facebook

Facebook's disaggregated storage and compute for Map/Reduce

Yun Jin

Intro

Facebook’s Monthly Active Users

Grew by 1.1B since 2010 monthly active users


Era of the the Desktop
Advance of the Mobile
Growth in Messenger
Videos vs pictures
360 video
Number of users in millions
1,750
1,500
1,250
1,000
750
500
250
0
Q3 '08 Q1 '09 Q2 '09 Q3 '09 Q4 '09 Q1 '10 Q2 '10 Q3 '10 Q4 '10 Q1 '11 Q2 '11 Q3 '11 Q4 '11 Q1 '12 Q2 '12 Q3 '12 Q4 '12 Q1 '13 Q2 '13 Q3 '13 Q4 '13 Q1 '14 Q2 '14 Q3 '14 Q4 '14 Q1 '15 Q2 '15 Q3 '15
© Statista 2016

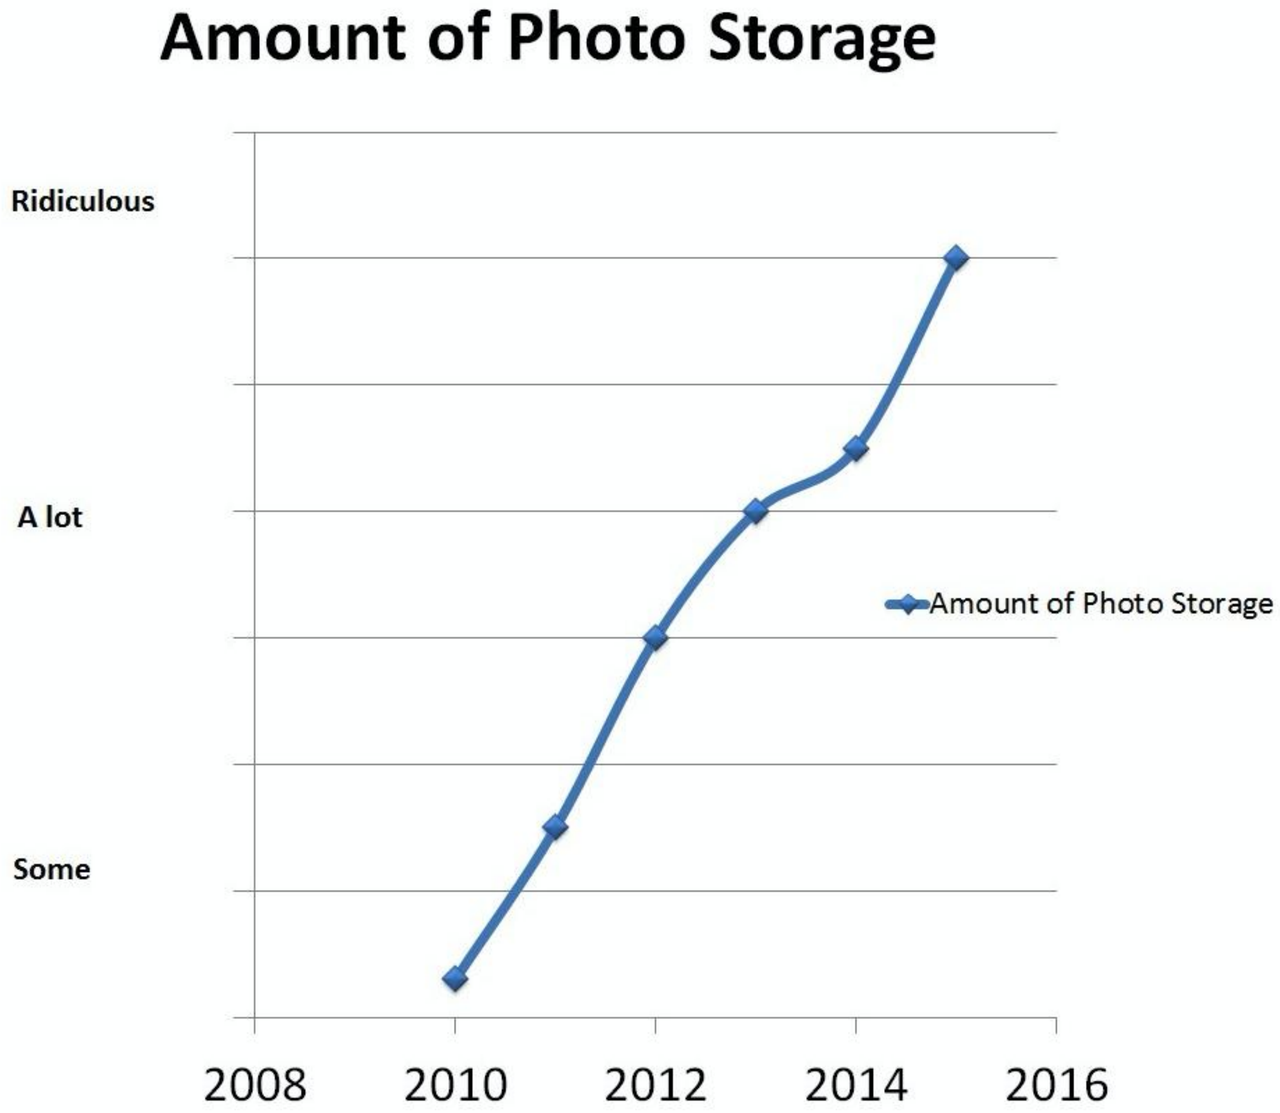
Amount of Photo Storage
Ridiculous
A lot
Some
Amount of Photo Storage
2008
2010
2012
2014
2016


Era of the the Desktop
Advance of the Mobile
Growth in Messenger
Videos vs pictures
360 video
Number of users in millions
1,750
1,500
1,250
1,000
750
500
250
0
Q3 '08
Q1 '09
Q2 '09
Q3 '09
Q4 '09
Q1 '10
Q2 '10
Q3 '10
Q4 '10
Q1 '11
Q2 '11
Q3 '11
Q4 '11
Q1 '12
Q2 '12
Q3 '12
Q4 '12
Q1 '13
Q2 '13
Q3 '13
Q4 '13
Q1 '14
Q2 '14
Q3 '14
Q4 '14
Q1 '15
Q2 '15
Q3 '15
© Statista 2016

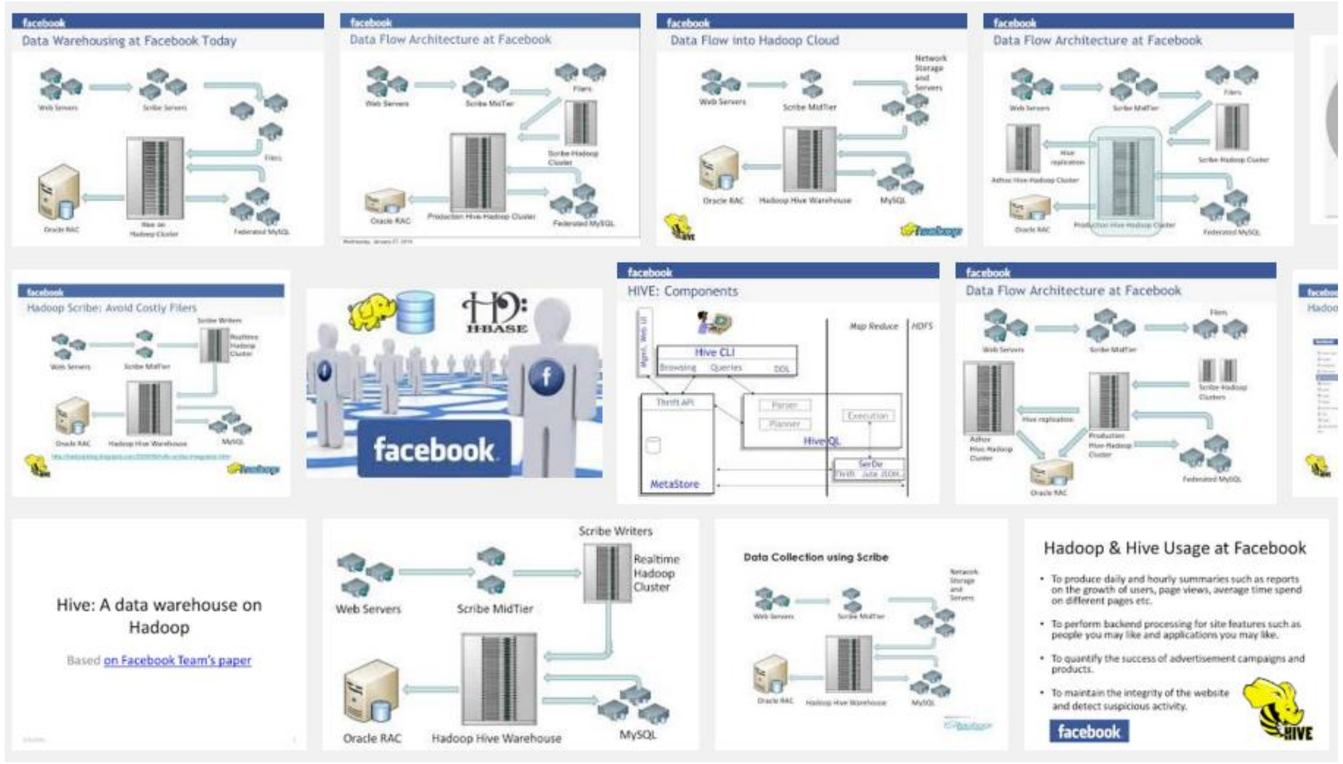
facebook
Data Warehousing at Facebook Today
facebook
Data Flow Architecture at Facebook
facebook
Data Flow into Hadoop Cloud
facebook
Data Flow Architecture at Facebook
facebook
Hadoop Scribe: Avoid Costly Filers
facebook
HIVE: Components
facebook
Data Flow Architecture at Facebook
Hive: A data warehouse on Hadoop
Based on Facebook Team's paper
Scribe Writers
Realtime Hadoop Cluster
Web Servers
Scribe MidTier
Oracle RAC
Hadoop Hive Warehouse
MySQL
Data Collection using Scribe
Hadoop & Hive Usage at Facebook
• To produce daily and hourly summaries such as reports on the growth of users, page views, average time spend on different pages etc.
• To perform backend processing for site features such as people you may like and applications you may like.
• To quantify the success of advertisement campaigns and products.
• To maintain the integrity of the website and detect suspicious activity.
facebook
HIVE


Era of the the Desktop
Advance of the Mobile
Growth in Messenger
Videos vs pictures
360 video
Number of users in millions
1,750
1,500
1,250
1,000
750
500
250
0
Q3 '08 Q1 '09 Q2 '09 Q3 '09 Q4 '09 Q1 '10 Q2 '10 Q3 '10 Q4 '10 Q1 '11 Q2 '11 Q3 '11 Q4 '11 Q1 '12 Q2 '12 Q3 '12 Q4 '12 Q1 '13 Q2 '13 Q3 '13 Q4 '13 Q1 '14 Q2 '14 Q3 '14 Q4 '14 Q1 '15 Q2 '15 Q3 '15
© Statista 2016

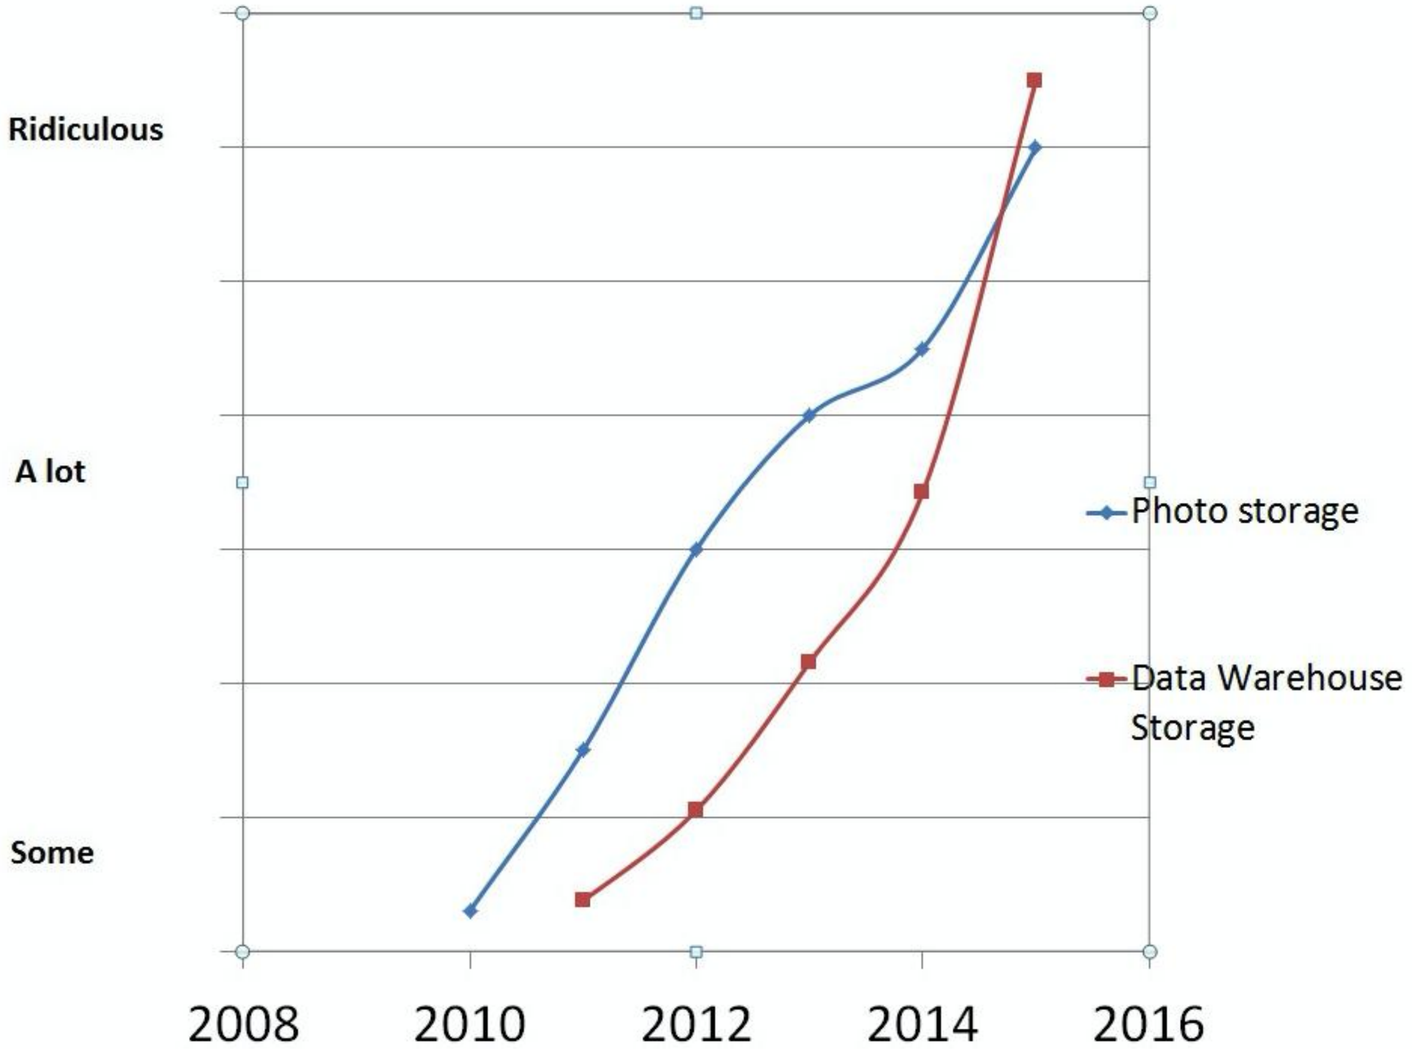
Ridiculous
A lot
Some
2008
2010
2012
2014
2016
Photo storage
Data Warehouse Storage

So what is an Exabyte?

- 1 Exabyte == 1000 Petabytes
- 1 Petabyte == 1000 Terabytes
- 1 Exabyte = ~250,000 4 TB drives

- **250k drives stacked fl**
  >30 times taller than Seattle Space Needle

X 30 times!

