

New Use Cases and New Technology Breathe Life into an Older Platform

Flape – The Next Generation

Storage Alchemy LLC




Key Industry Themes



TCO/ROI



***Big Data
Analytics***



Convergence



SSD/Flash



Mobile



Cloud

Backdrop for Data Access

- ◆ Data growth is unavoidable
 - ◆ As much as 65% YoY
- ◆ $24 \times 7 \times 365 = \text{datacenter reality}$
 - ◆ $24 \times 7 \times 365 = \text{data access requirement}$
- ◆ New requirements to keep more data for longer periods
 - ◆ Cost for management, media, bandwidth multiply
- ◆ Globalization
 - ◆ 100% data availability / access



Data Islands to Data Continents

Transactional Databases and Analysis

(application data, block level, virtualization)

- I/O intensive
- Random read/write
- Large/Small files
- Modest storage growth
- Steady growth rates
- Mission Critical
- Block-level virtualization
- Structured data (mostly)

Analysis



20GB

Database



8MB

Persistent Data

Files, Data Protection and Archive Data

(user data, file level, abstraction layer)

- Large files
- Very large storage
- Infrequent access
- Event driven
- Reference content
- Business vital
- Created but not modified
- Data accumulation
- Data integrity
- Long-term retention
- Explosive growth

Backup



10MB

Replication



20MB

Maps



60MB

Video



300MB

Imaging



48GB

Document



80KB

Transactional data

Persistent data

Backup

DR Copy

Amount of data in the typical enterprise

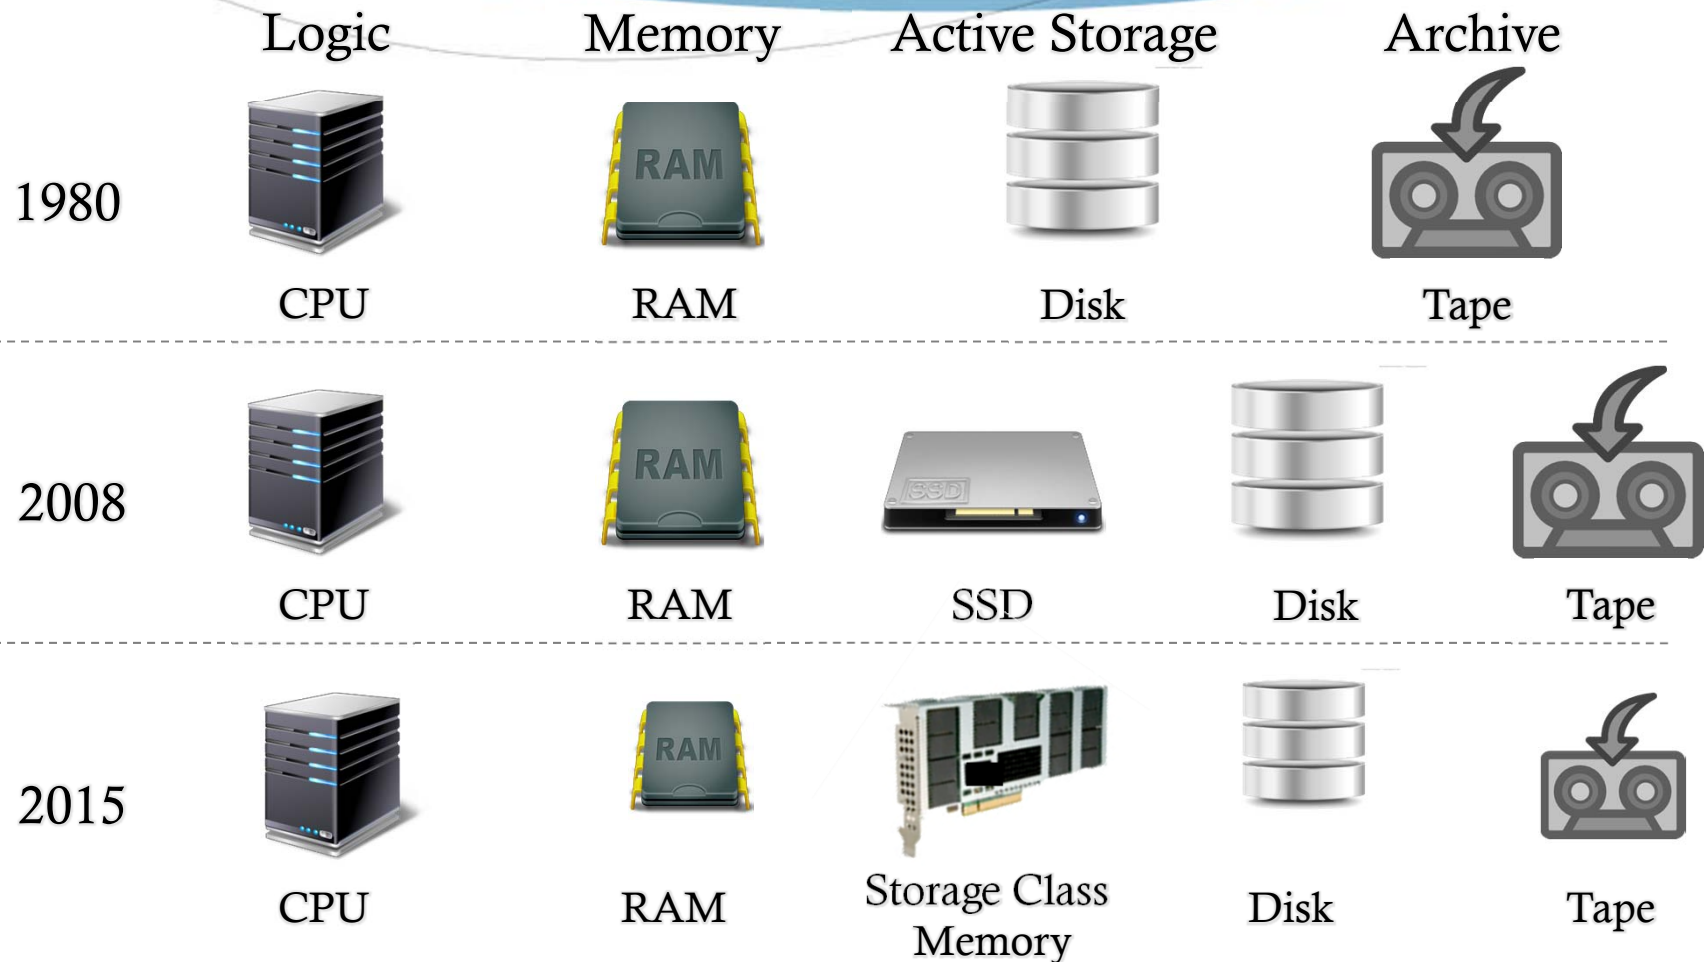
Within 30 days the majority of data becomes “persistent data”.

