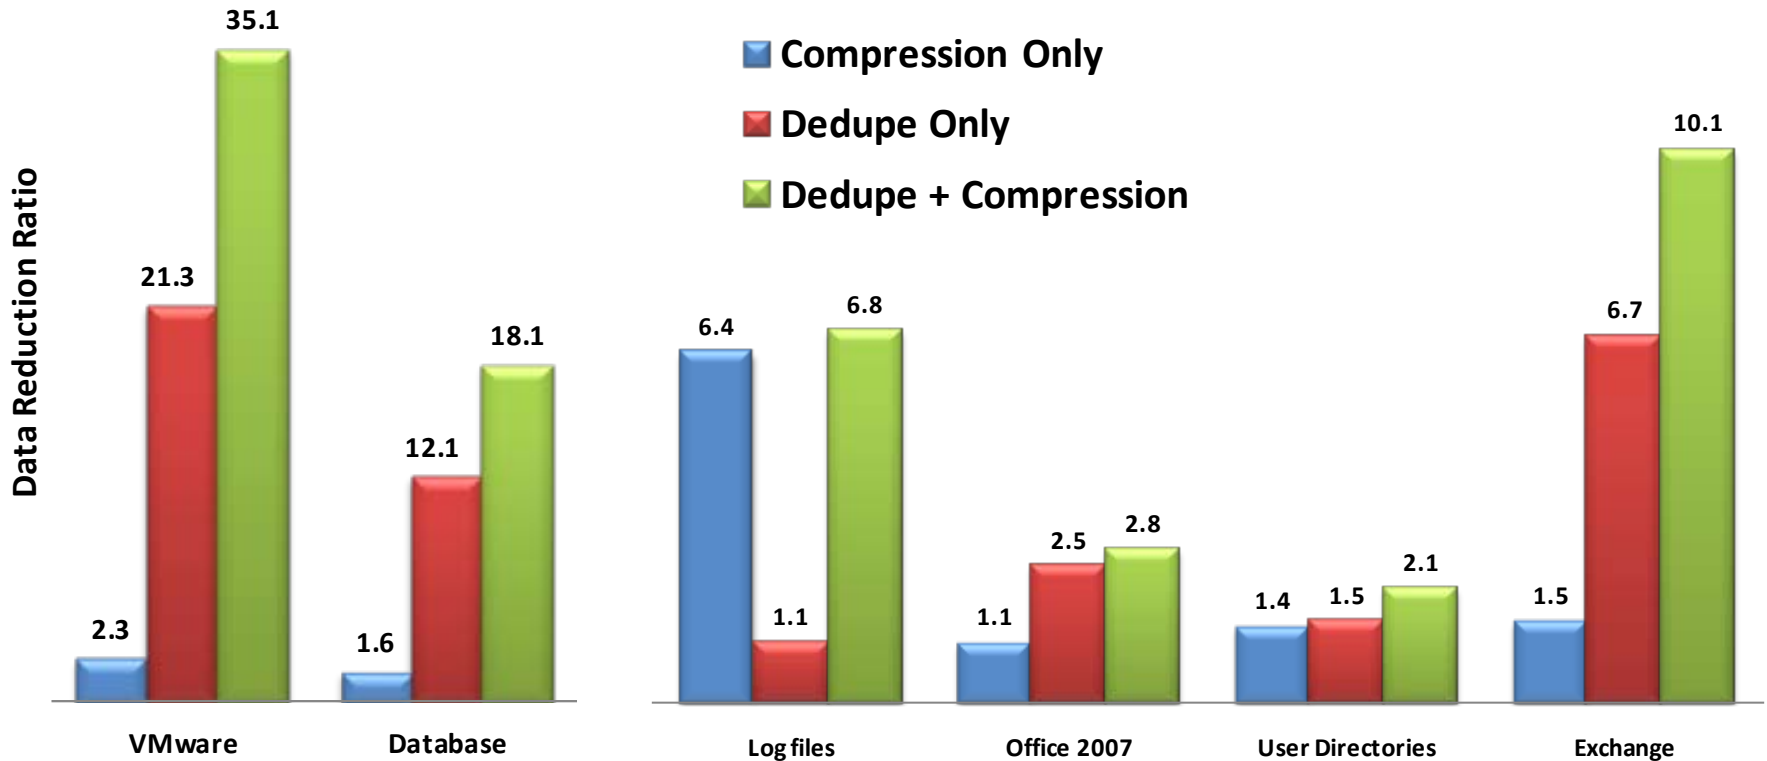
- ❑ Dedupe then compress (easiest), or

- ❑ Compress then dedupe (requires compressed format segmentation)