Data Scientists are DW customers

Data Scientist and DW

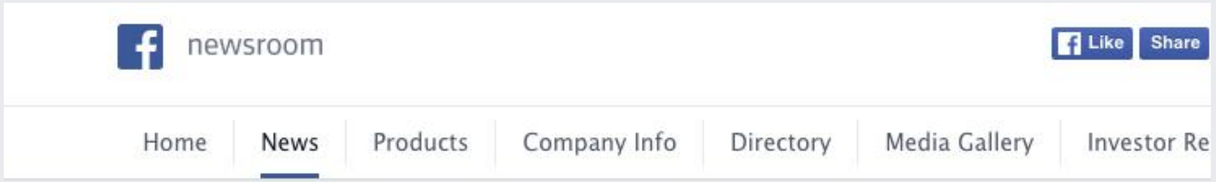

April 4, 2016

Using Artificial Intelligence to Help Blind People 'See' Facebook

Data Scientist and DW

Research at Facebook

Home Our Research Academic Programs Publications Blog People

Core Data Science

Home Publications People Blog

Automatic Alterr

newsroom

Like Share

Home News Products Company Info Directory Media Gallery Investor Re

April 4, 2016

Using Artificial Intelligence to Help Blind People 'See' Facebook

Research at Facebook

Home Our Research Academic Programs Publications Blog People

NFL Fan Friendships on Facebook

Blog

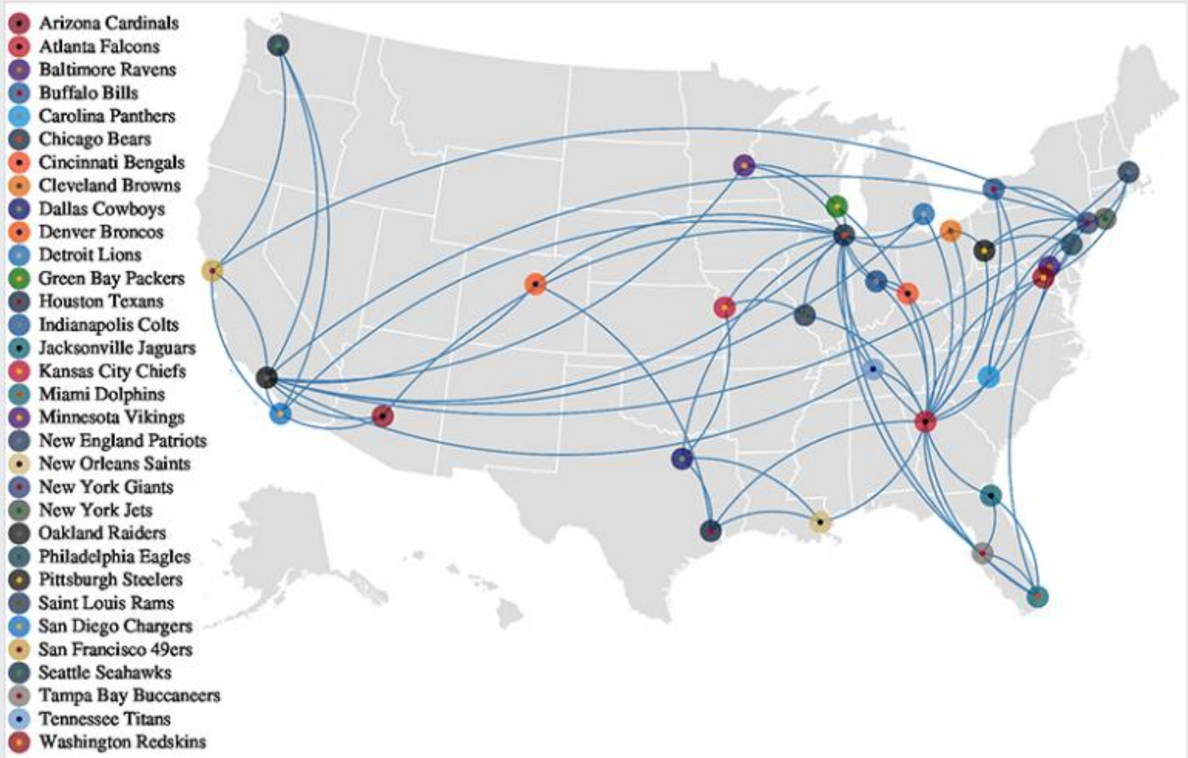

Data Scientist and DW

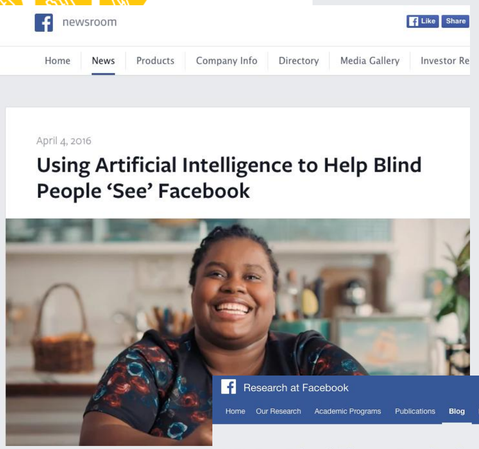

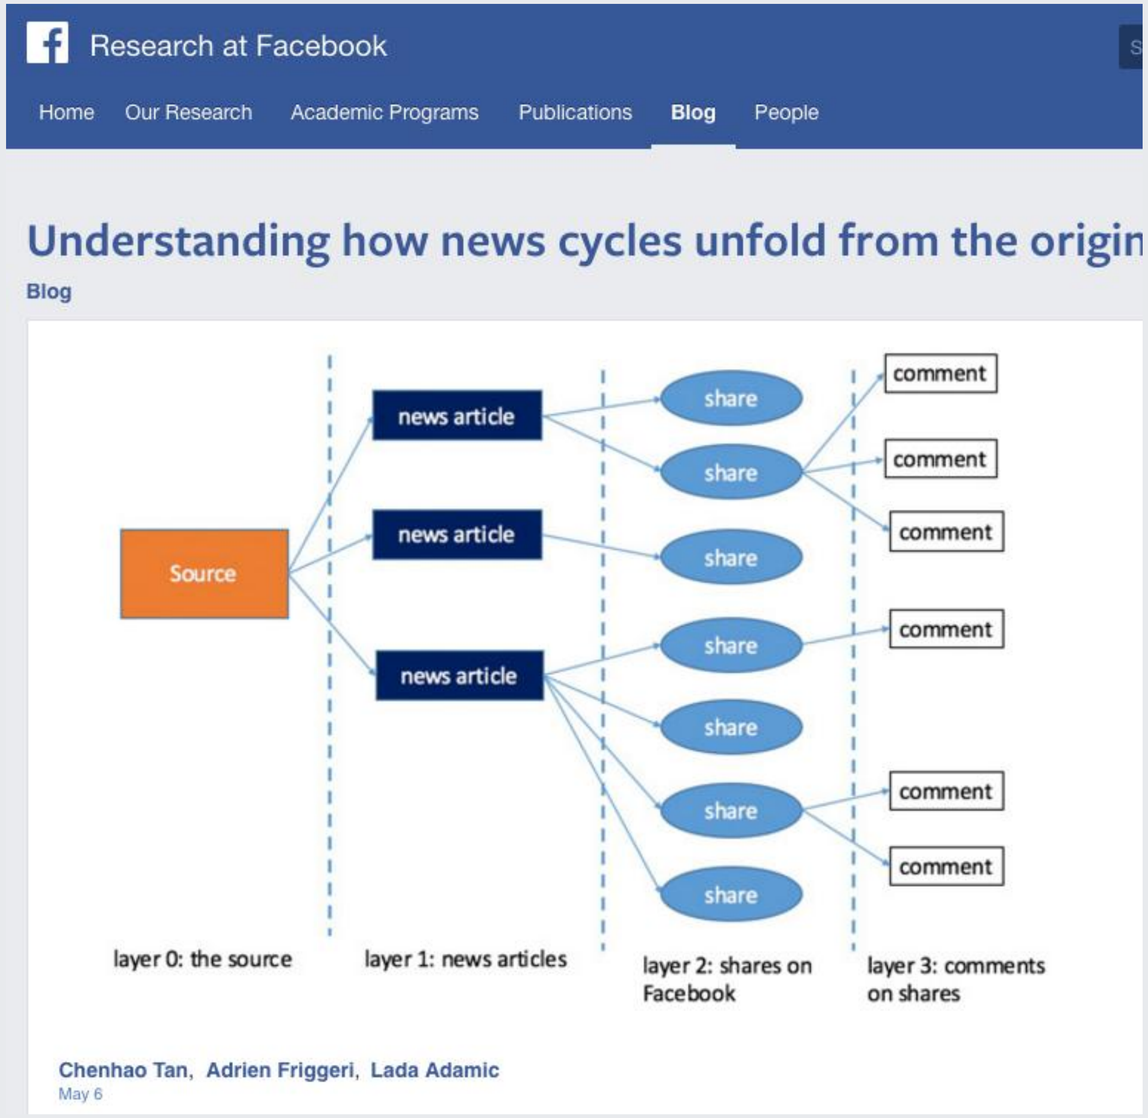

Data Science@FB: dynamic, growing and hard to predict

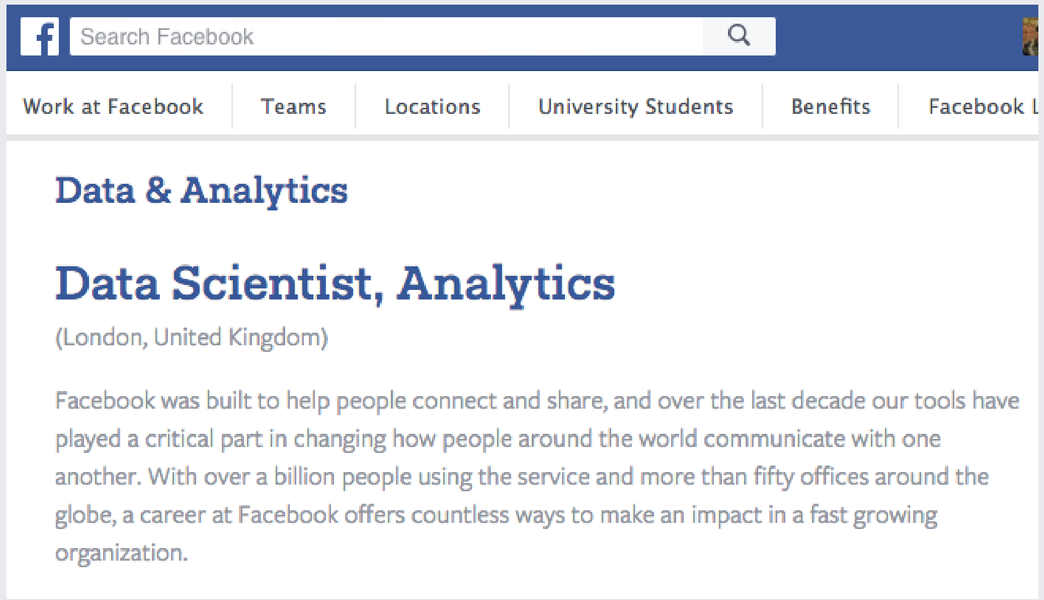

Data Warehouse at Facebook

Data Warehouse at Facebook

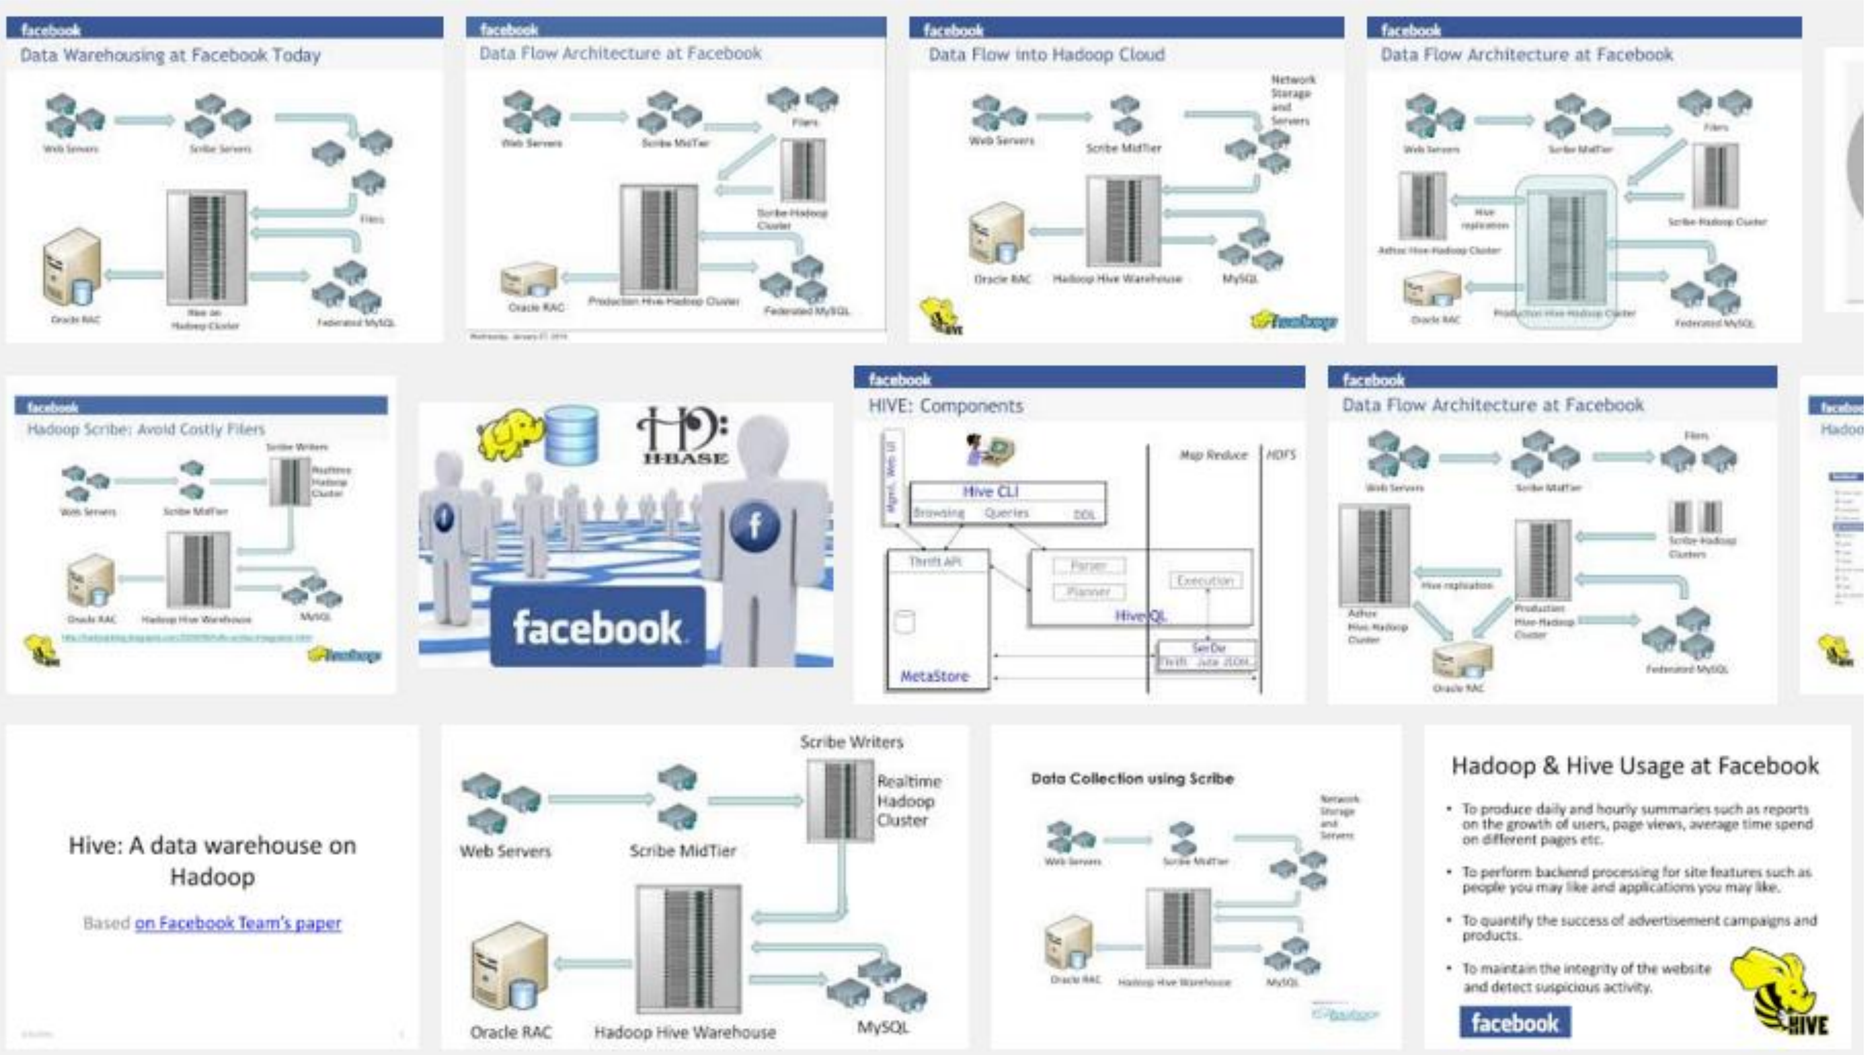

DW @FB: Hive on Hadoop (HDFS)

DW @FB: Hive on Hadoop (HDFS)

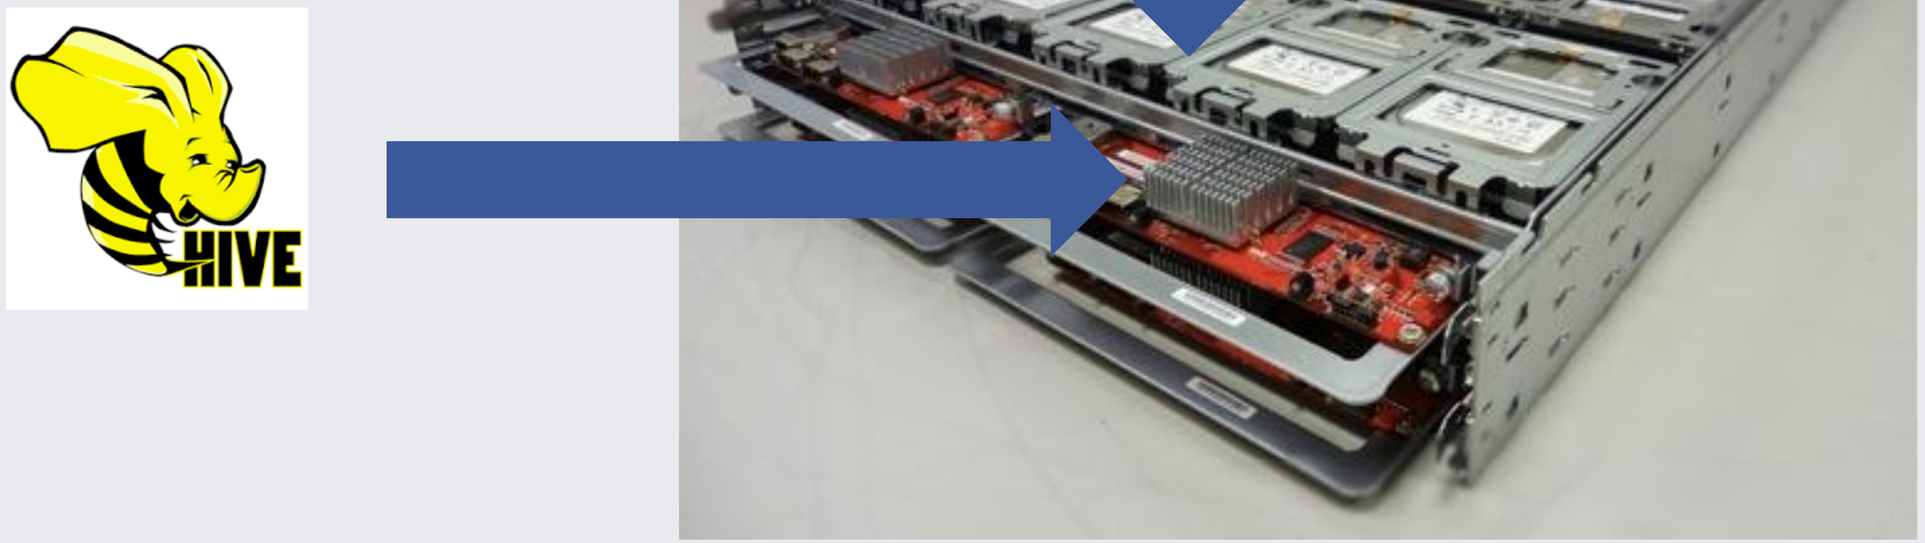

DW @FB: Hive on Hadoop (HDFS)

Permanent and Temp Storage

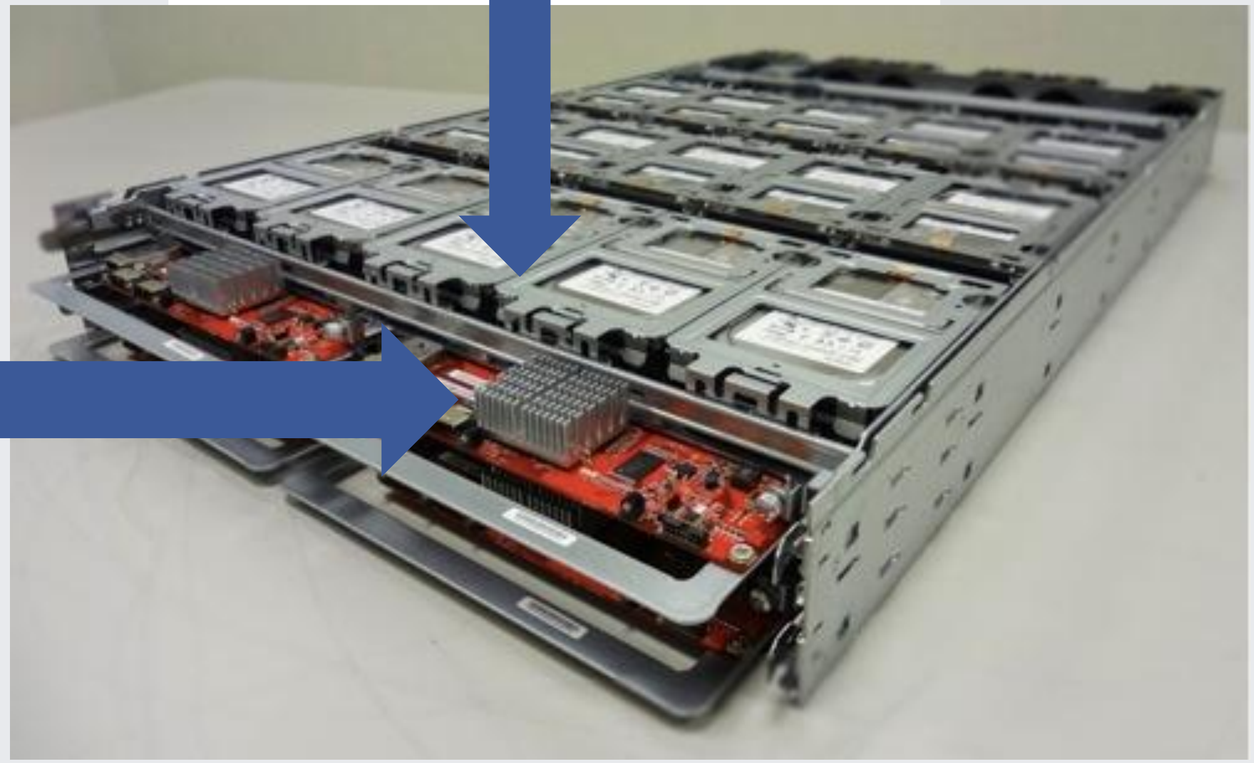

Dynamic demand vs. static resource allocation

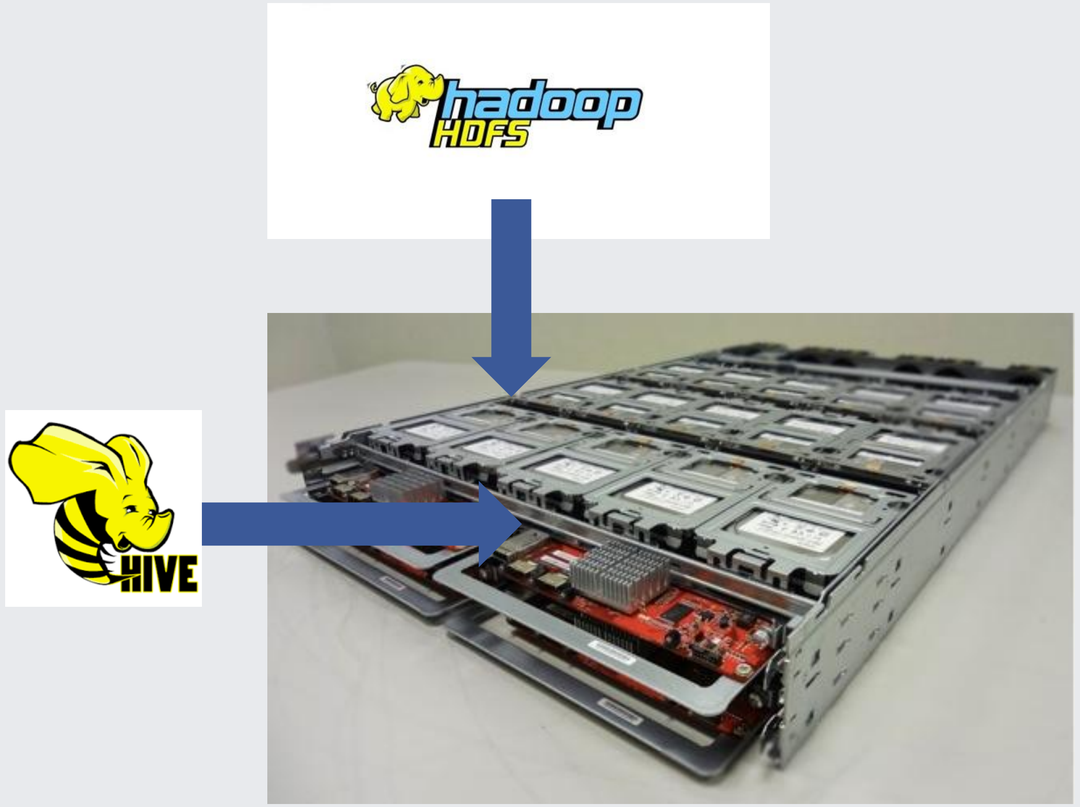

	Specification	Quantity
Storage	4 TB SAS HDD	15
CPU	Intel Xeon 20 core	2
Network	10 Gbps	1

Efficiency at hyper scale

	Specification	Quantity
Storage	4 TB SAS HDD	15
CPU	Intel Xeon 20 core	2
Network	10 Gbps	1

Elasticity
Splitting DC for Storage and Compute

	Specification	Quantity
Storage	4 TB SAS HDD	15
CPU	Intel Xeon 20 core	2
Network	10 Gbps	1

Compute

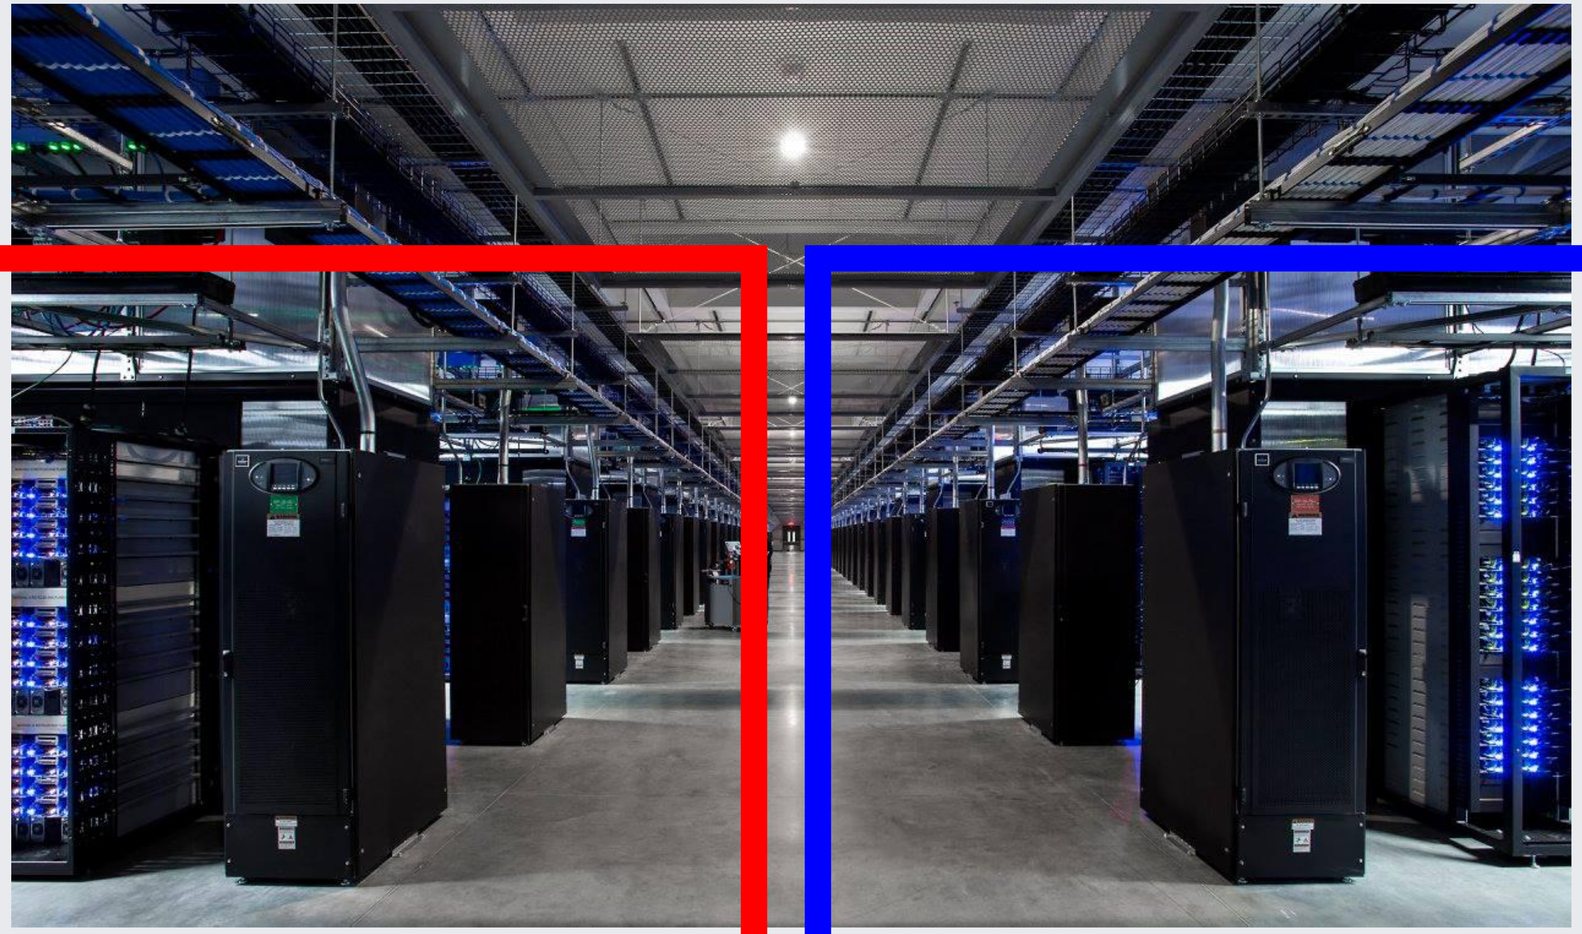

Storage

Elasticity
Splitting DC for Storage and Compute

	Specification	Quantity
Storage	4 TB SAS HDD	15
CPU	Intel Xeon 20 core	2
Network	10 Gbps	1