What Customers Want...

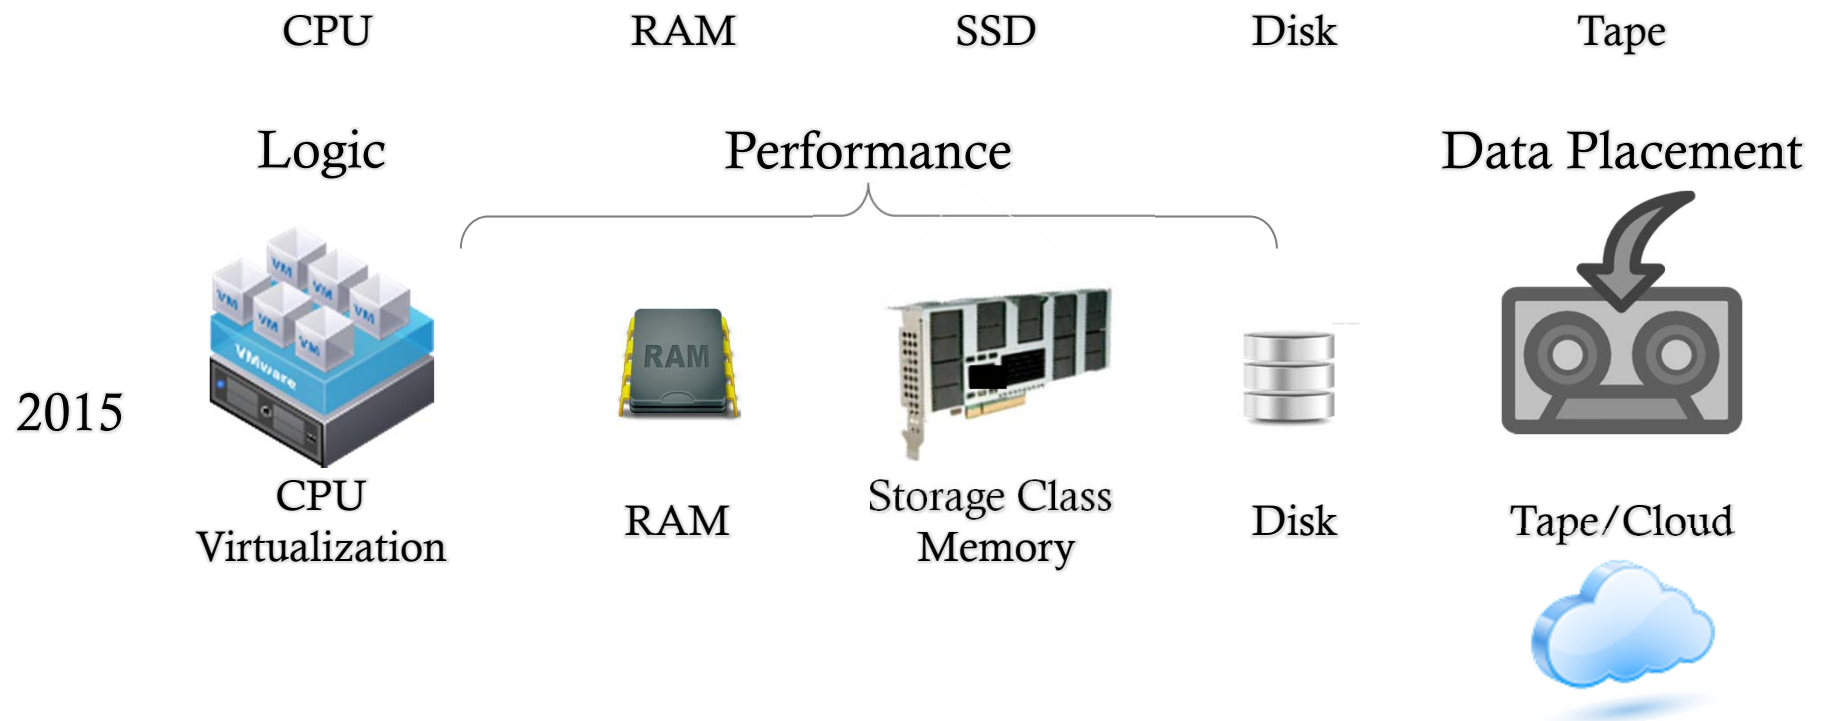
- ◆ Cost Effective
 - ◆ Storage
 - ◆ Software
 - ◆ Including maintenance
- ◆ High Performance
 - ◆ Meet business requirements
- ◆ High Reliability
 - ◆ Very long term retention (100 yrs.)
- ◆ Easy Data Management
 - ◆ PB scale repositories



THE SHIFTING ROLE OF STORAGE TECHNOLOGY



THE SHIFTING ROLE OF STORAGE TECHNOLOGY



New / Interesting Architectures “Flape”



Flash

- ◆ **Best overall \$/IOP for performance**


- ◆ Flash bandwidth performance CAGR 43%
- ◆ Lowers license costs by 38%
- ◆ Lower utilities by 75%
- ◆ Lower operational costs 35%



- ◆ **Best overall for long term data retention**

- ◆ 5x lower cost
- ◆ 100,000x more reliable than disk
- ◆ Bandwidth CAGR 30%
- ◆ Areal density CAGR 30%
- ◆ Data retrieval rates 4x disk*