Compression and Deduplication

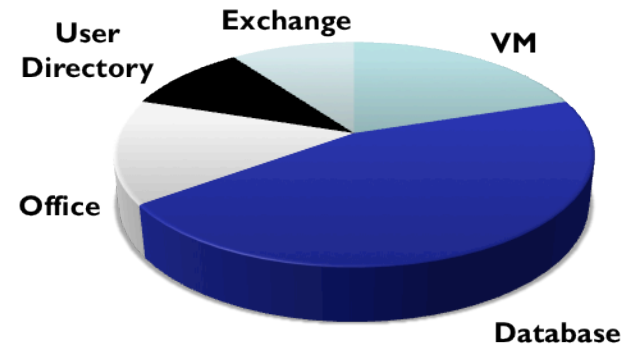
- Compression (e.g. LZ) identifies "micro" duplicates
- Deduplication identifies "macro" duplicates



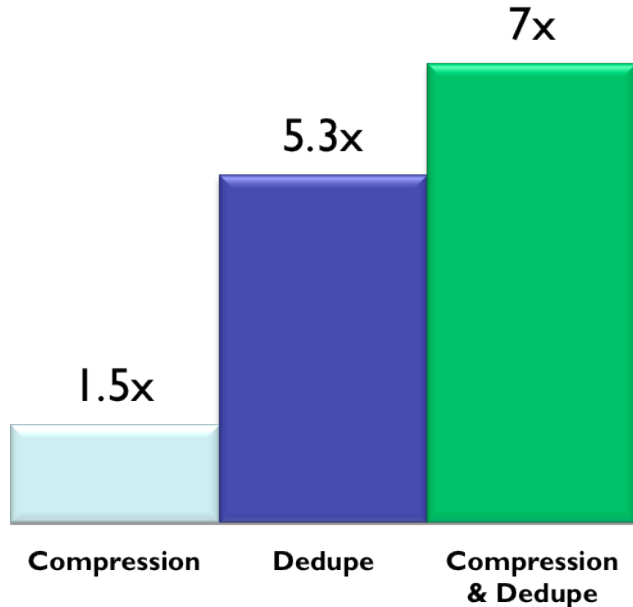
Primary Data Efficiency Impact

- For IPB Enterprise Primary environment

1 PB Enterprise Mix



Reduction Rate



Conclusion:

- Compression is good
- Dedupe is better: >3x data reduction over wider data set range
- Dedupe + Compression is best of breed


Deploying Primary Optimization

- ❑ Where does it run?
 - ❑ Integrated into Storage
 - ❑ Intermediary Appliance
 - ❑ Host Software