Compute

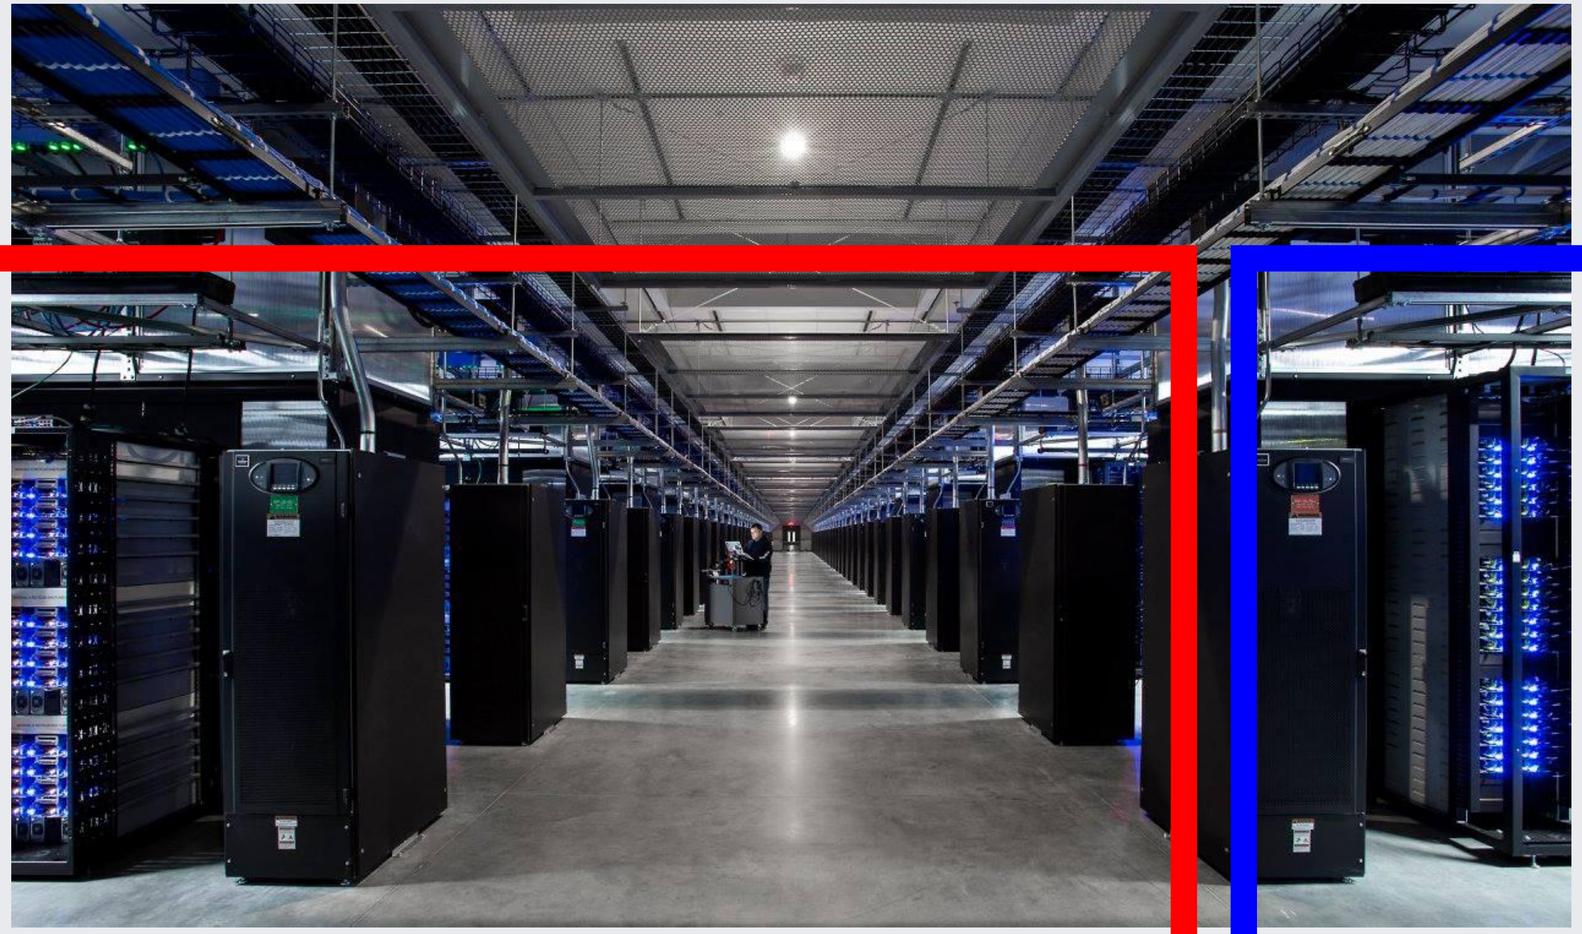

Storage

Elasticity
Splitting DC for Storage and Compute

	Specification	Quantity
Storage	4 TB SAS HDD	15
CPU	Intel Xeon 20 core	2
Network	10 Gbps	1

Compute

Storage

Elasticity

Compute

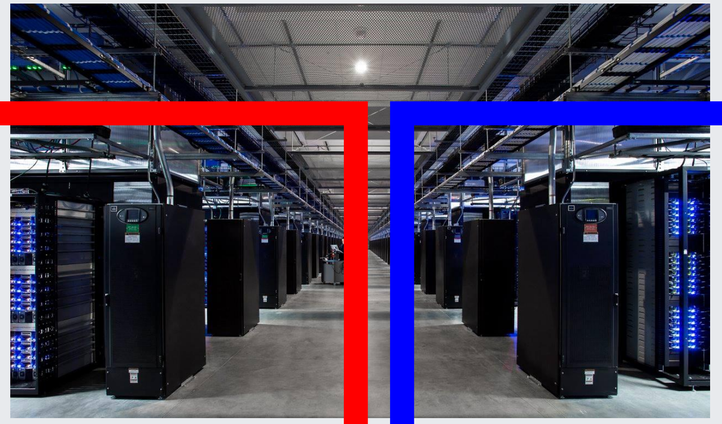

- Add more compute if/when needed
- Upgrade to latest Intel’ s CPUs
- Replace compute every 3 years or sooner

Storage

- Keep HDDs longer than 3 years
- Grow storage capacity independent from compute
- Migrate to 8-10 TB HDDs

facebook

Storage and Compute separation

Is that possible?

Compute and storage: local vs. remote

Compute

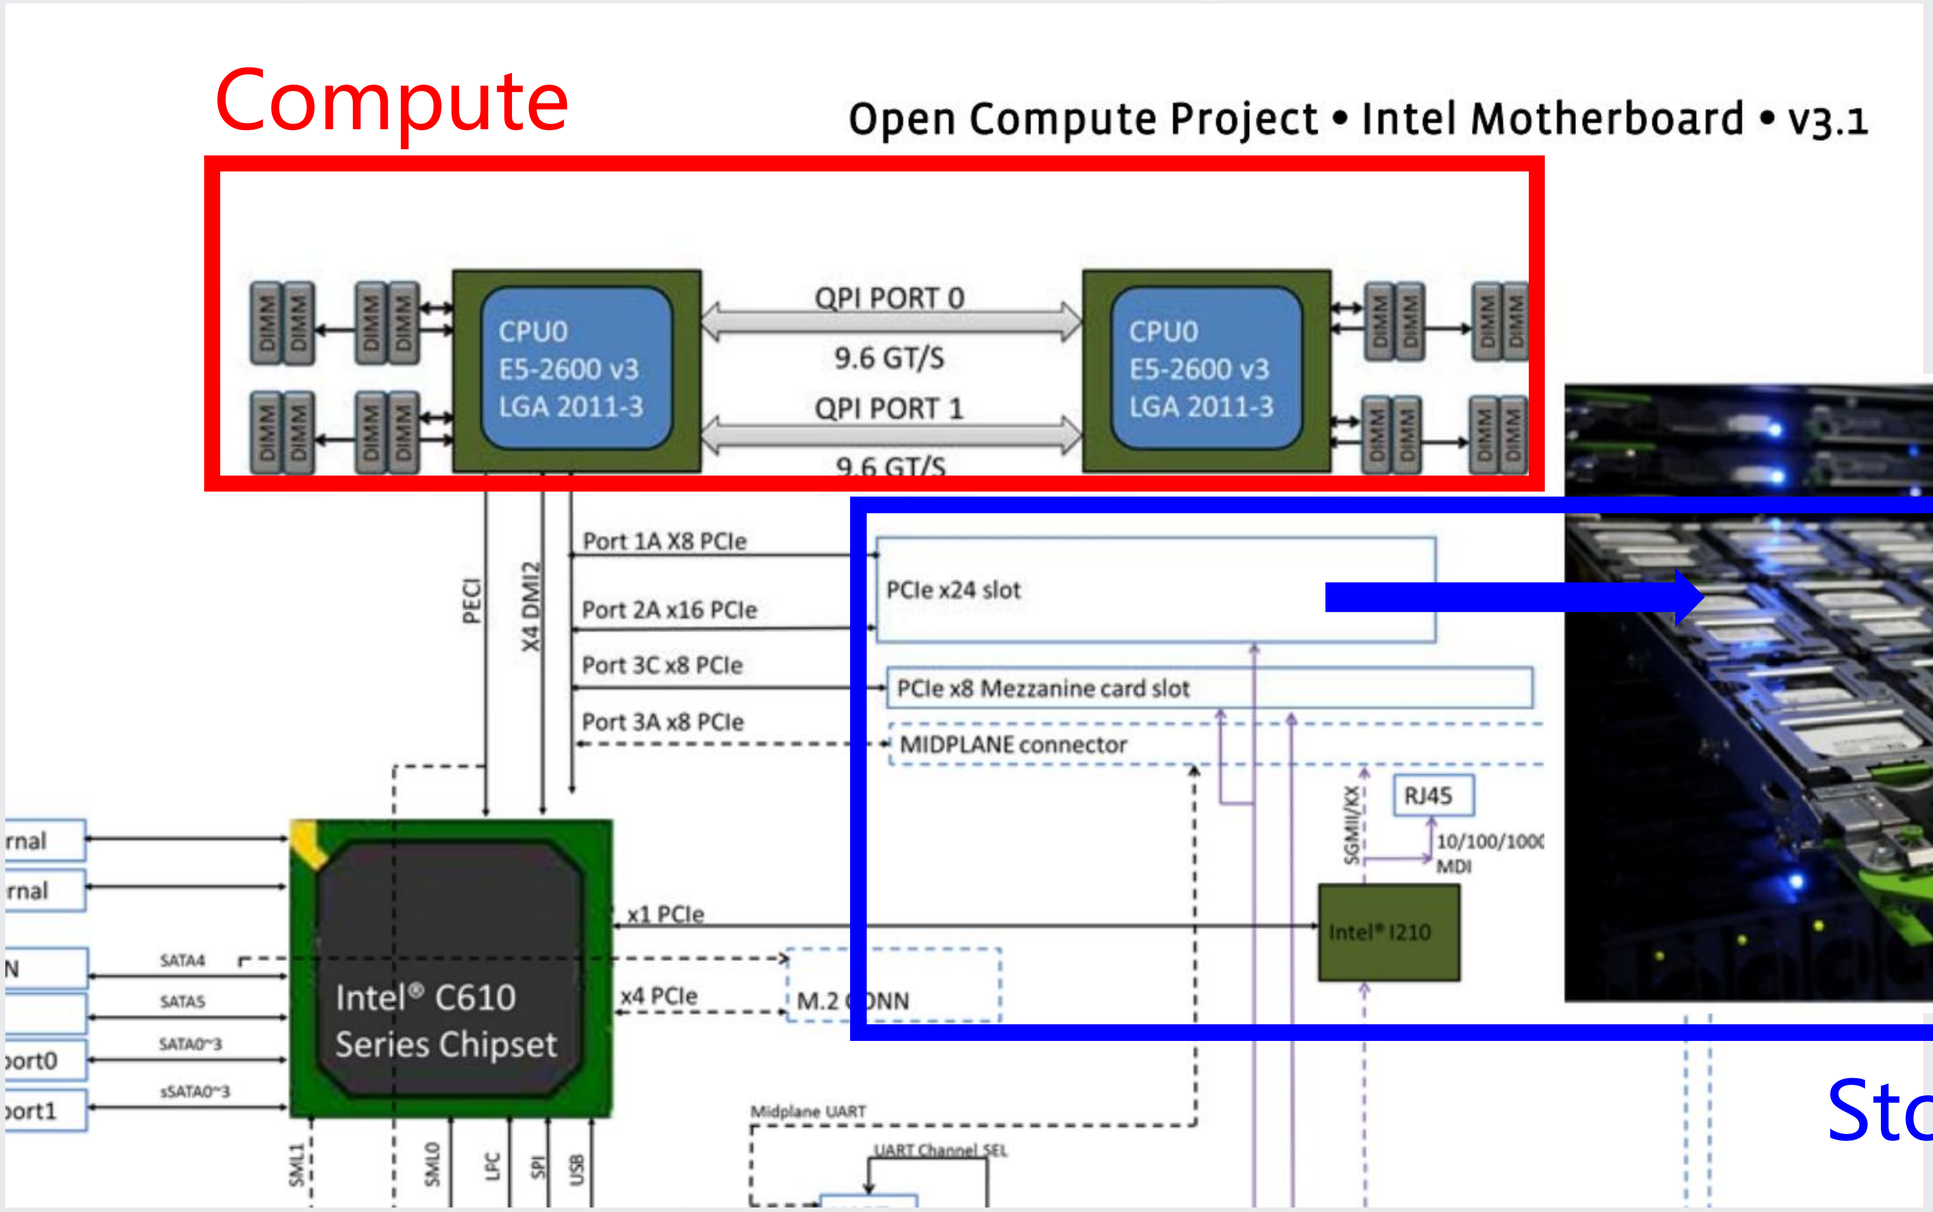

Storage

Last ~25 years of HDD evolution

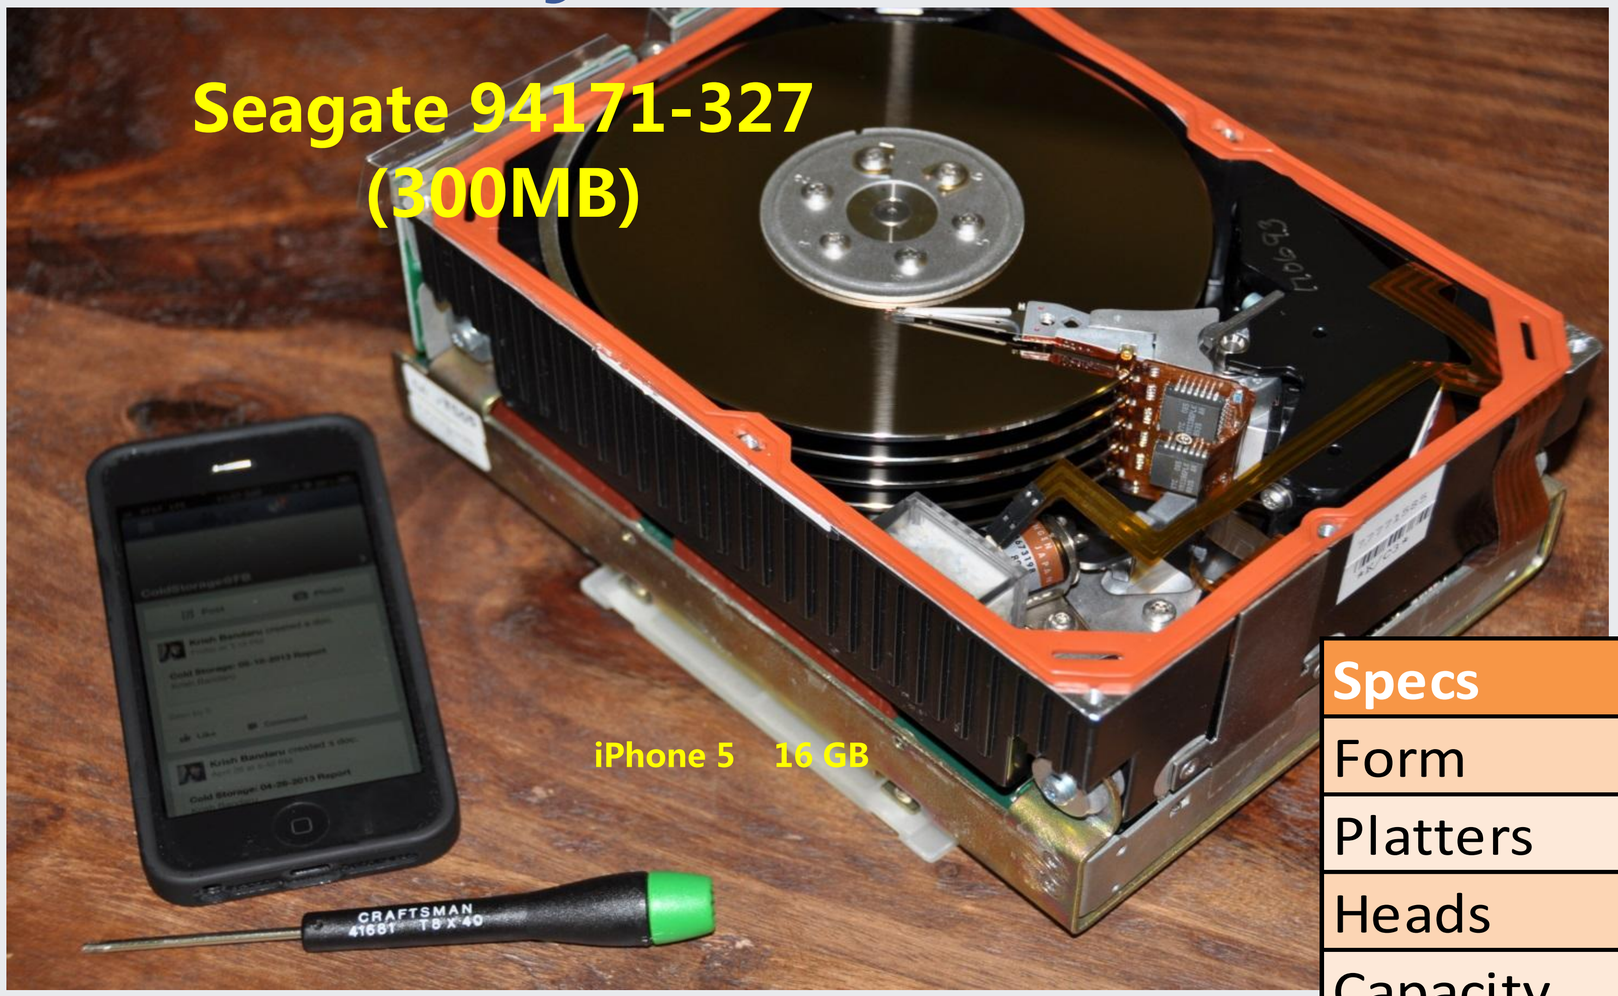

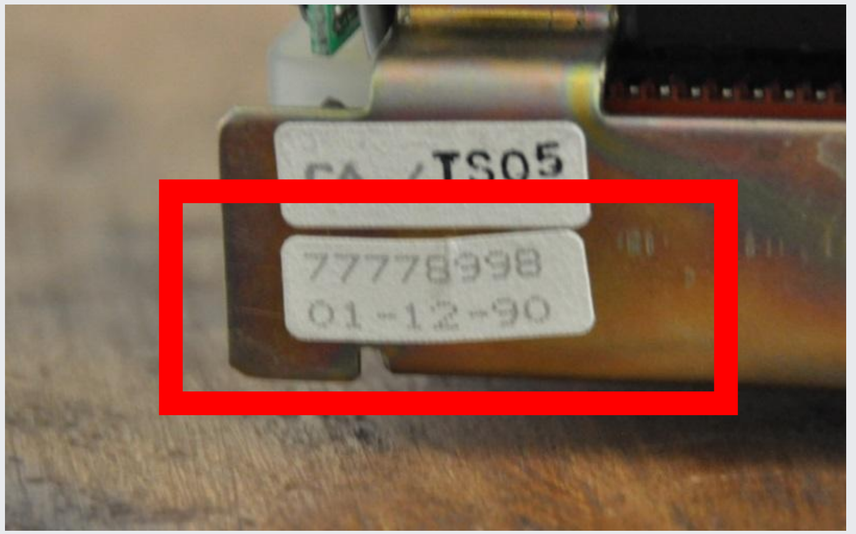

Specs	Value
Form	3.5"
Platters	5
Heads	9
Capacity	300MB
Interface	SCSI
Seek time	17ms
Data transfer rate	1 MB/sec

Last ~25 years of HDD evolution

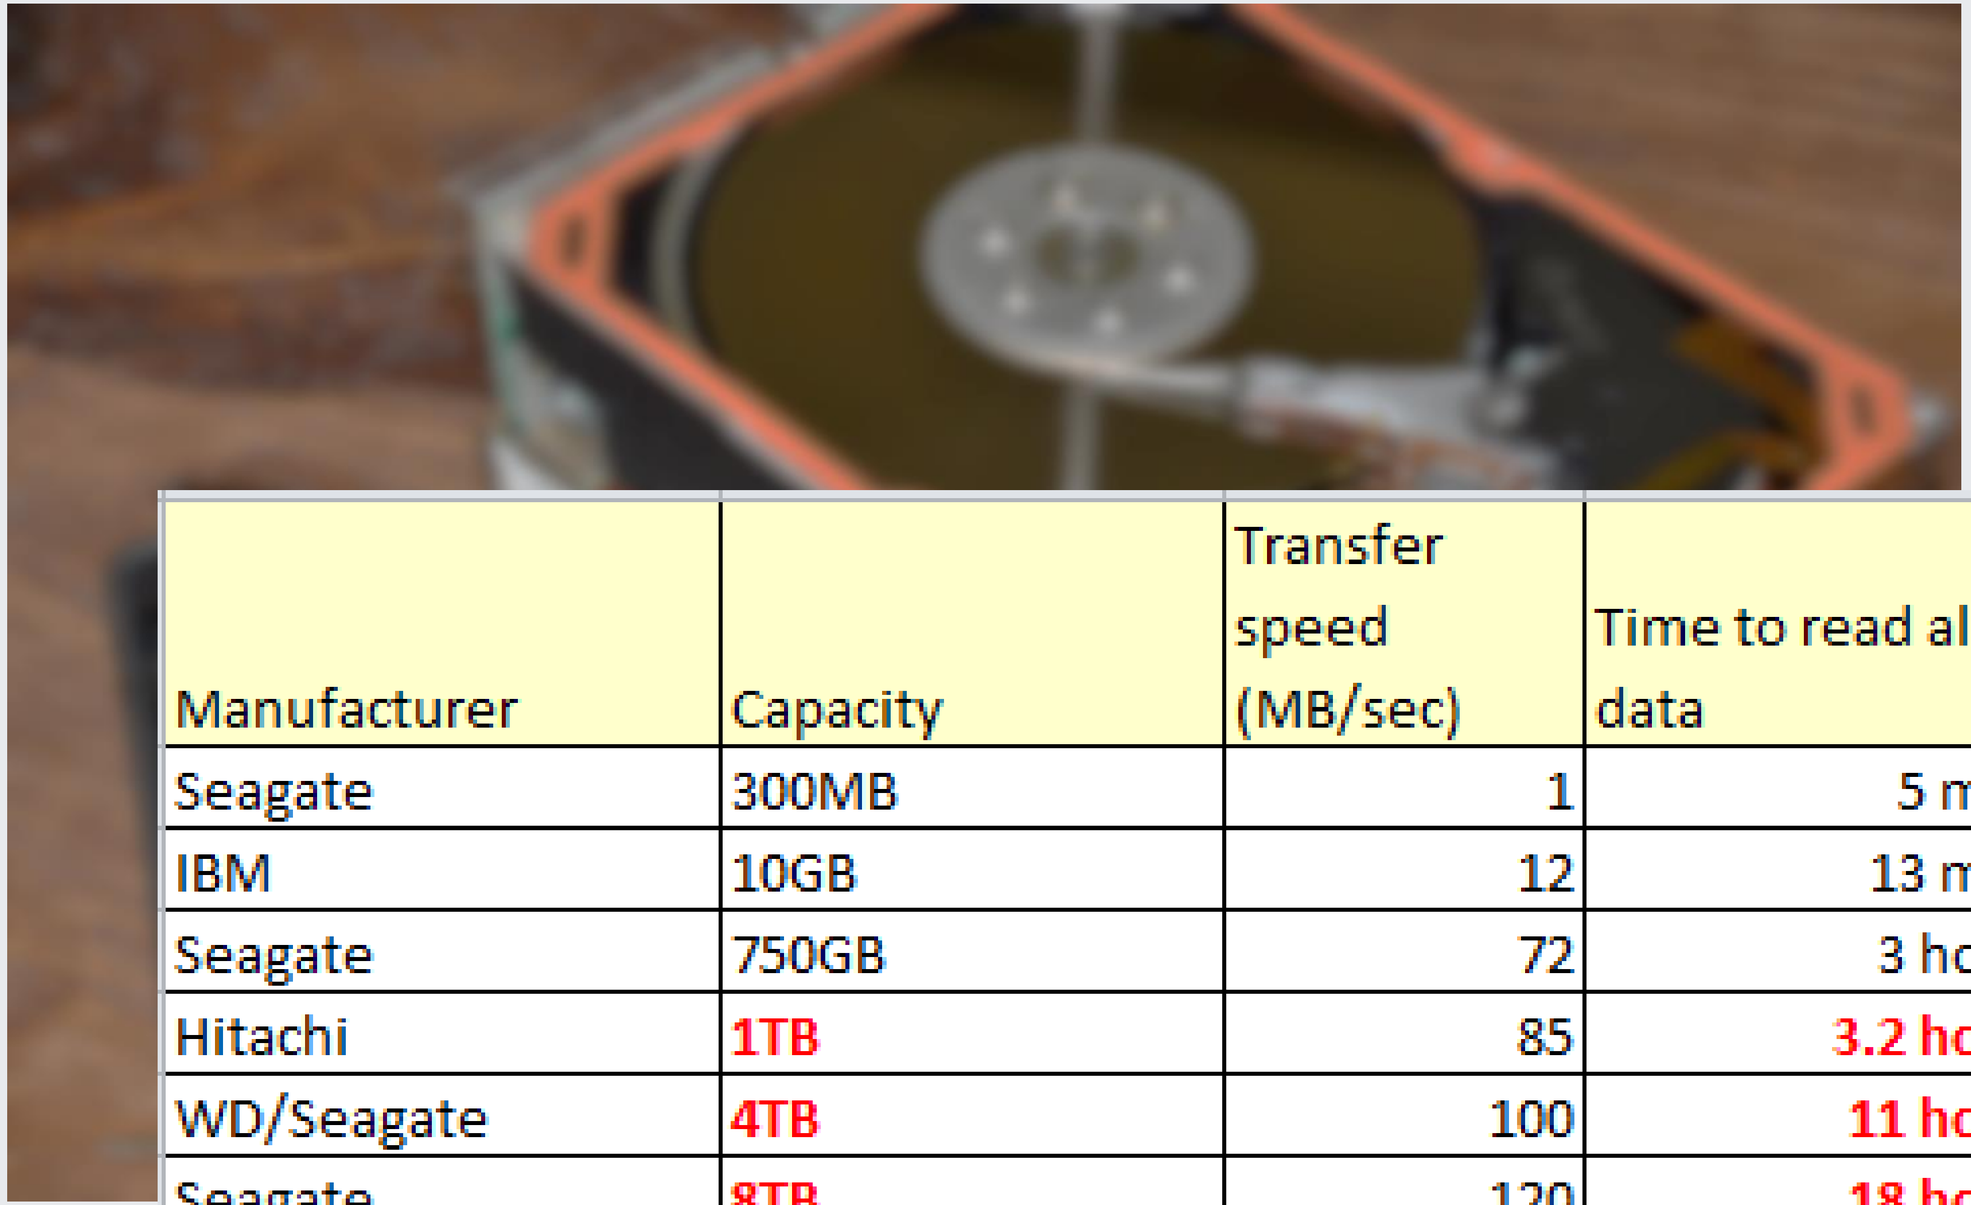

Manufacturer	Capacity	Transfer speed (MB/sec)	Time to read all data	Year
Seagate	300MB	1	5 mins	1990
IBM	10GB	12	13 mins	1998
Seagate	750GB	72	3 hours	2006
Hitachi	**1TB**	85	**3.2 hours**	2007
WD/Seagate	**4TB**	100	**11 hours**	2012
Seagate	**8TB**	120	**18 hours**	2014