*=TTFB



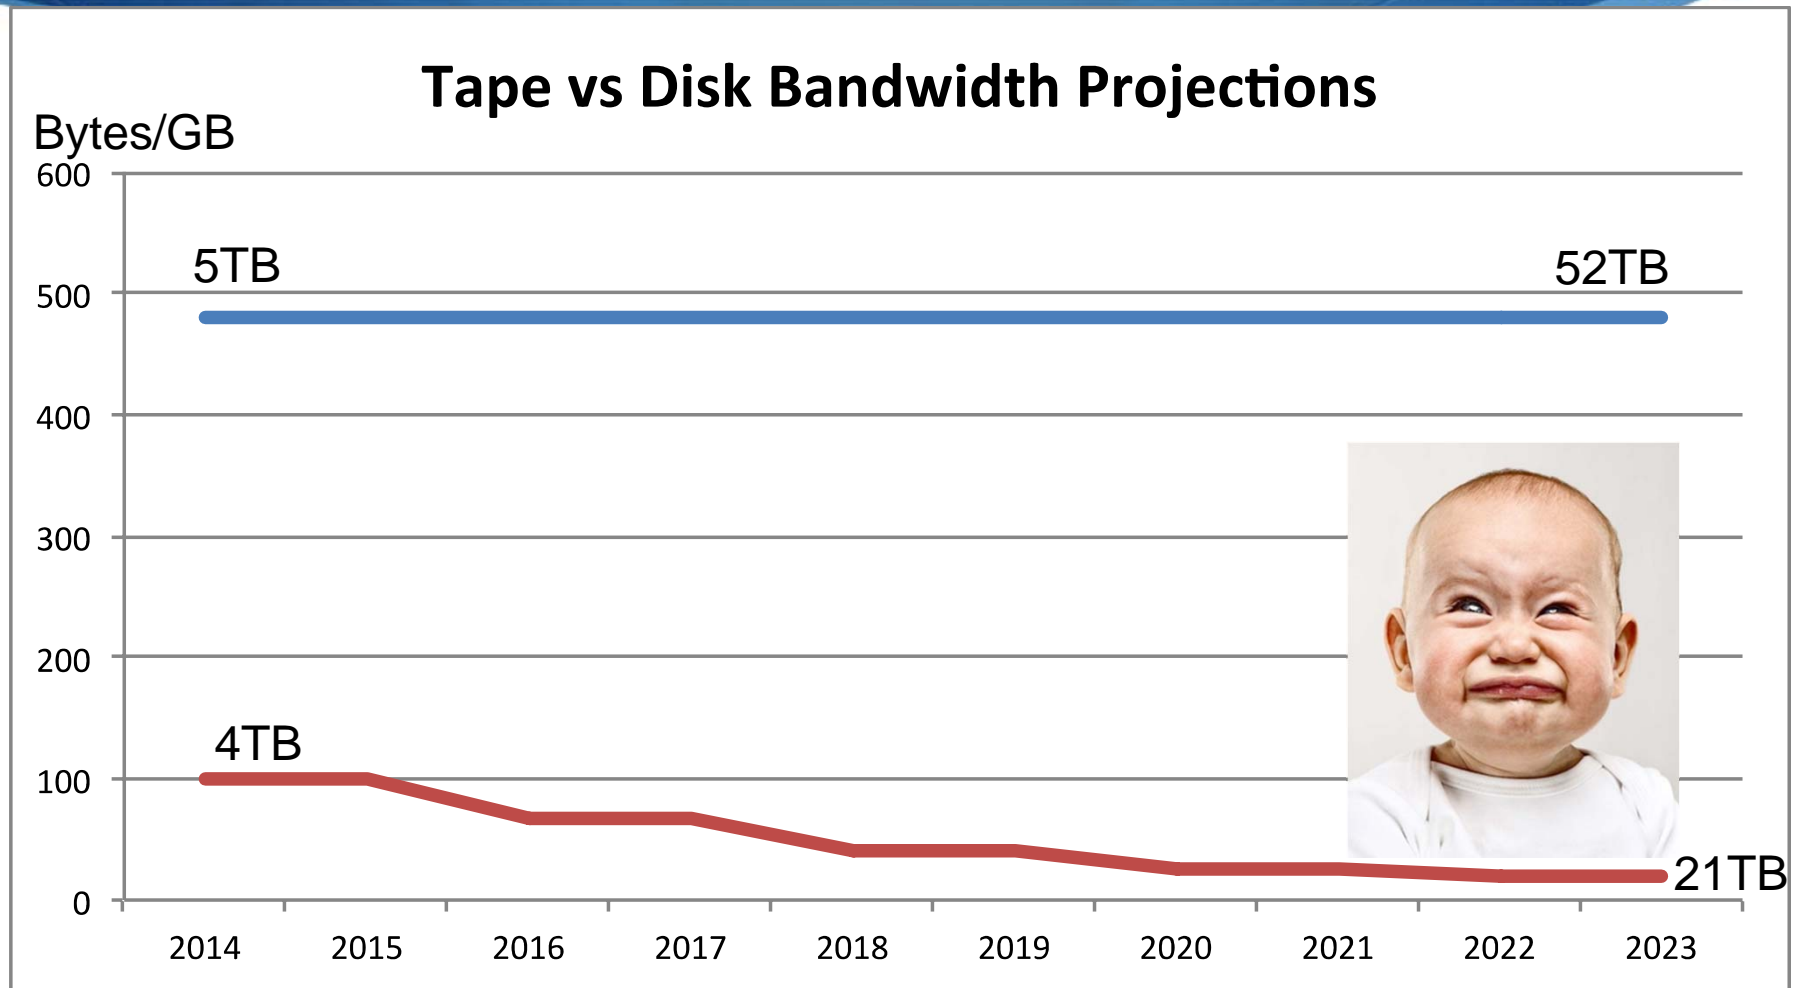
Tape

IDC Predicts – TAM

| Market | 2014 | 2015 | 2016 | Total |
|-----------------------------------|------------------|------------------|------------------|------------------|
| Cloud – Public / Private | 6,800 PB | 10,200 PB | 14,400 PB | 31,400 PB |
| Long-term Preservation - Capacity | 4,900 PB | 9,200 PB | 16,500 PB | 30,600 PB |
| Scale-out NAS – Capacity | 1,600 PB | 2,600 PB | 4,225 PB | 8,425 PB |
| Archive – Capacity | 8.2 PB | 15 PB | 27.8 PB | 51 PB |
| Total Capacity | 13,300 PB | 22,015 PB | 35,125 PB | 70,476 PB |

- 70 Exabytes will be required to store these three data categories by 2016
- All numbers reflect disk capacities – *HOWEVER* IDC says that tape could and should play in this market
 - LTFS is an enabler to making tape much more significant in these data storage categories
 - Today LTFS is in it's infancy stage which is why IDC doesn't cover these markets with tape predictions
 - LTFS has been approved by the LTO Consortium as an industry standard
- Tape should garner 10% (at a minimum) of this market space = **700PB**

Tape Technology is Improving over Time




*Bandwidth off the media

Source: Wikibon: 2014

Comparative Costs of Flape

Total 10-year Cumulative **Hardware** Cost for 1PB Long-term Retention CAGR 30%



Source: © Wikibon 2014, from Numerous Sources including Analysts, Consultants, IBM & Oracle.

Key Observations & Recommendations

- ◆ Without a “write to flash first” architecture, the cost of software and hardware at scale become prohibitive