- ❑ When
 - ❑ Inline
 - ❑ Post-process
 - ❑ Parallel

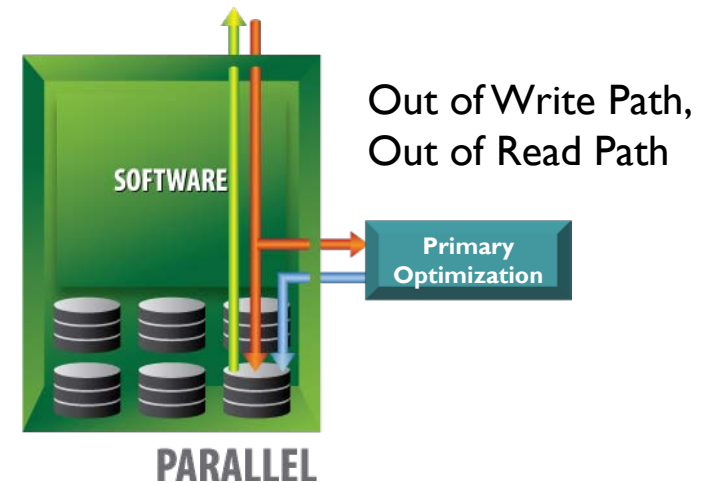

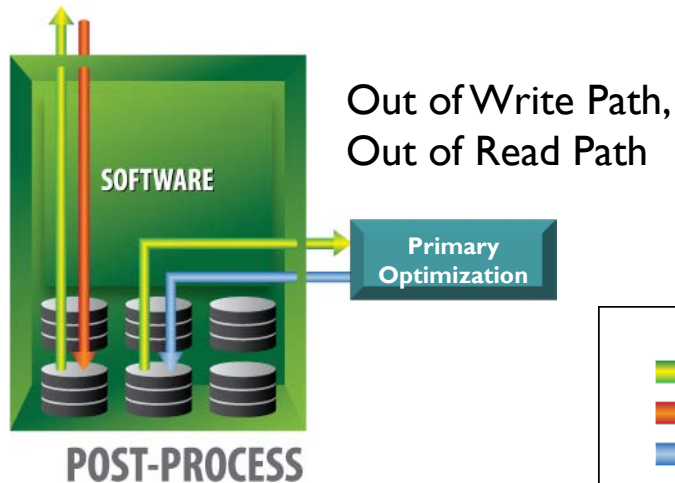
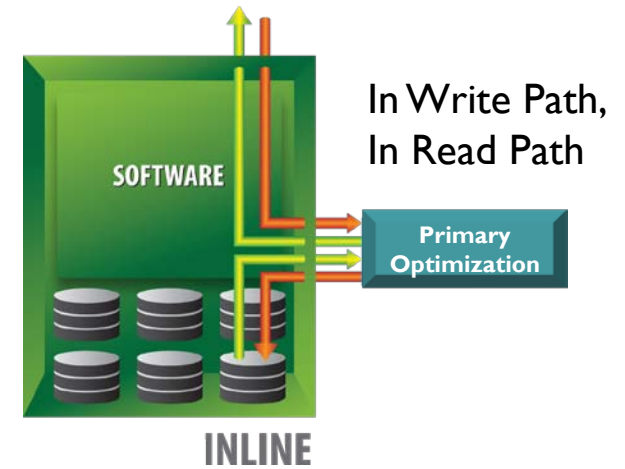
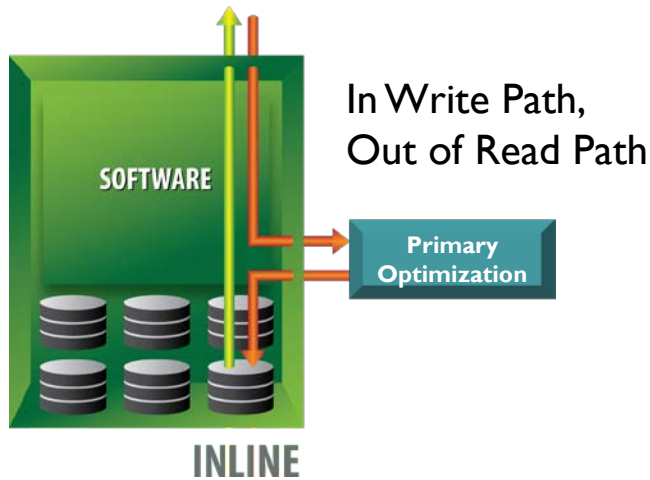
Intermediary Appliance

- ❑ Storage optimization runs on a separate hardware device
- ❑ All data passes through the appliance on read and write
- ❑ Benefits:
 - ❑ Brings storage optimization to legacy platforms
- ❑ Challenges:
 - ❑ Additional hardware expense
 - ❑ Introduces bottleneck to all I/O operations
 - ❑ Appliance can mask functionality
 - ❑ Failure can affect availability
 - ❑ Data lock-in to optimization appliance technology



- ❑ Storage optimization takes place on application host
- ❑ All data passes through the software on read and write
- ❑ Benefits:
 - ❑ Brings storage optimization to legacy platforms
 - ❑ No additional hardware cost or complexity
- ❑ Challenges
 - ❑ Difficult to implement with shared storage
 - ❑ Consumes host CPU and memory resources
 - ❑ Data lock-in to specific optimization technology

Deploying Deduplication



- ❑ Data continues to grow exponentially
- ❑ Primary Storage Optimization technologies save you money
 - ❑ (This includes thin provisioning if you're not doing that yet)
- ❑ Compression + Deduplication is best
- ❑ Different integration models have different tradeoffs
 - ❑ Cost
 - ❑ Performance
 - ❑ Data Savings
 - ❑ Operational impact (availability, reliability, etc.)
- ❑ In the end, optimization will move into the storage