Last ~25 years of Ethernet

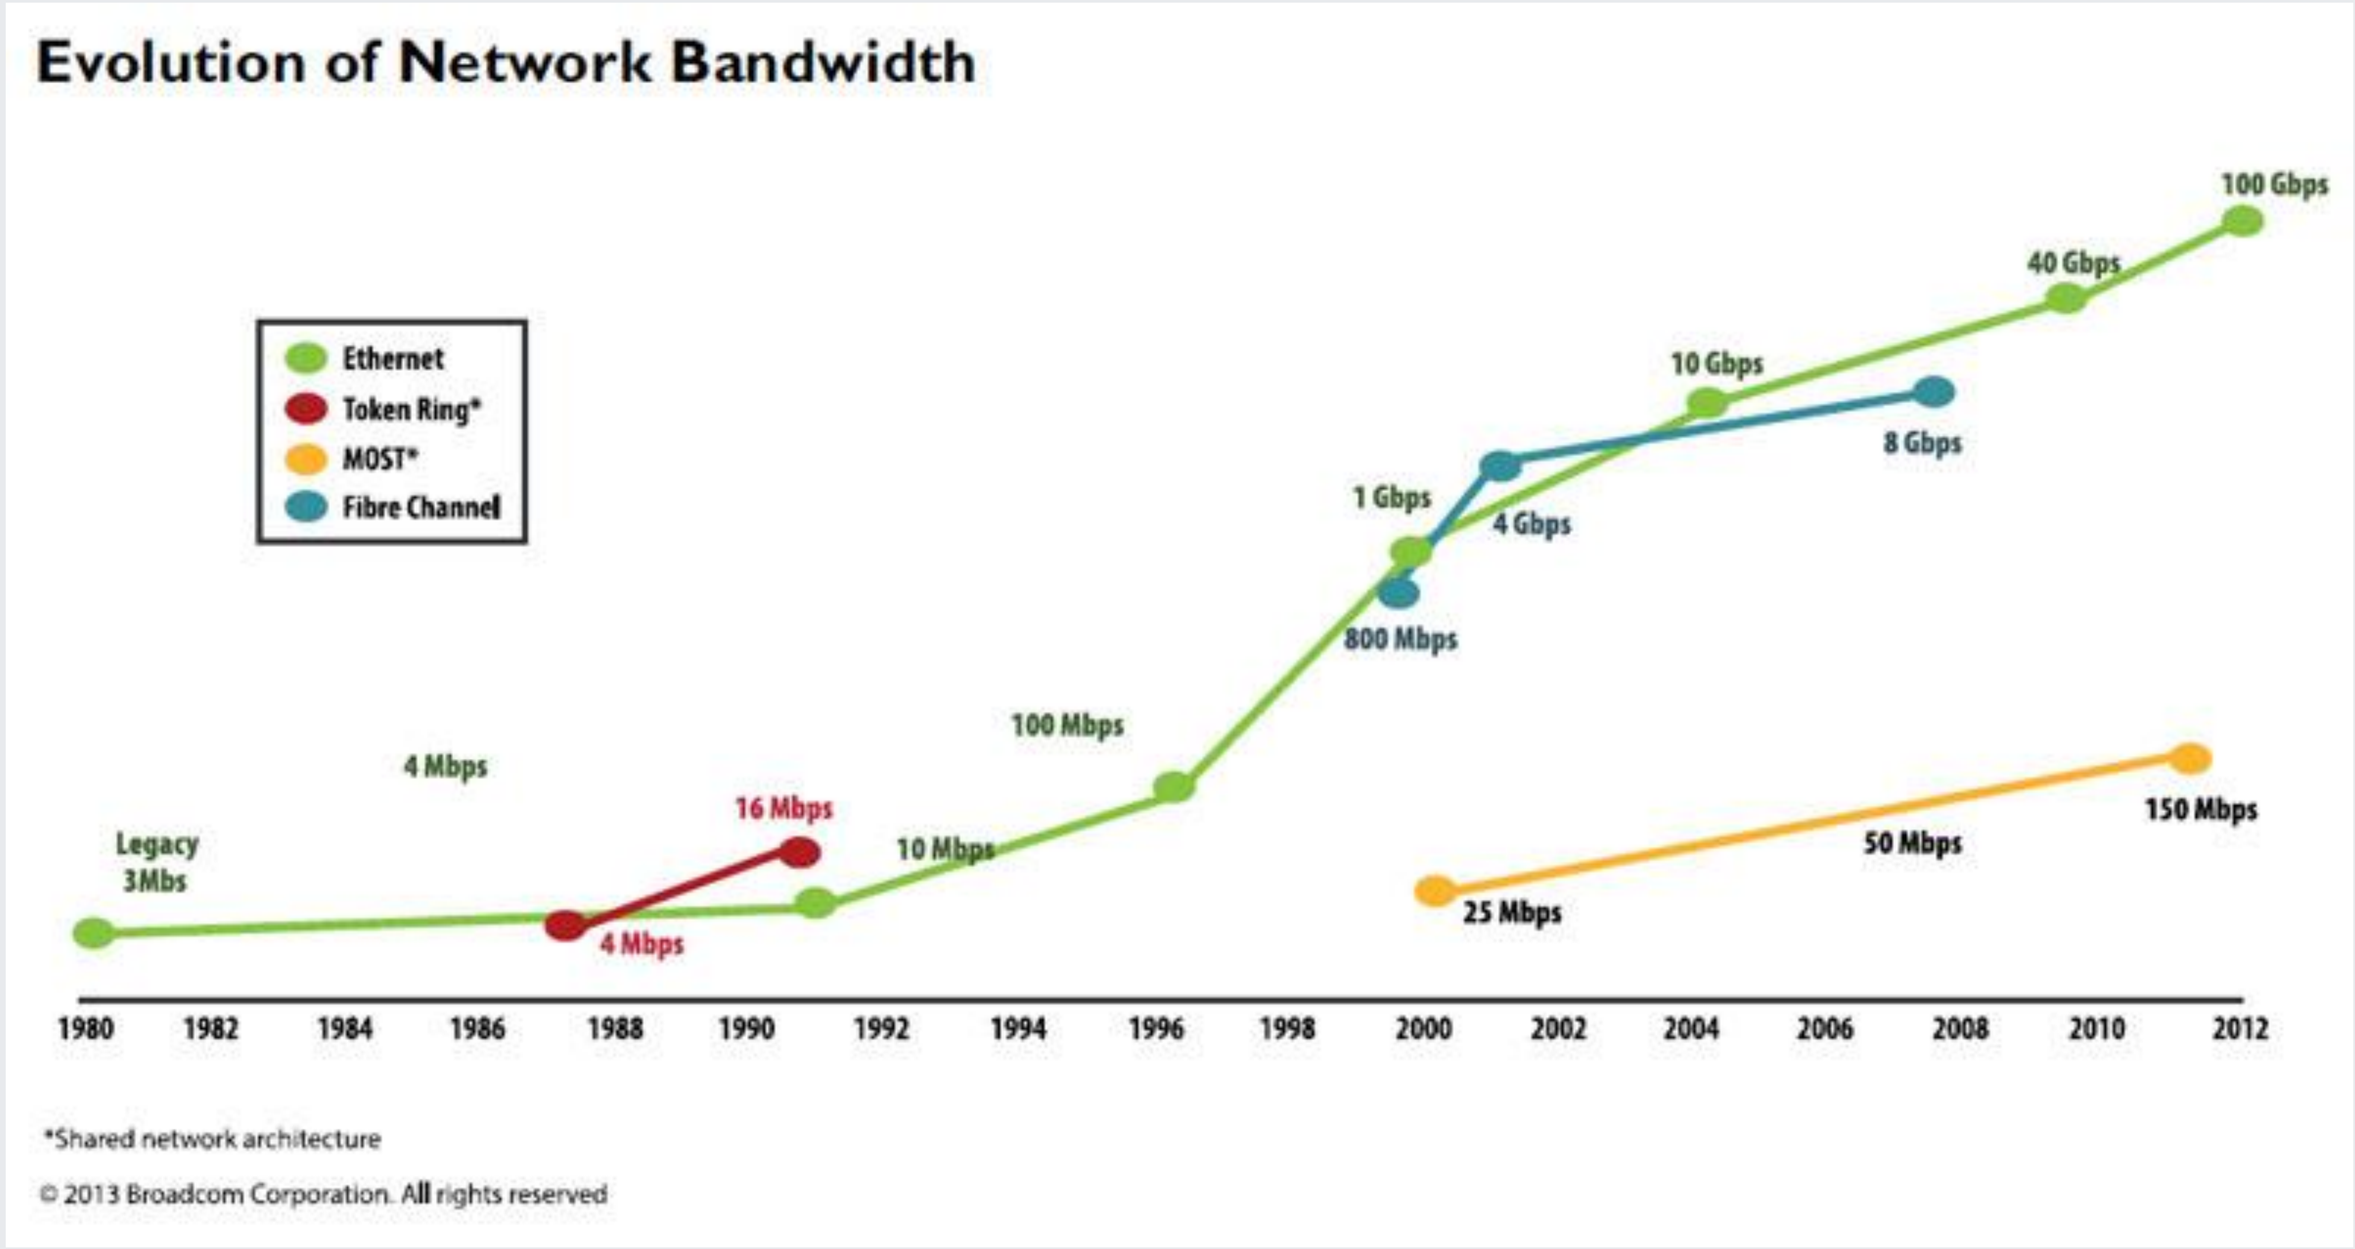

Last ~25 years of HDD vs. Ethernet

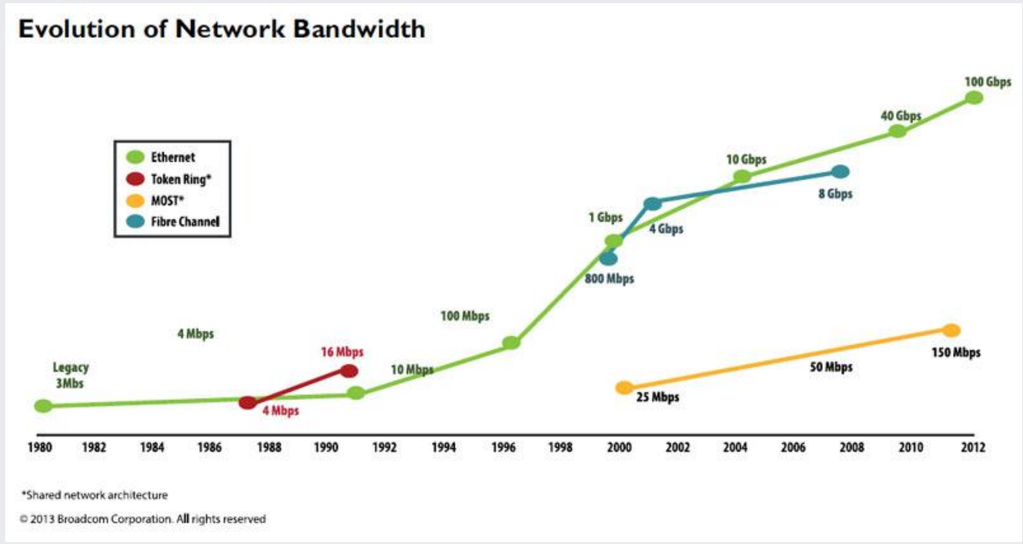

Manufacturer	Capacity	Transfer speed (MB/sec)	Time to read all data	Year
Seagate	300MB	1	5 mins	1990
IBM	10GB	12	13 mins	1998
Seagate	750GB	72	3 hours	2006
Hitachi	1TB	85	3.2 hours	2007
WD/Seagate	4TB	100	11 hours	2012
Seagate	8TB	120	18 hours	2014

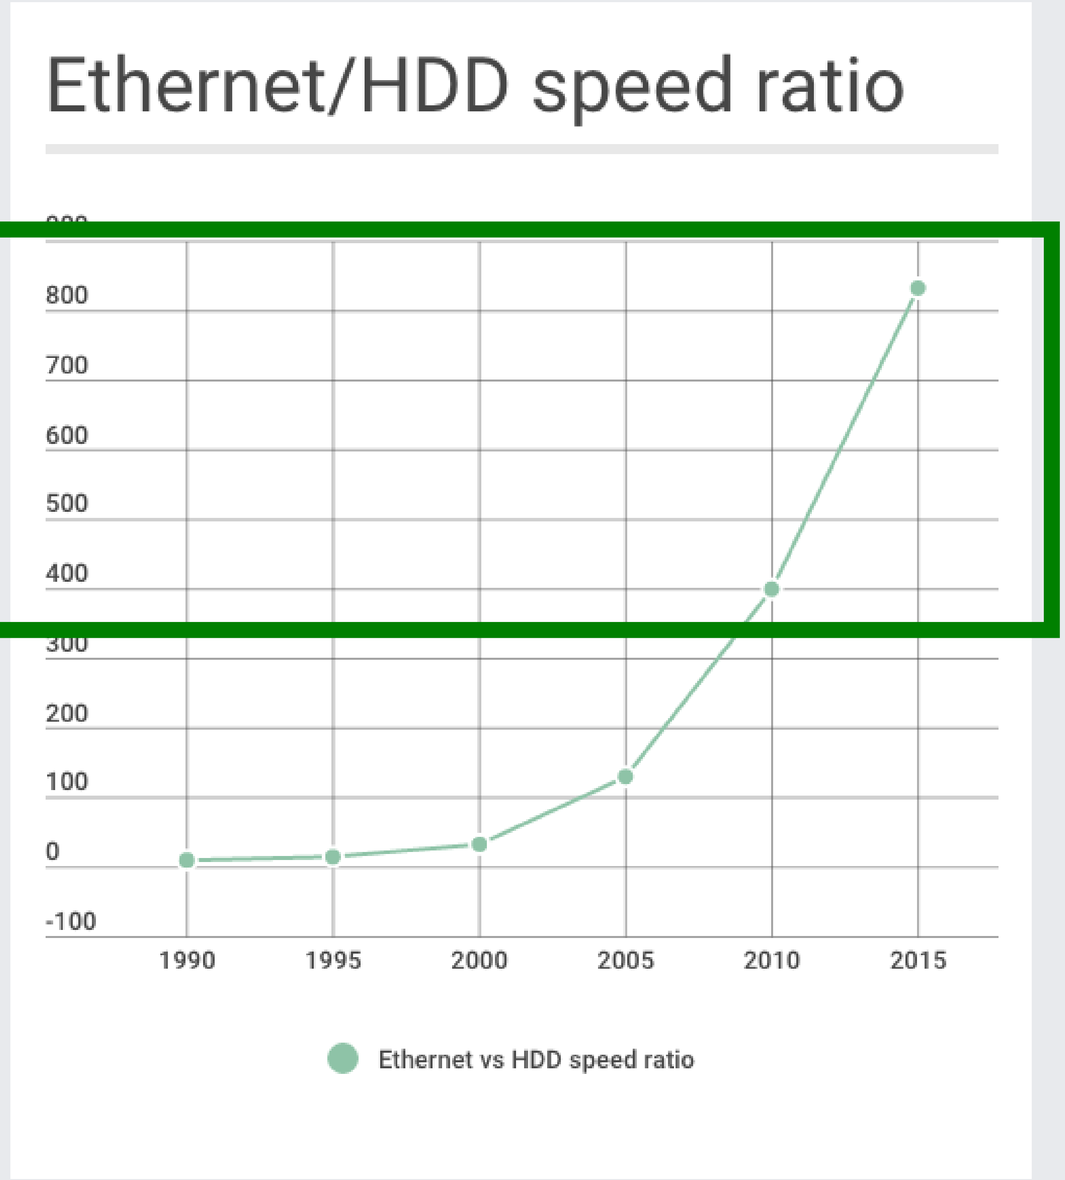

When Ethernet is faster than HDDs
And Data Center network is ‘flat’