- ◆ Today, need to understand the limitations of the available ‘storage services’
- ◆ This ripples thru to opex because the amount of effort required at scale to tune and manage the system
- ◆ Traditional arrays today don’t scale well- i.e. Ones that use “Flash Cache” (reads) and SSD (writes). The problem w/SSD is you have to manually allocate the volume – too hard and too expensive
- ◆ In general, the more integration that can be done closer to the application...costs will be lower and value higher
- ◆ Software solutions w/ metadata support will win out

Key Tape Takeaways...



For many archiving apps – Tape is a much better investment than disk



There is no such thing as “fast disk”



Tape is improving with age – disk isn't



Combining flash w/tape (“Flape”) is a winner for archiving and deep archiving apps

Summary...

- ◆ Tape is far from dead!
- ◆ Tape's “sweet spot” is large objects (e.g. film clips, email blobs, etc.)
- ◆ Software exploitation is the key
 - ◆ Insist that ISVs accommodate tape for archiving apps
- ◆ Tape with flash will yield optimum business value
- ◆ The value will come from exploiting metadata through software

Good Articles to Read

- ◆ [http://wikibon.org/wiki/v/Rise of the Machines: The Rebirth of Tape](http://wikibon.org/wiki/v/Rise_of_the_Machines:_The_Rebirth_of_Tape)
- ◆ <http://wikibon.org/blog/why-tape-is-poised-for-a-george-foreman-like-comeback/>
- ◆ [http://wikibon.org/wiki/v/The Emergence of a New Architecture for Long-term Data Retention](http://wikibon.org/wiki/v/The_Emergence_of_a_New_Architecture_for_Long-term_Data_Retention)



THANK YOU