Compute

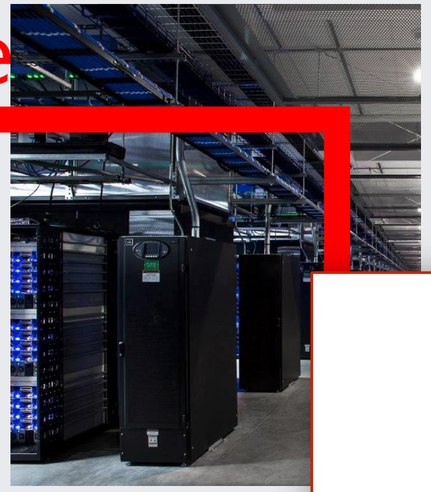

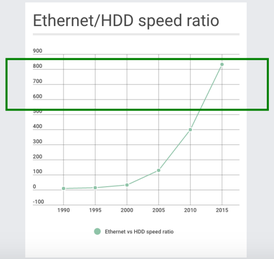

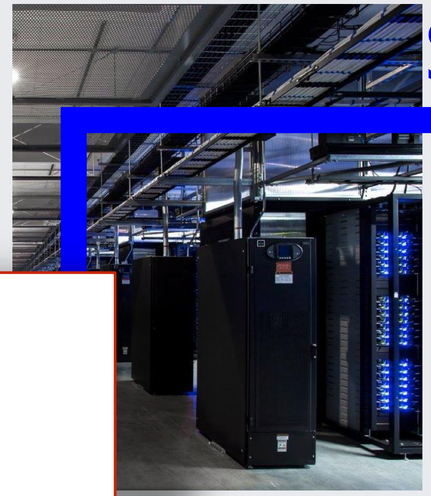

Storage

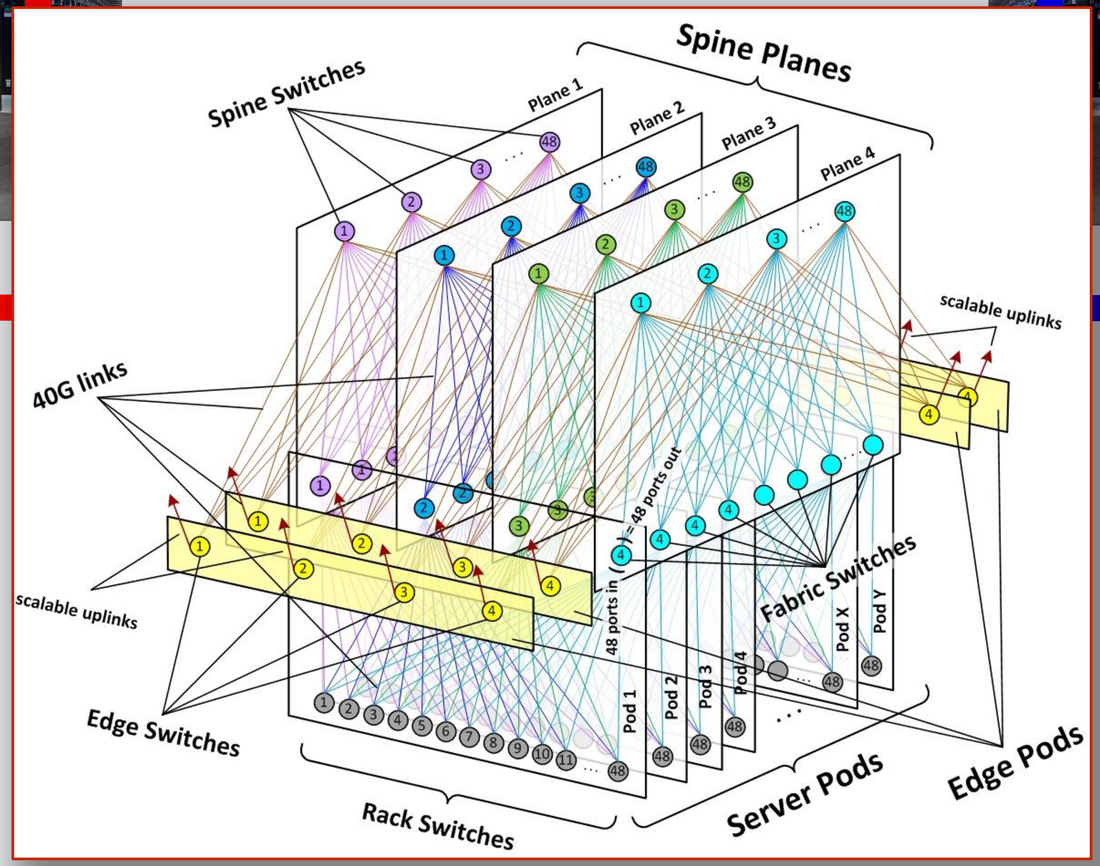

facebook

Warm Storage

Architecture

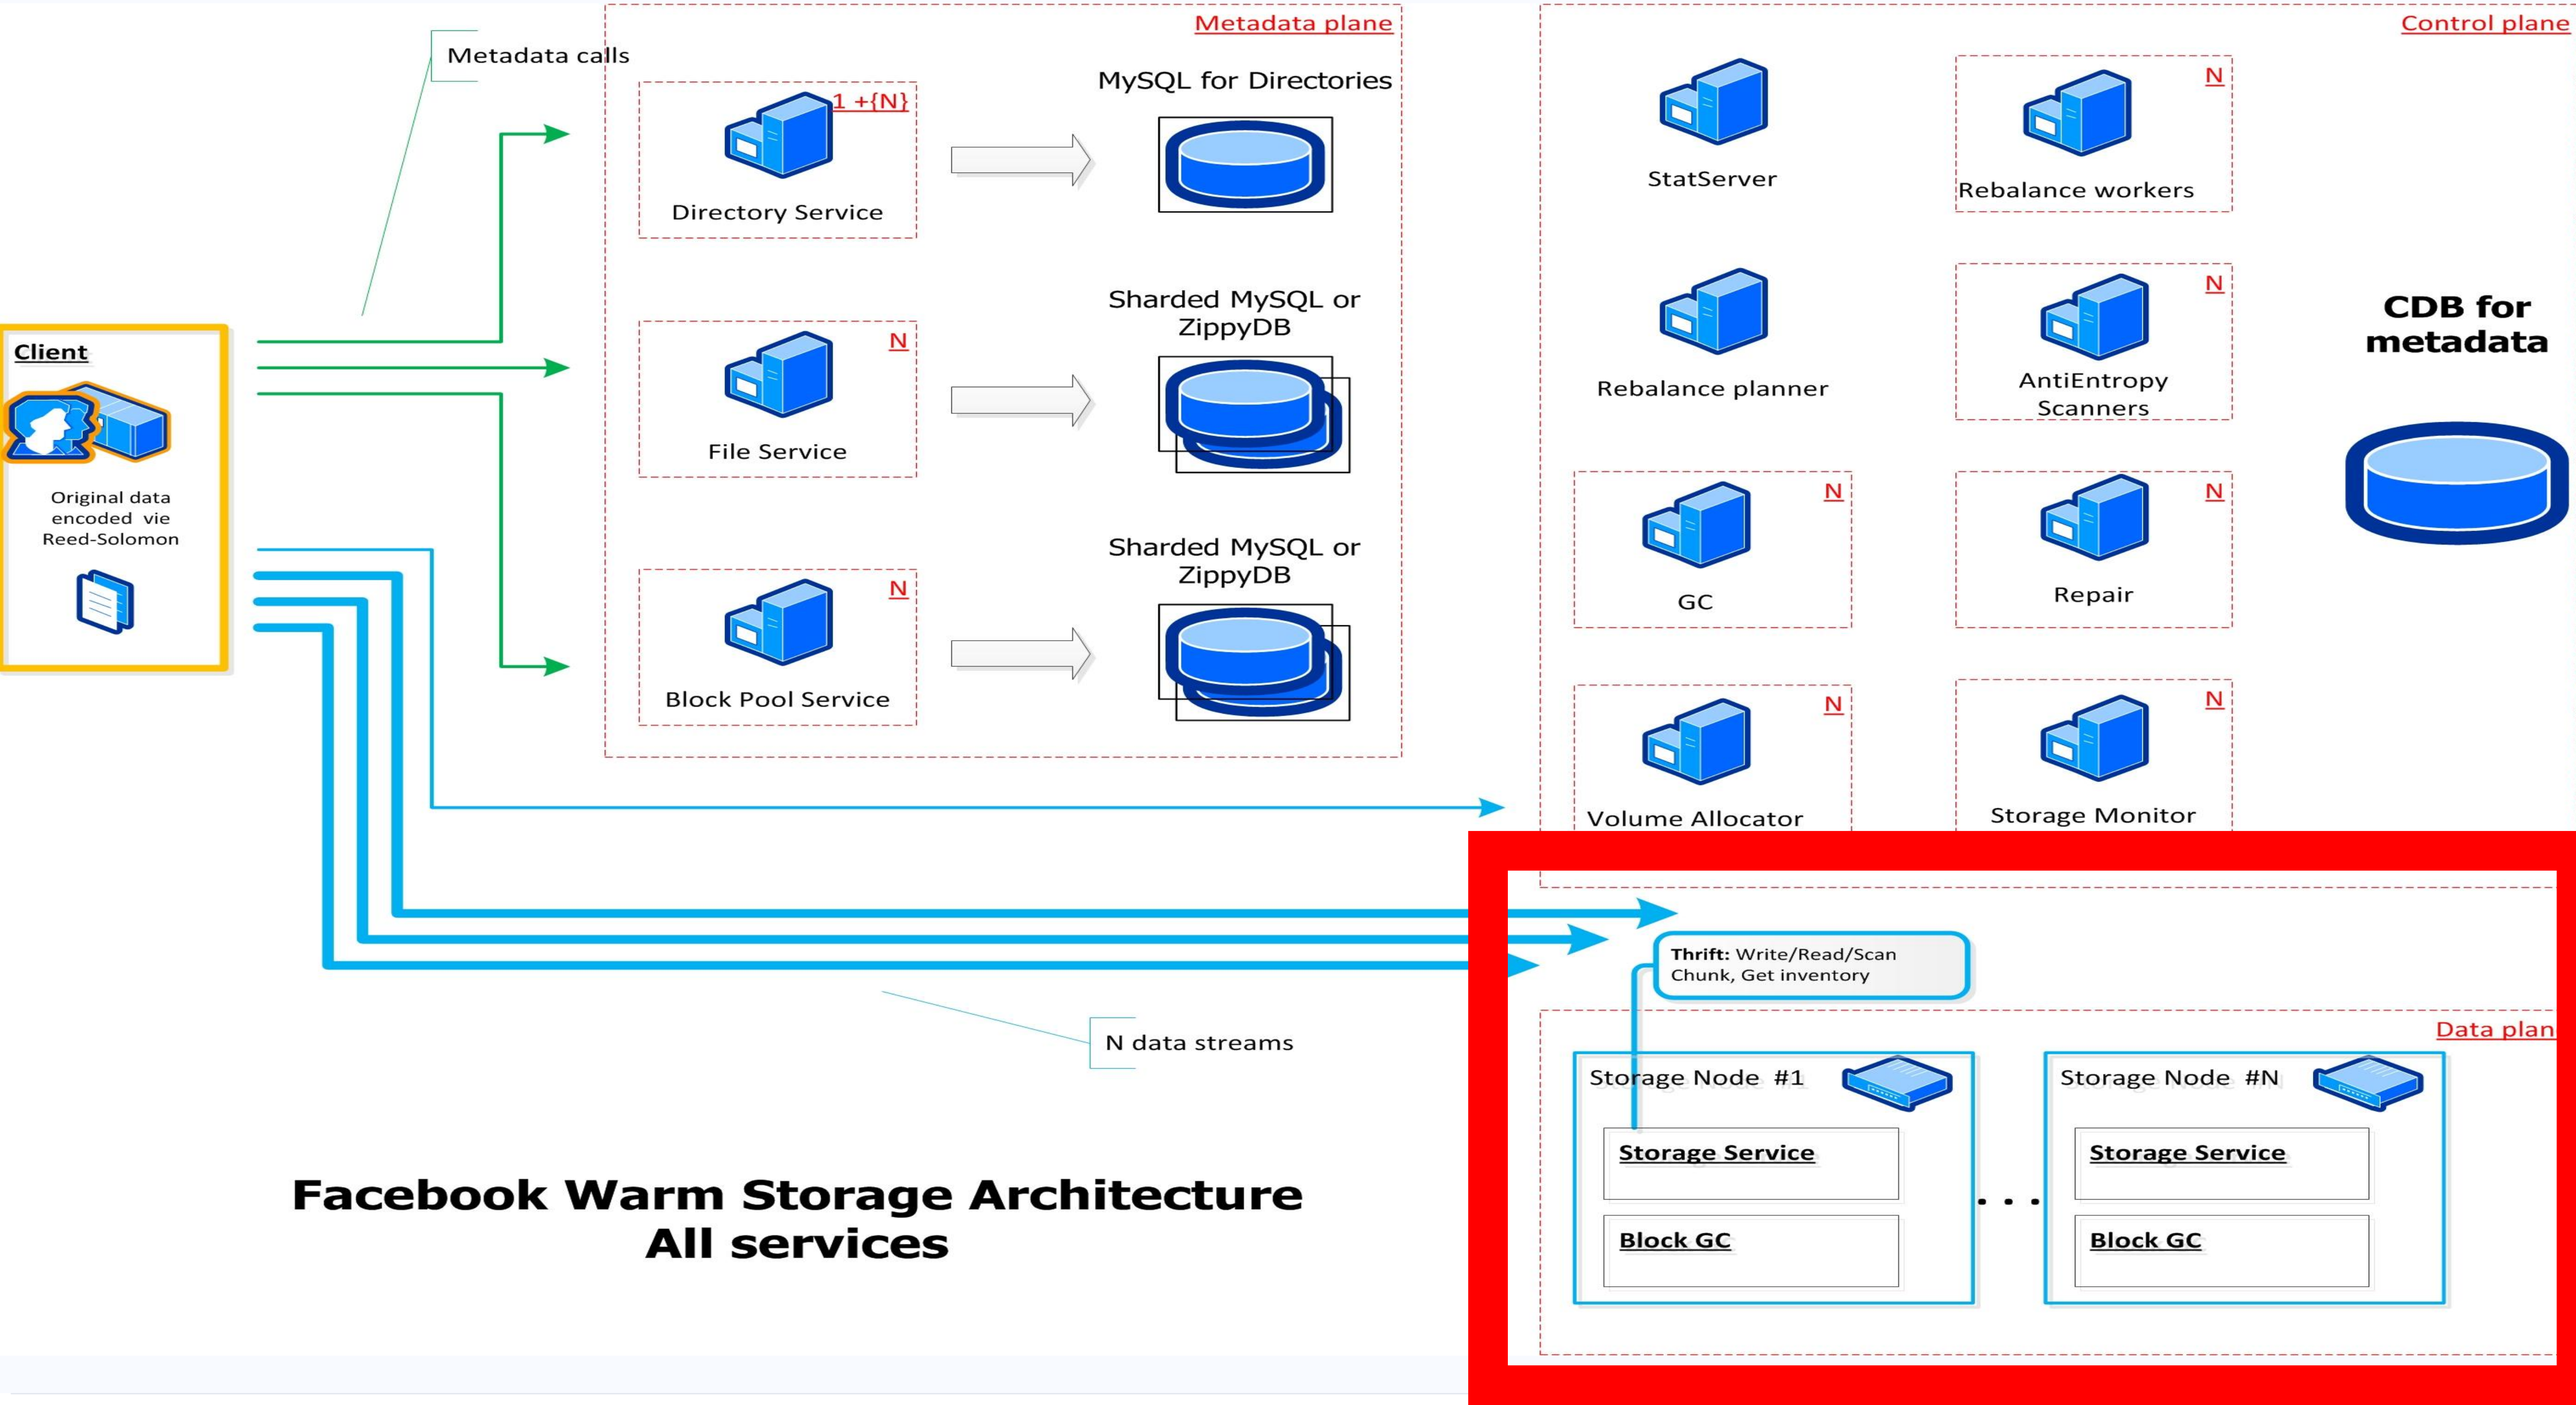

Facebook Warm Storage Architecture
All services

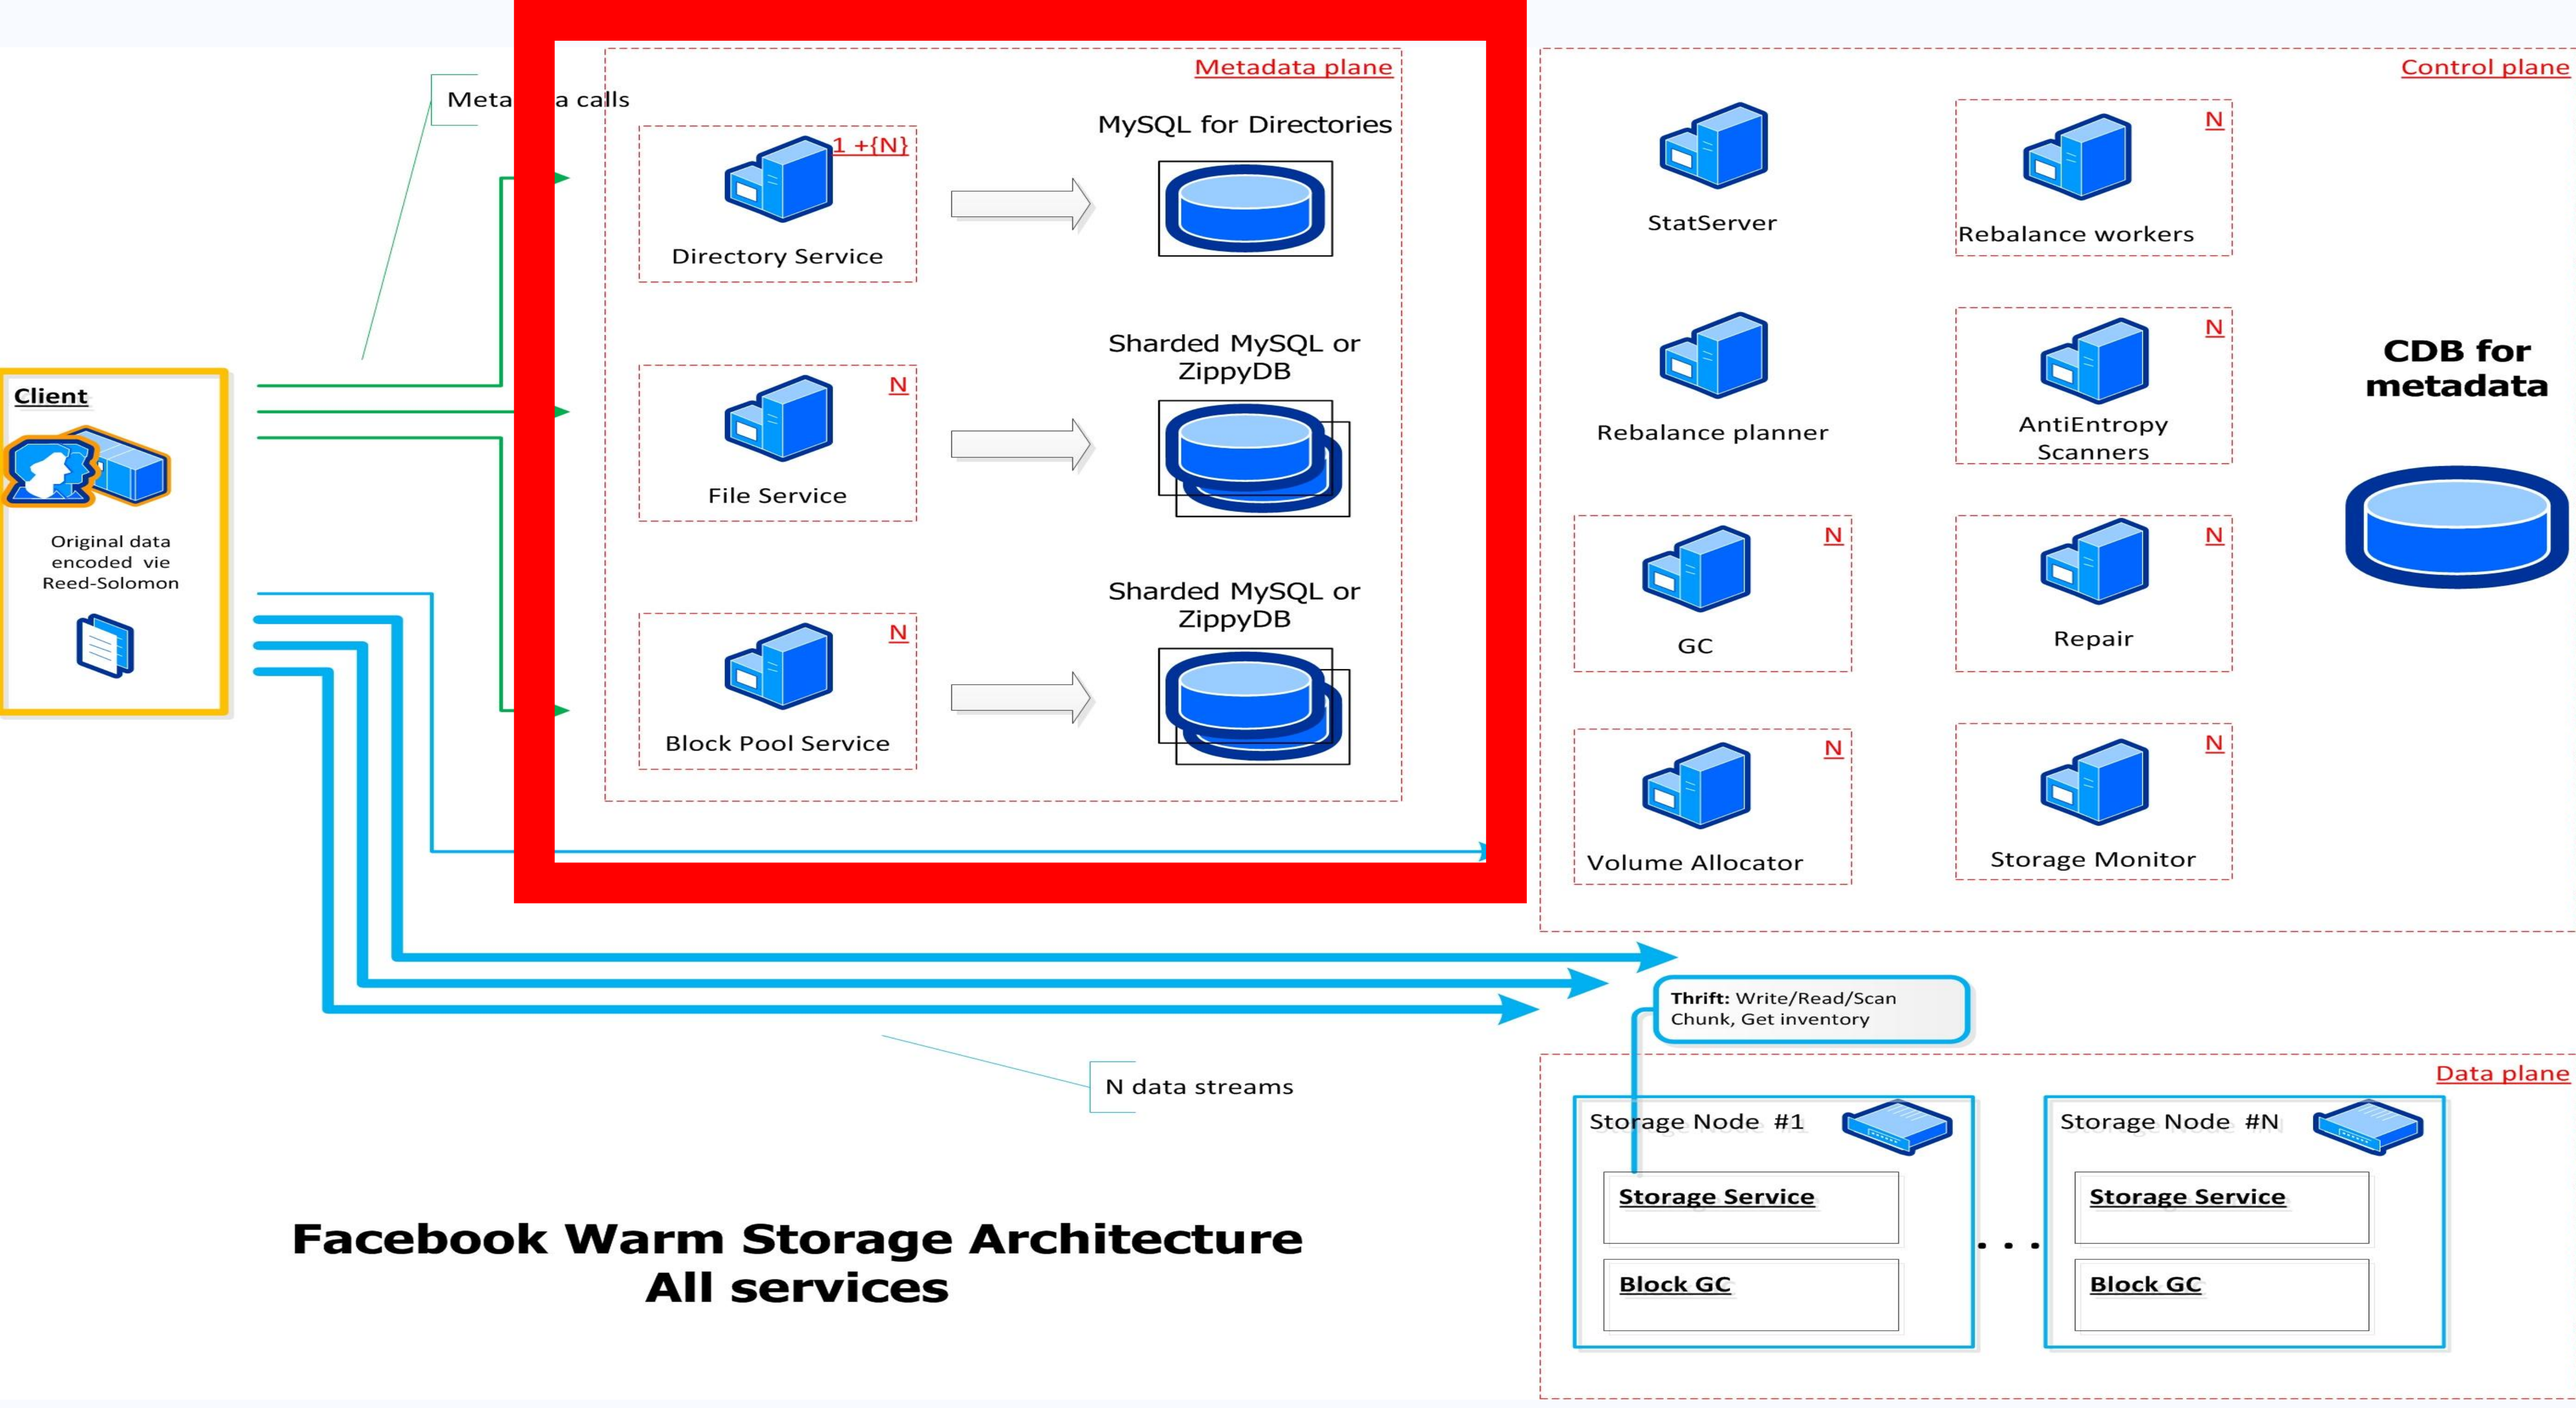

Facebook Warm Storage Architecture
All services

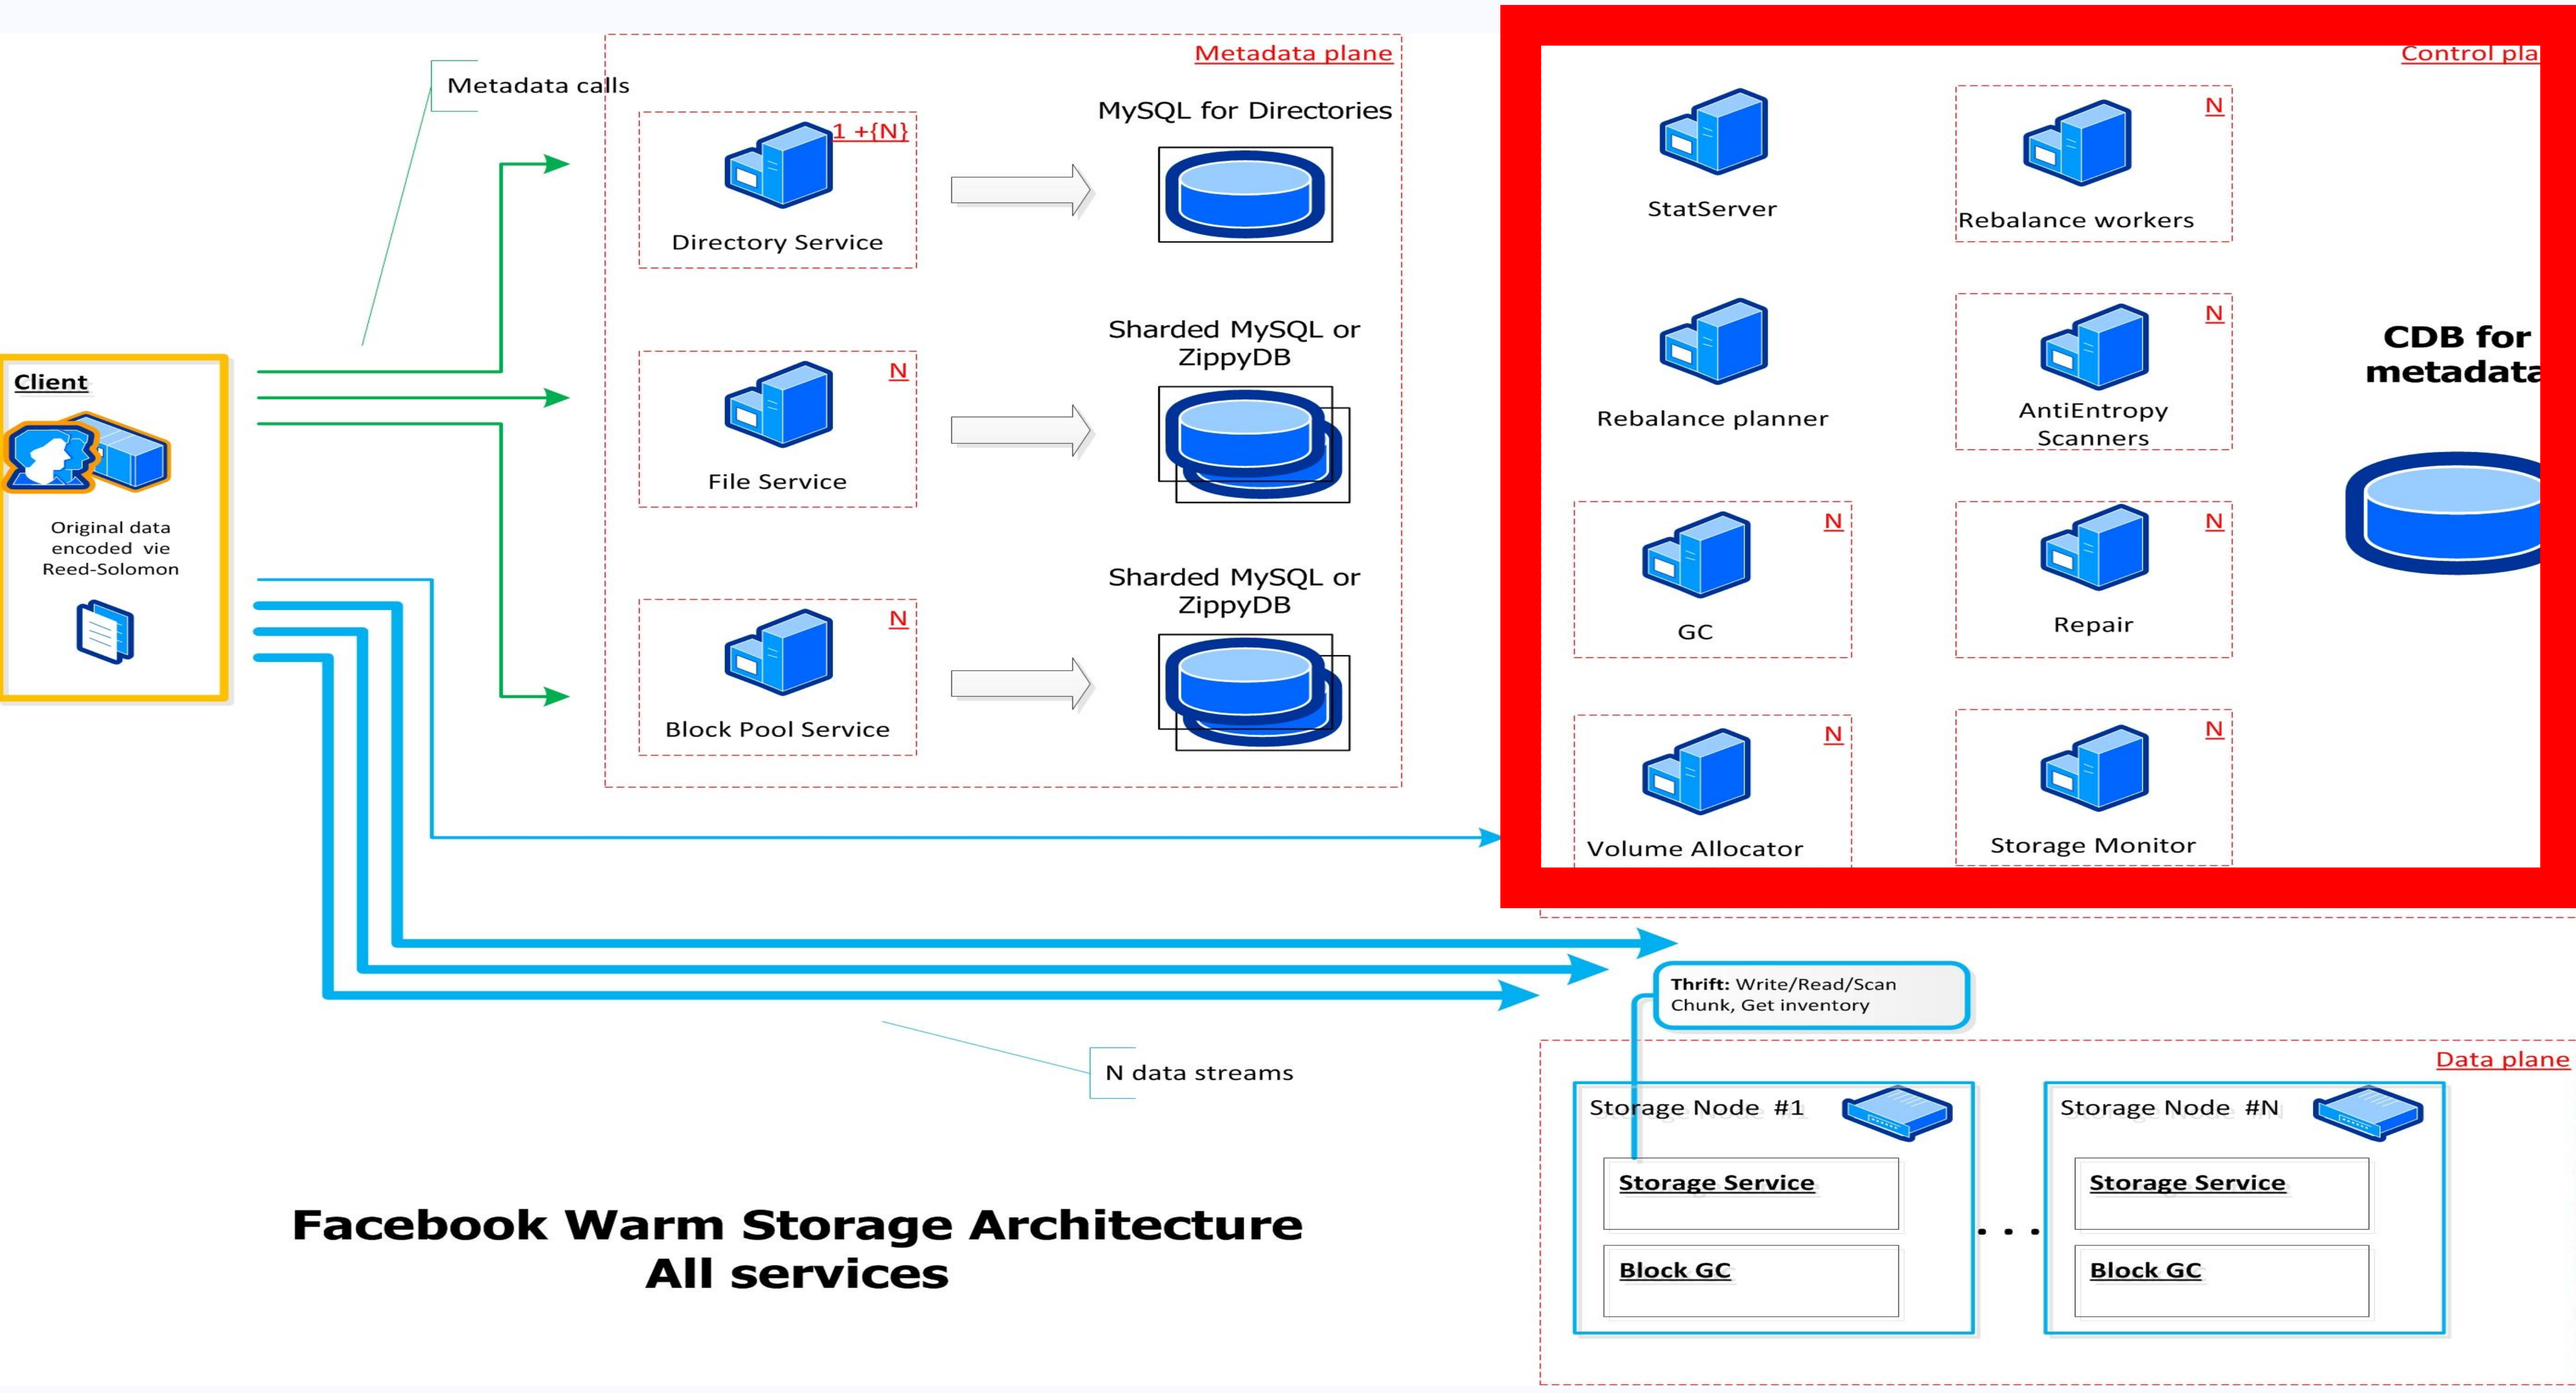

Facebook Warm Storage Architecture
All services

Storage Service IO Model

HDFS Data Node I/O Model

- HDFS Disk Thrashing
  - 4k – 8K I/O threads

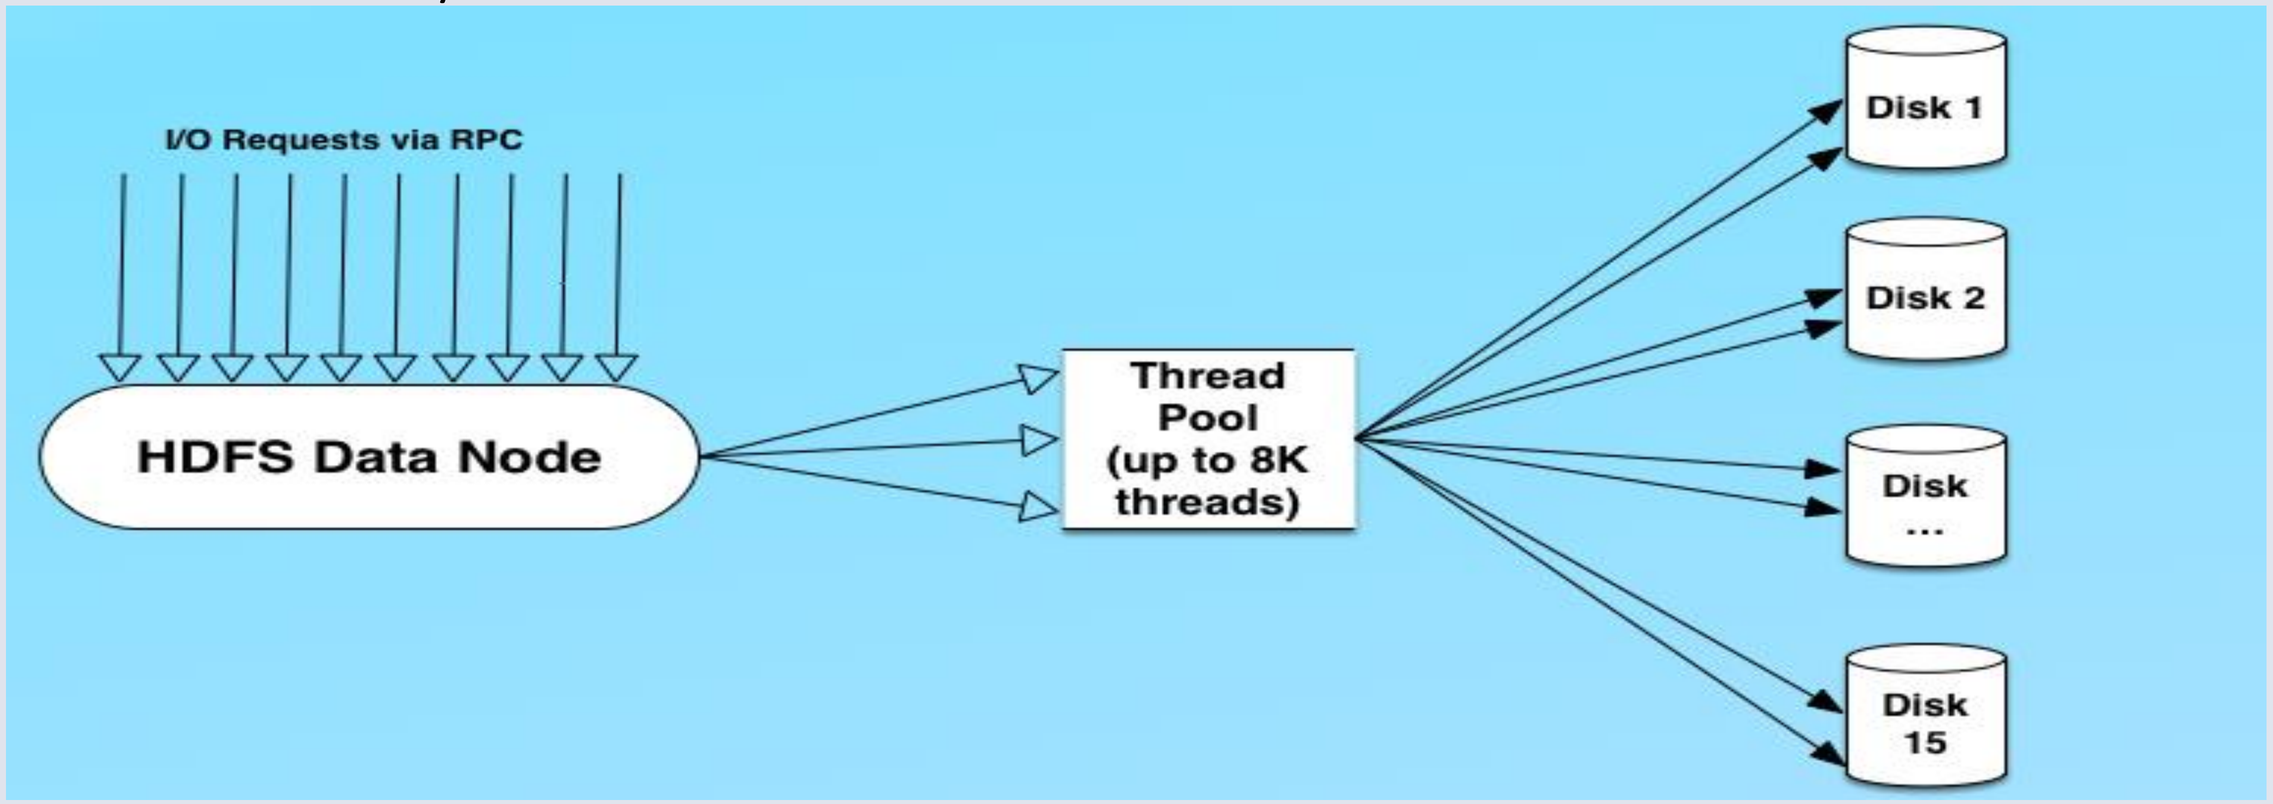

HDFS IOPS/Node – very large cluster

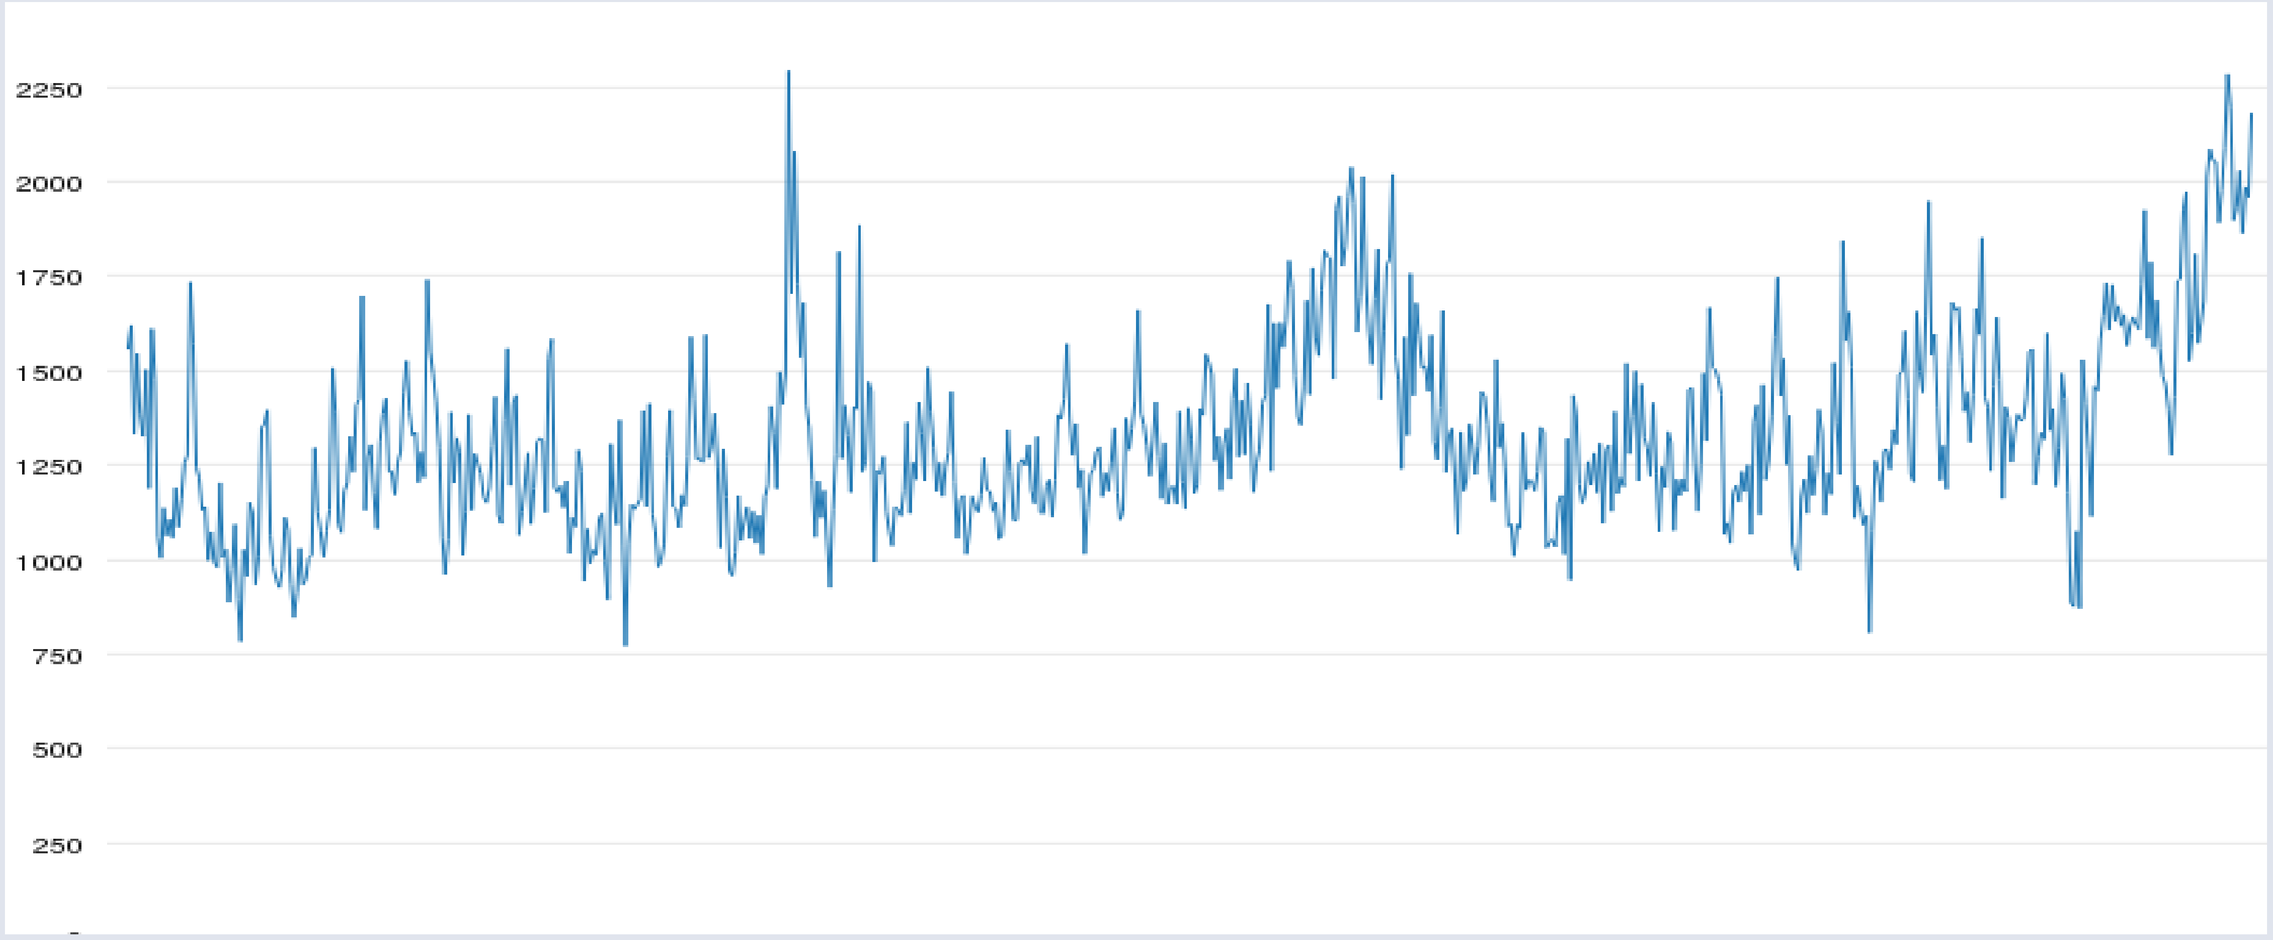

HDFS (MB/s per Node – 15 disks)

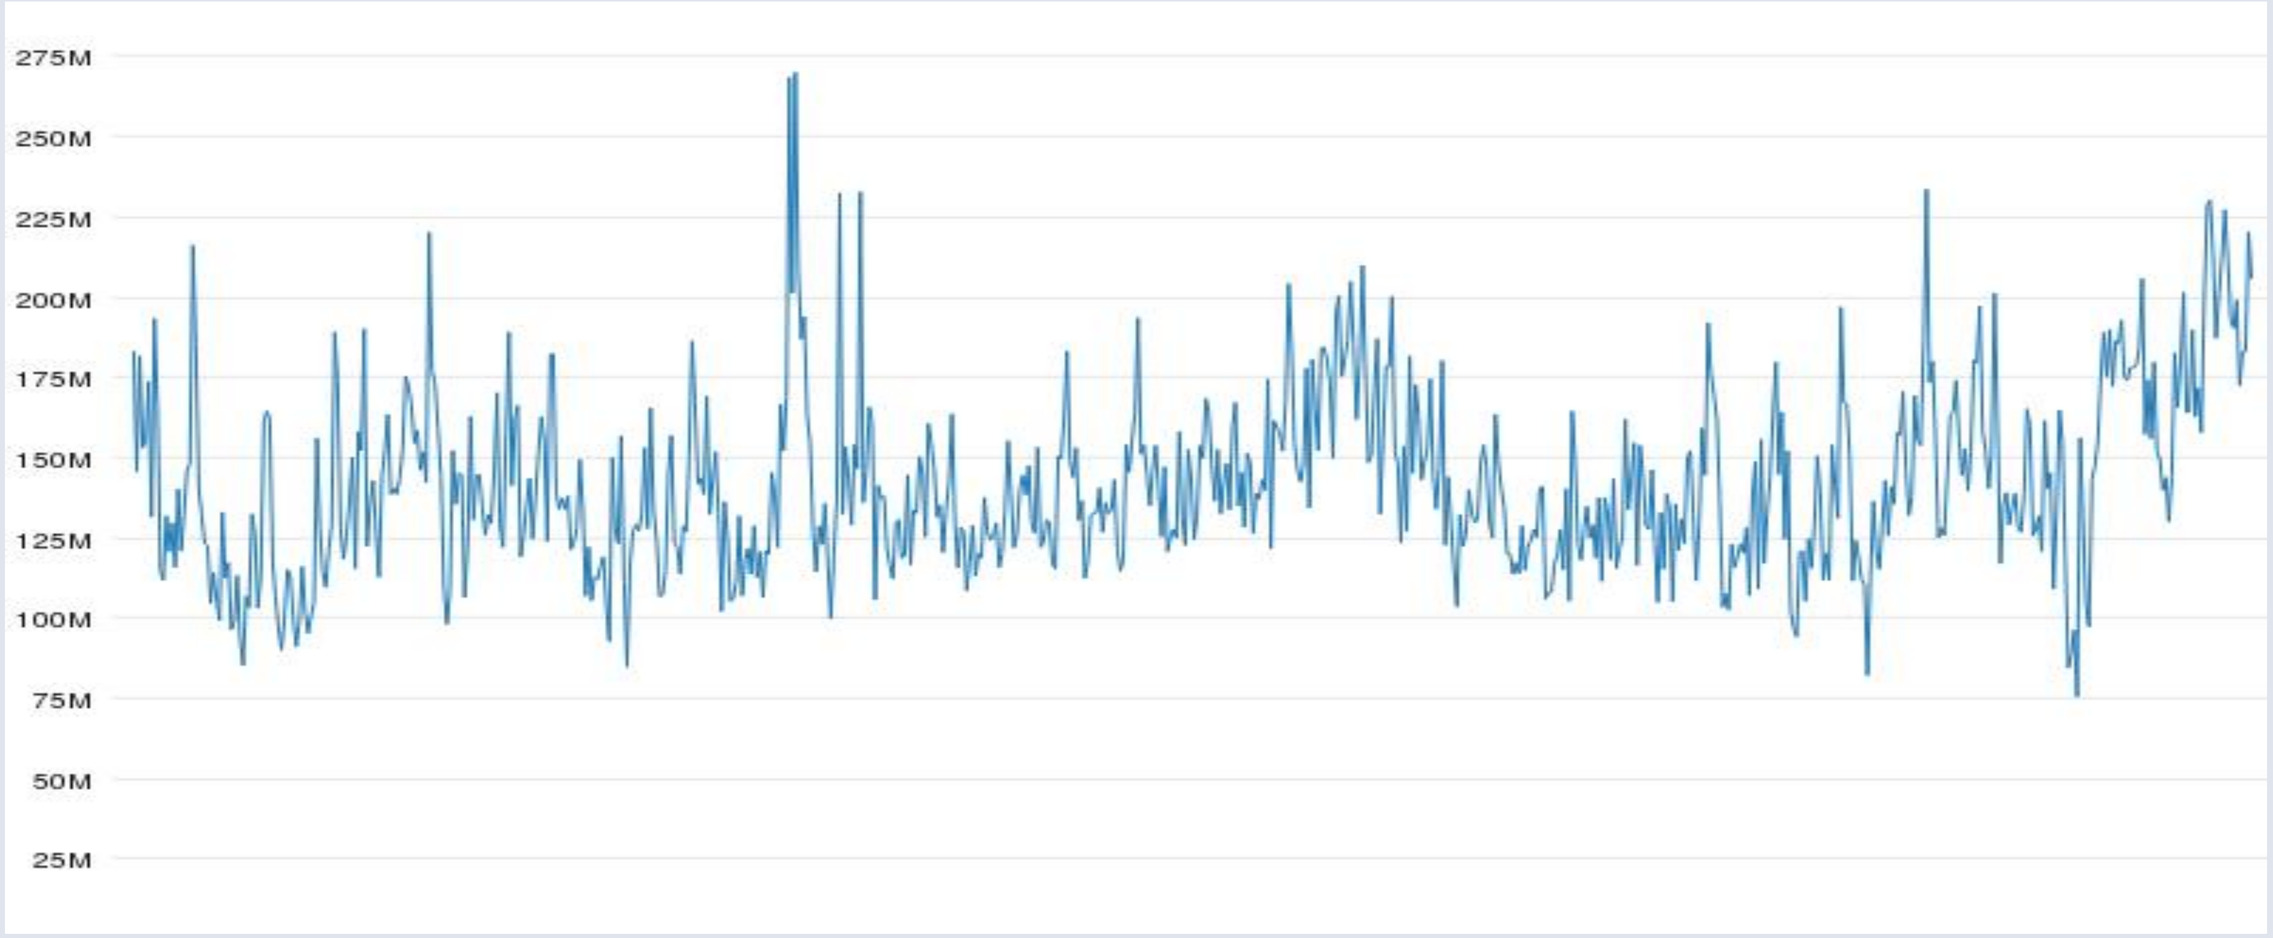

Storage Node Single Thread I/O Model

- 1 thread per disk
- I/O operations are pushed to a priority queue

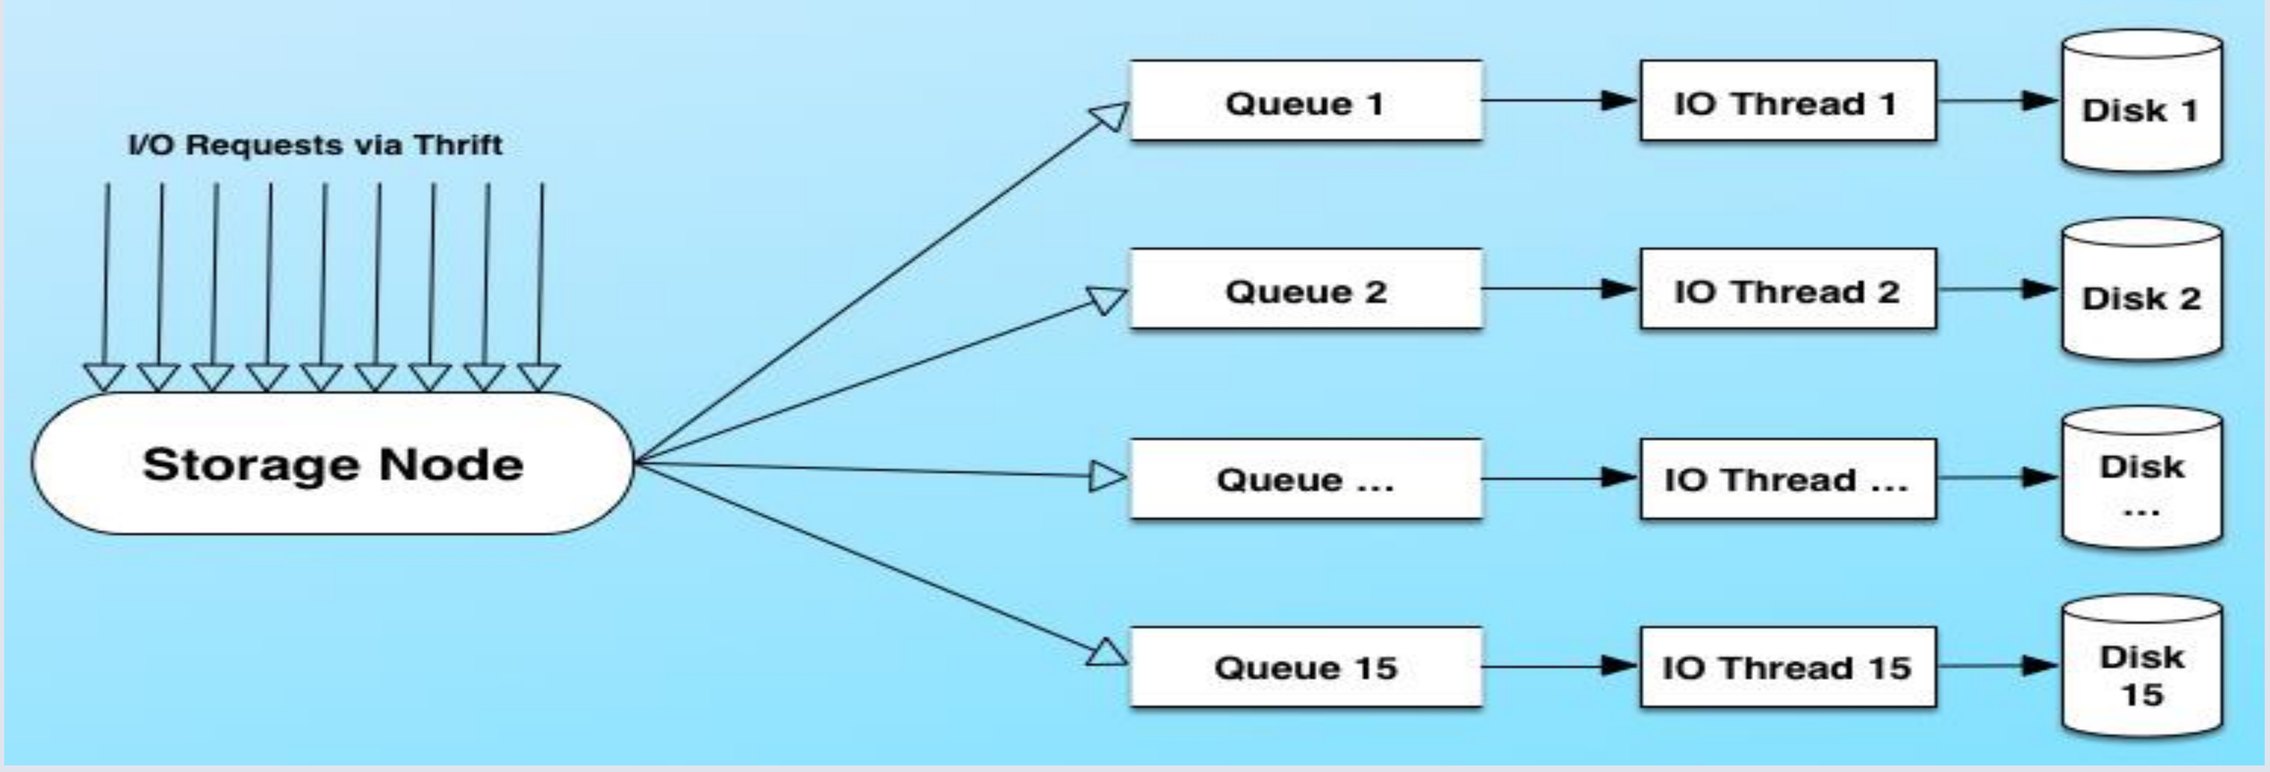

facebook

TempFS and Dis Aggregated storage

Co-located HDFS Map Reduce

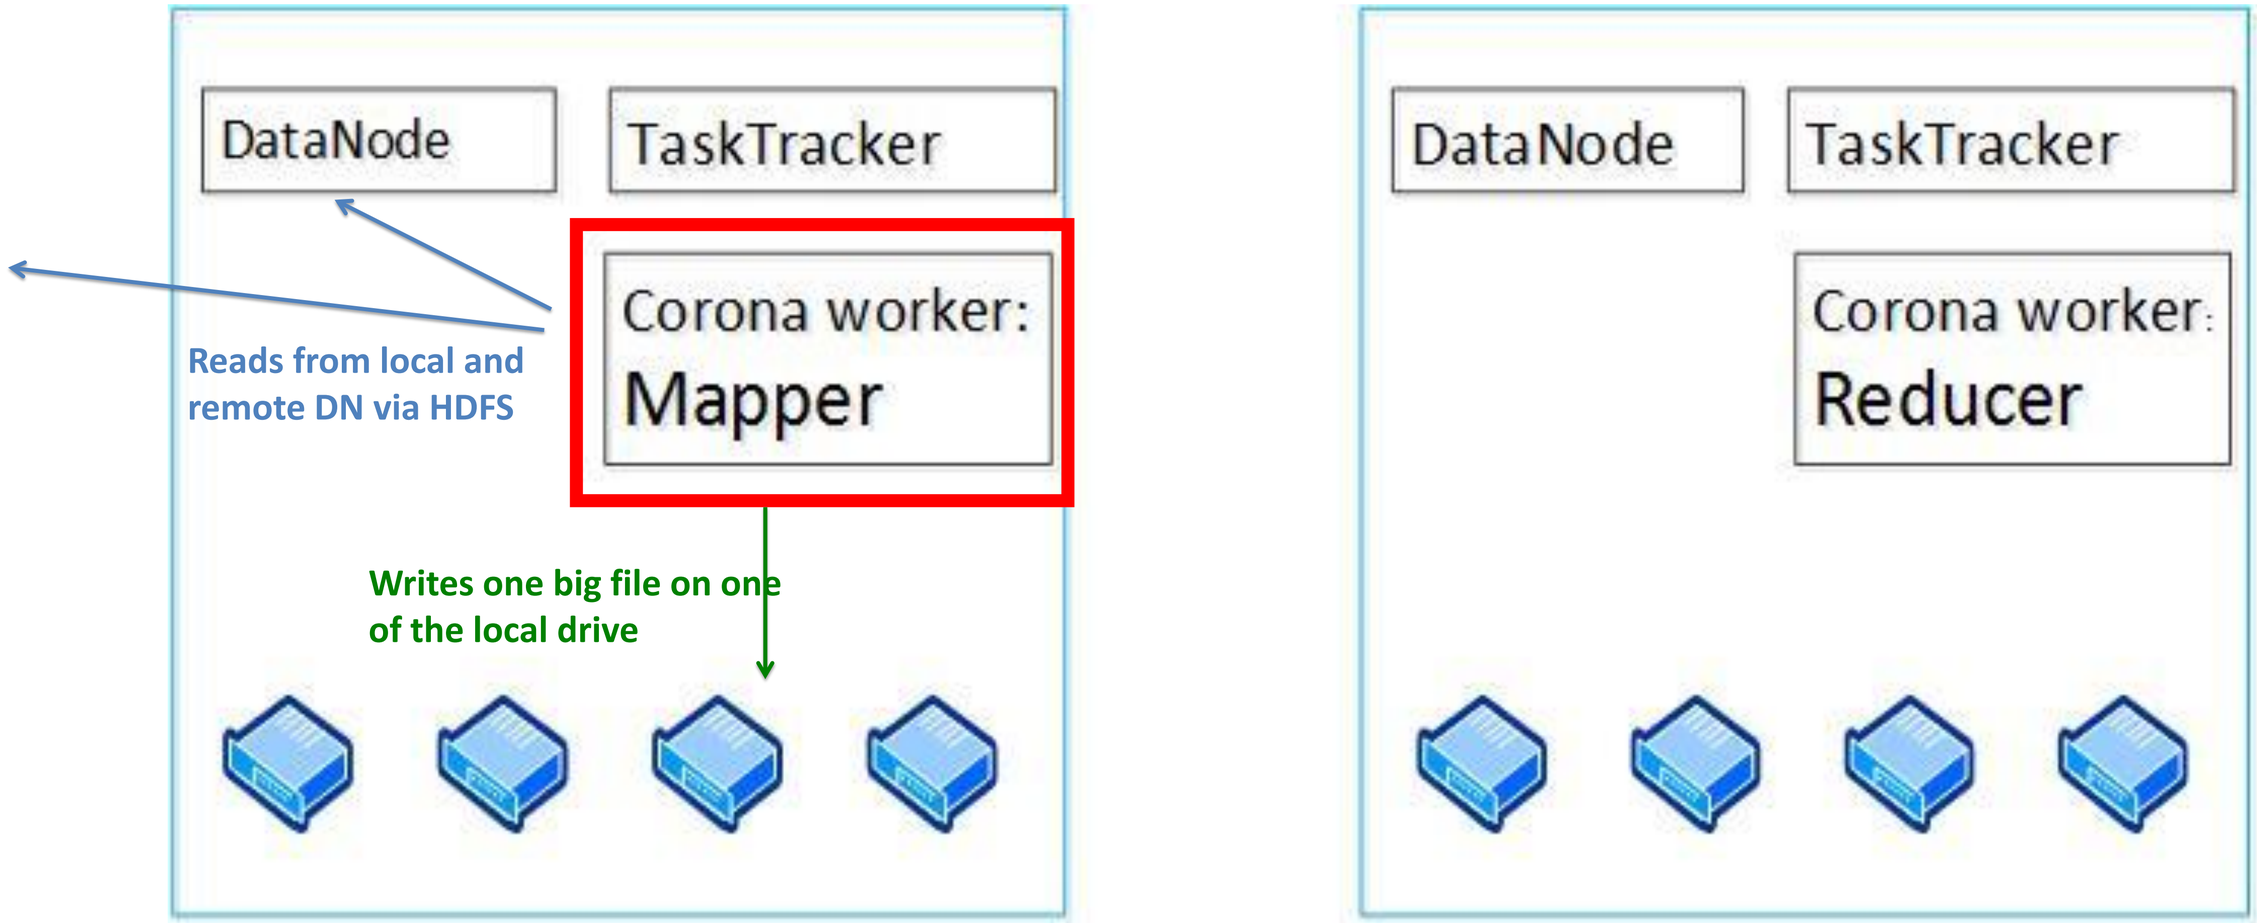

Co-located HDFS Map Reduce

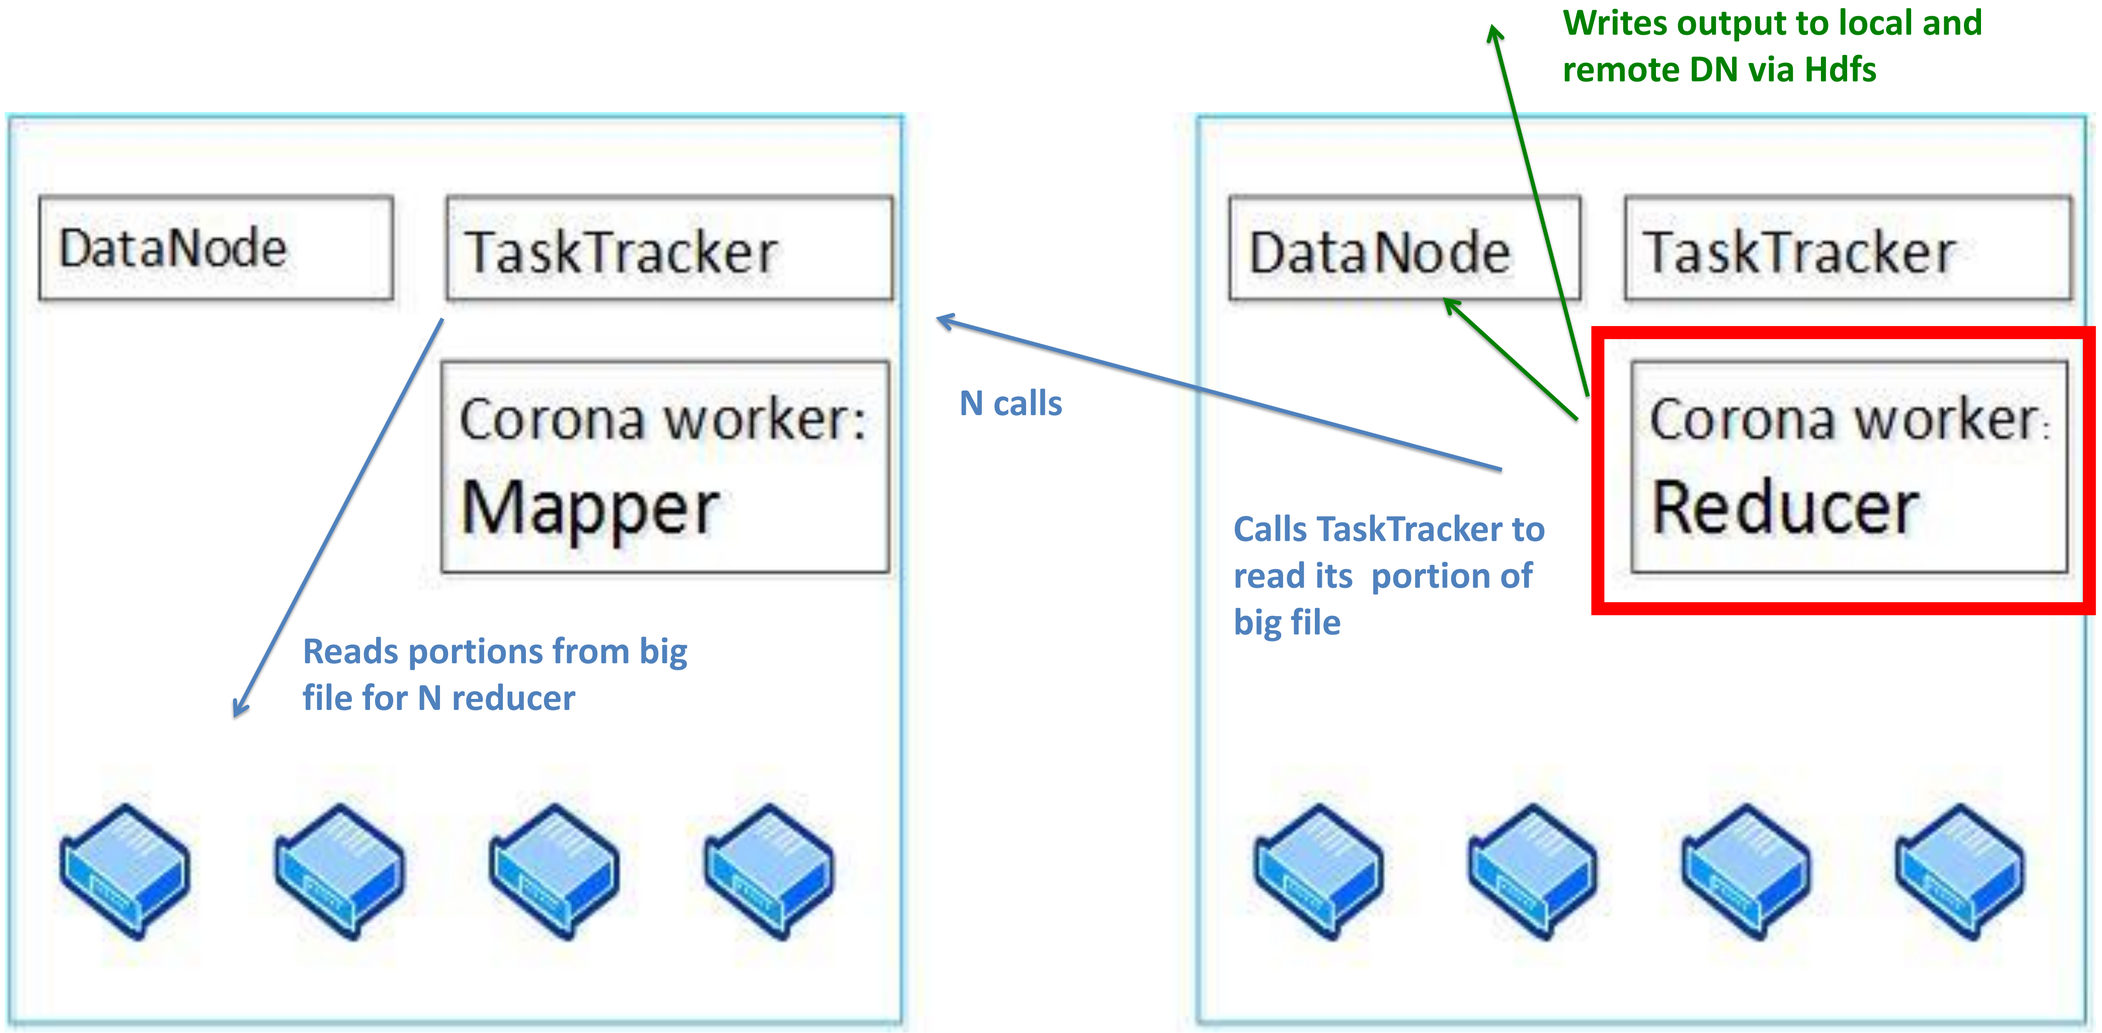

Disaggregated Map Reduce

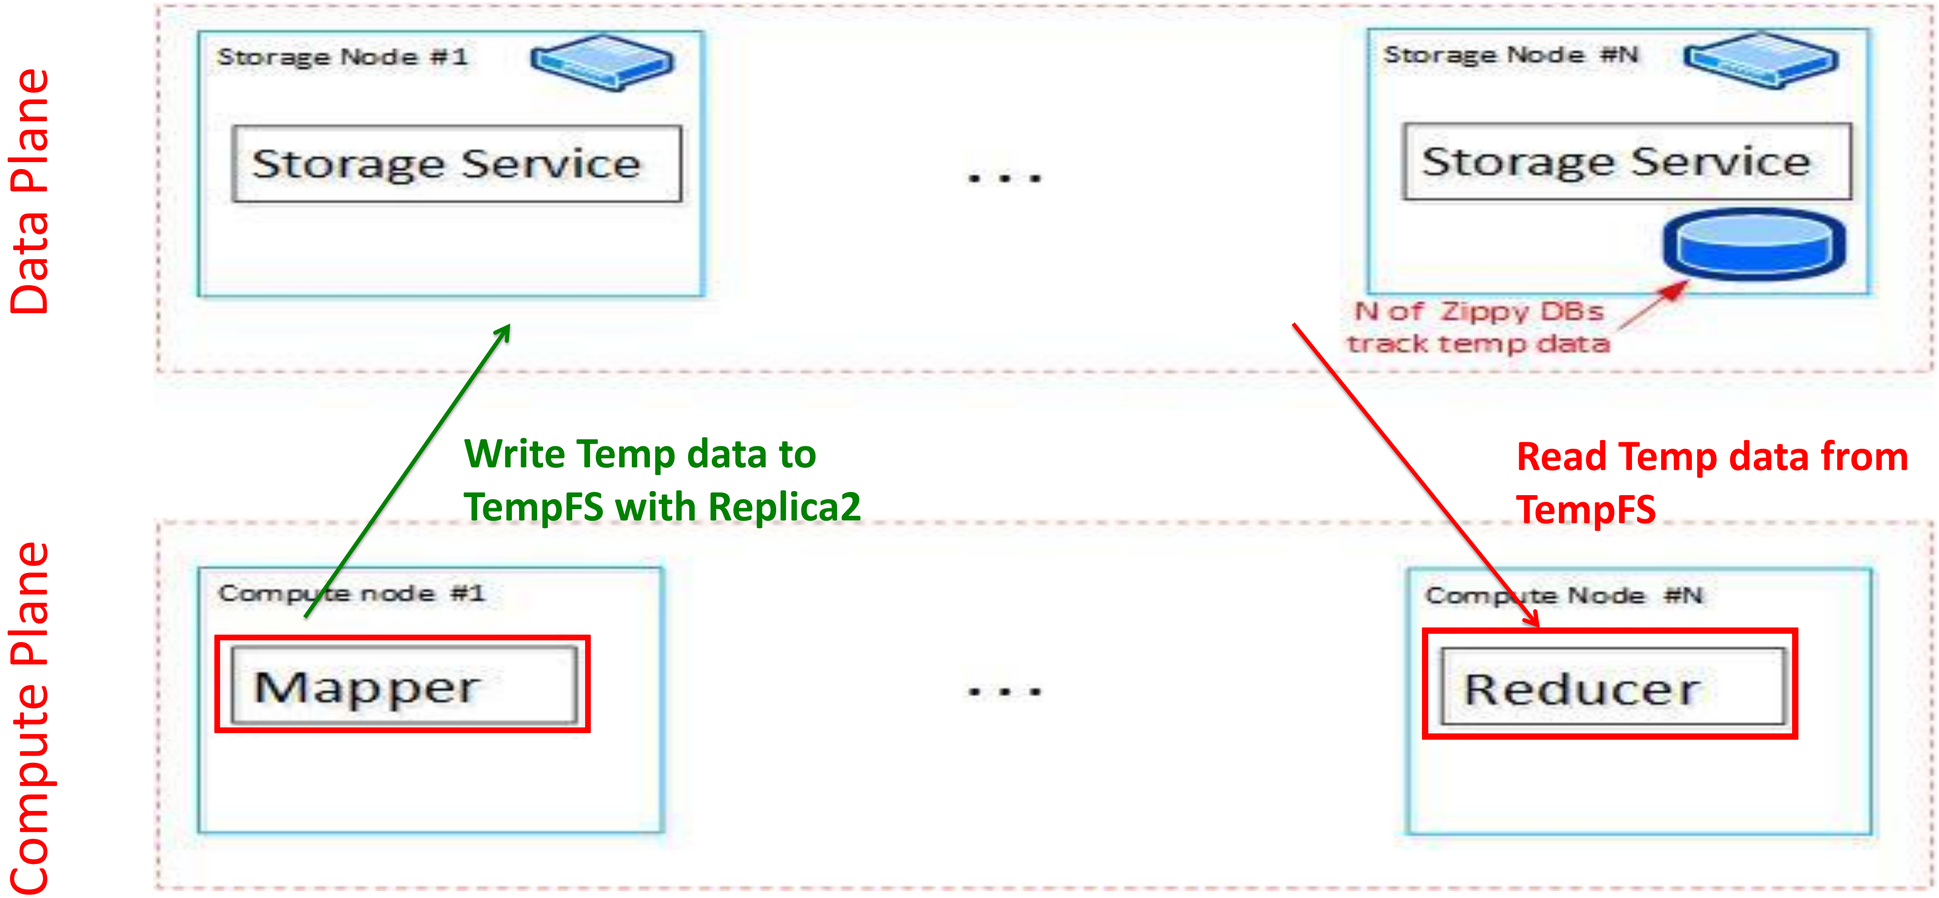

TempFS Architecture

Changes in Warm Storage Architecture for TempFS

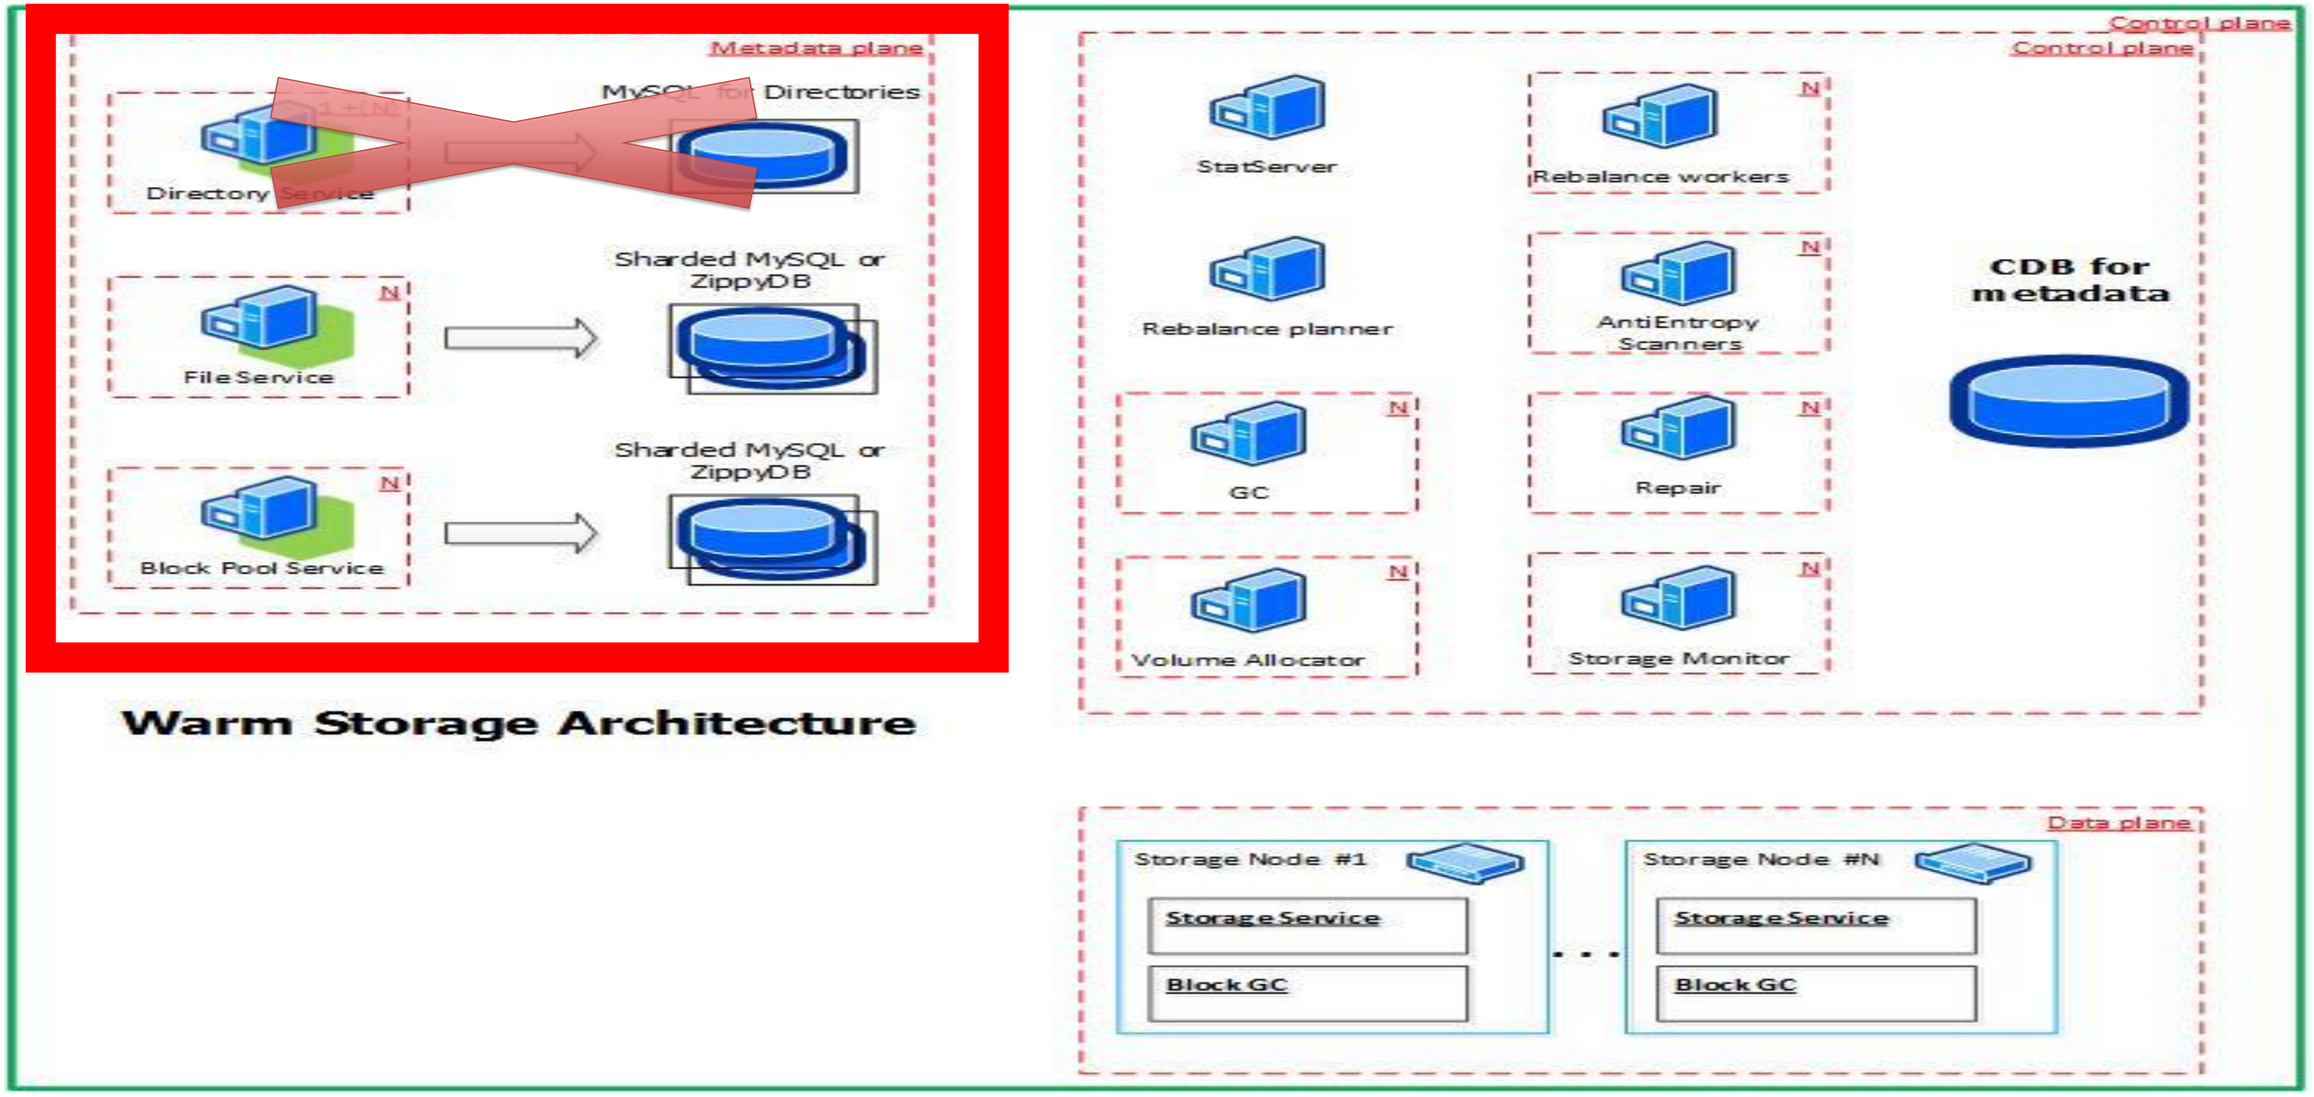

Changes in Warm Storage Architecture for TempFS

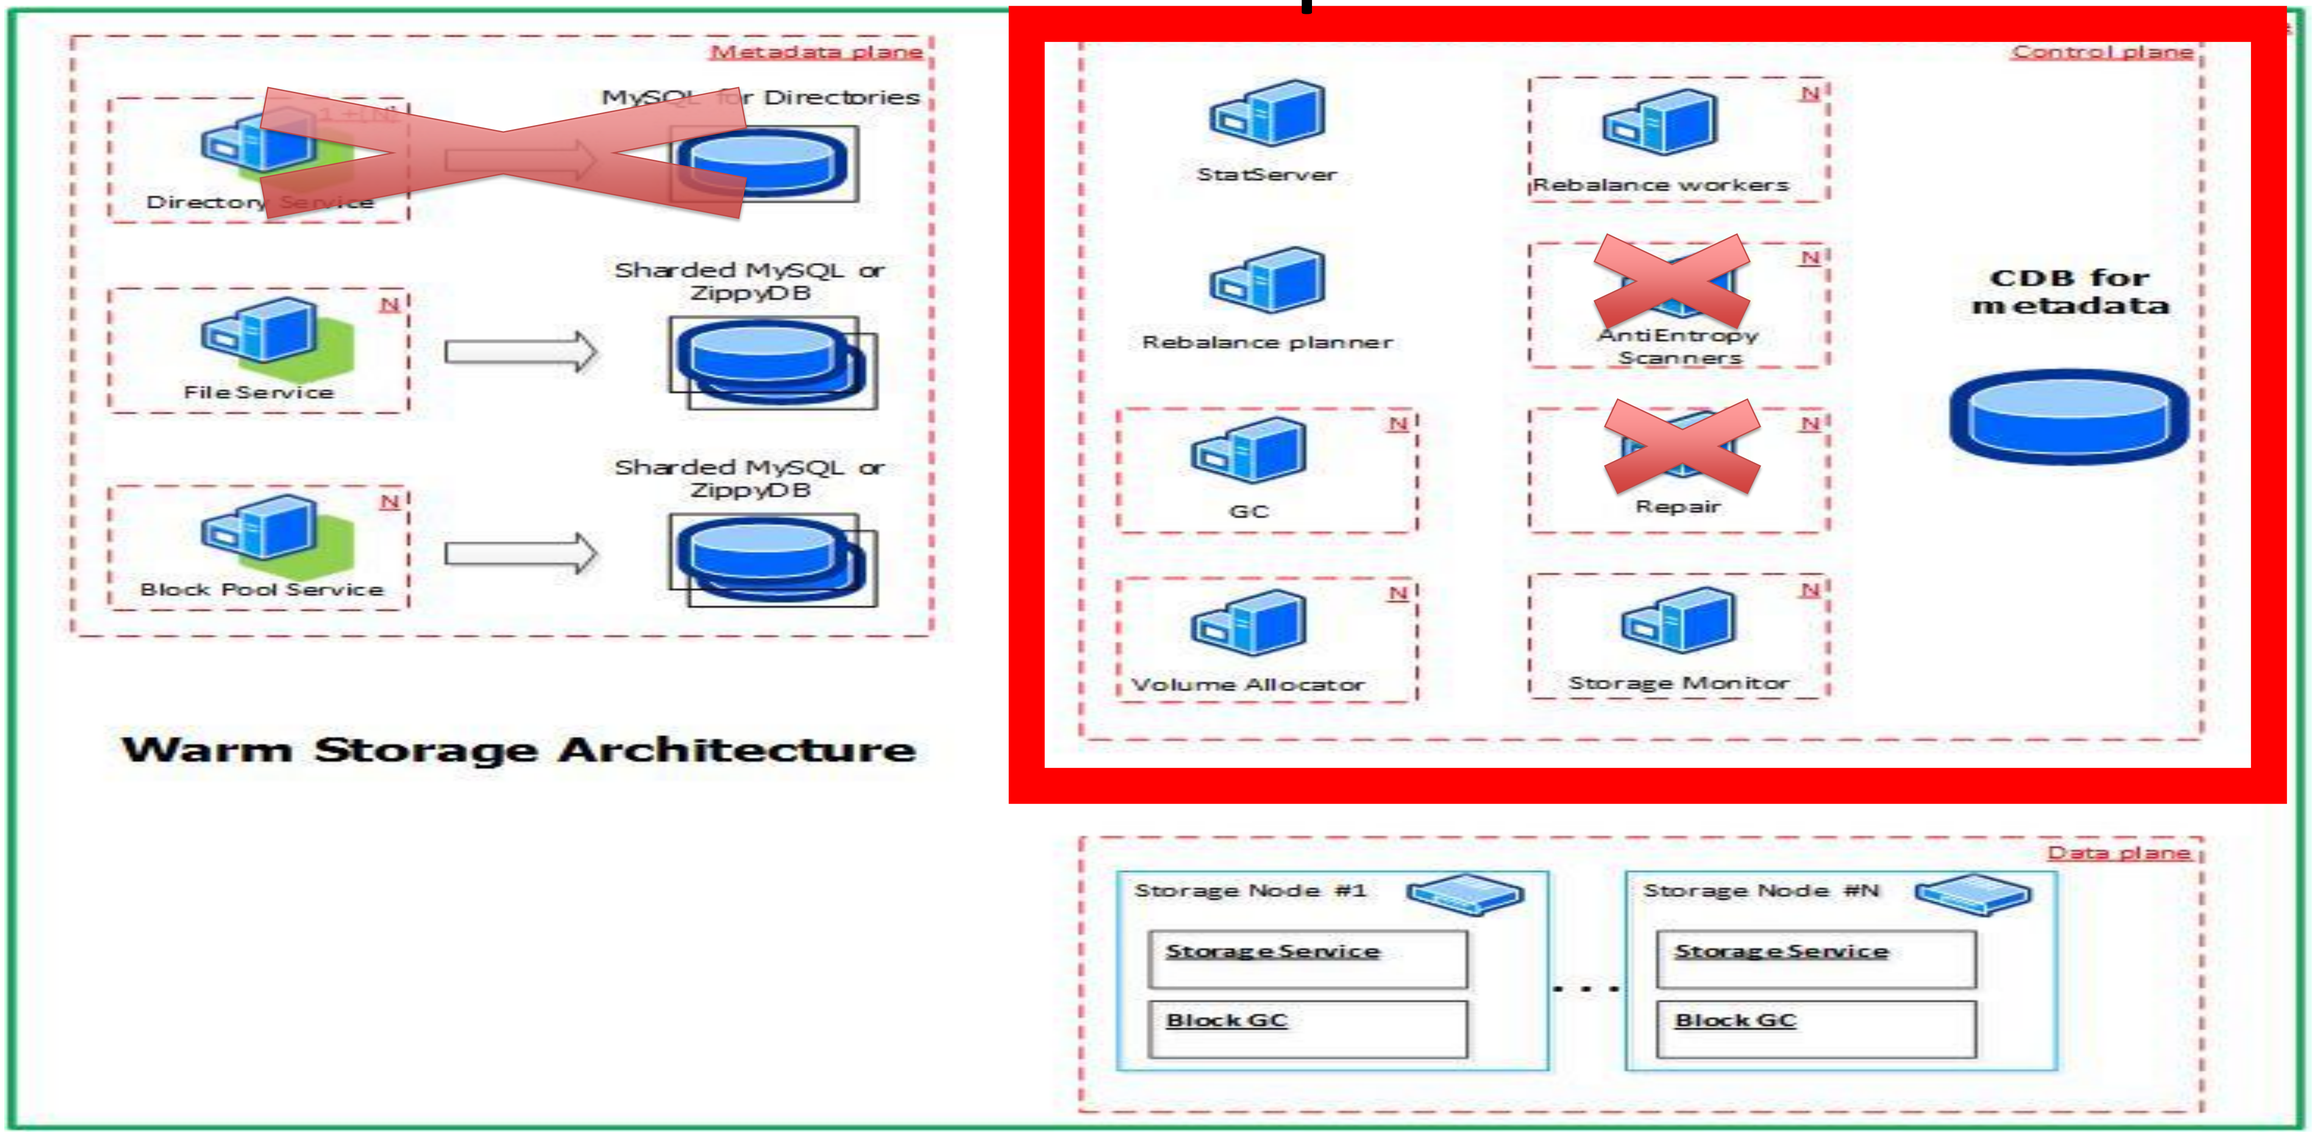

Changes in Warm Storage Architecture for TempFS

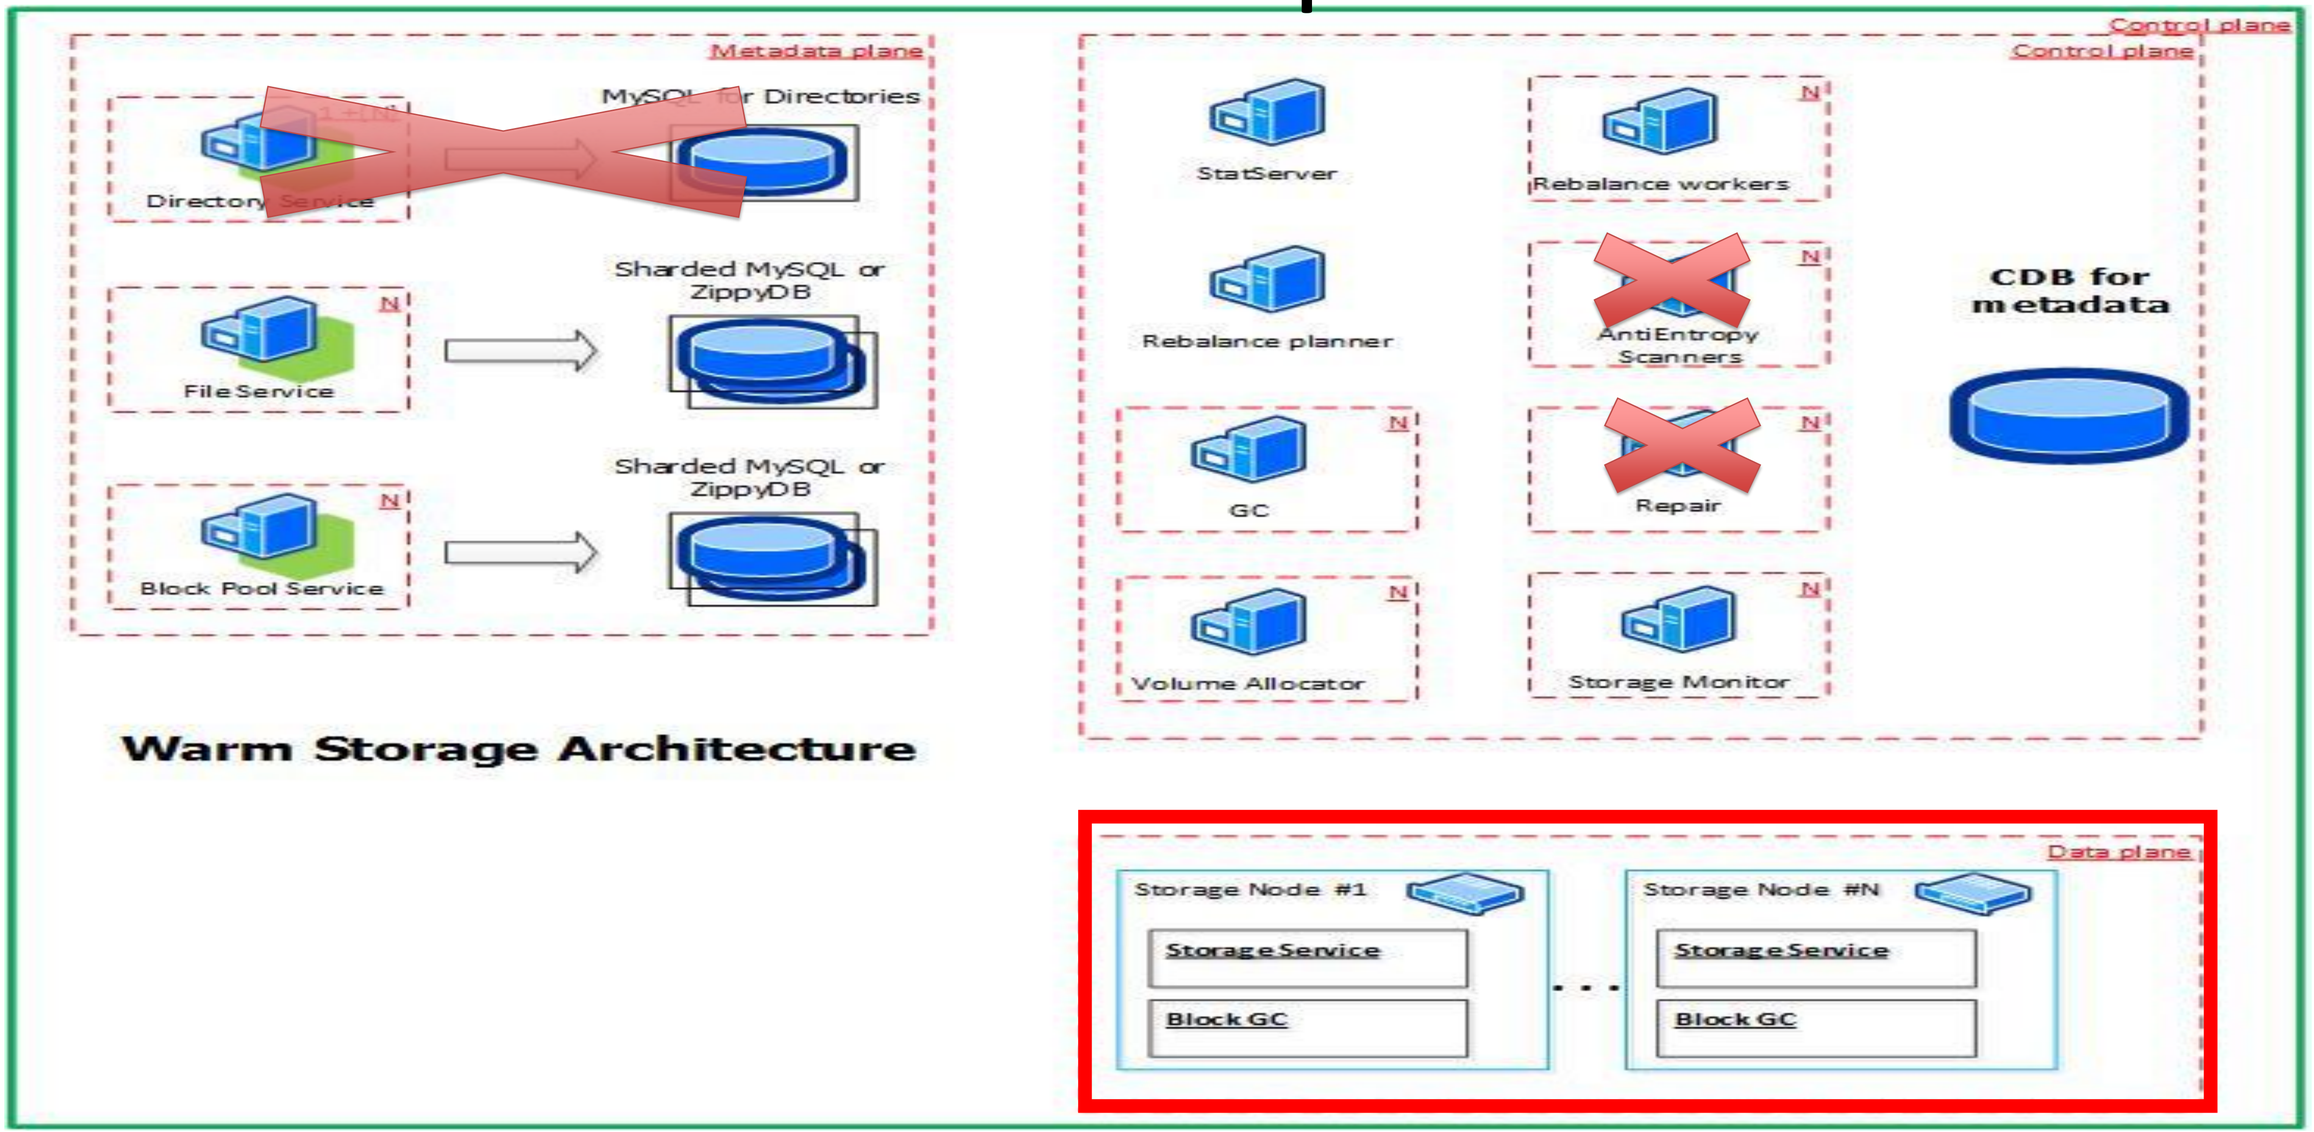

Future Work

- Streaming protocol
- Optimizing corona for Disagg

facebook

Conclusions

• Conclusions

1. The end is close for Moore’s and Kryder’s laws
2. But networking is still improving
3. Efficiency at hyper scale is hard

4. Storage and Compute separation gives better choices and helps with efficiency