Designing Winning Enterprise Storage

Curt Beckmann
Principal Architect
Brocade

Today's Presentation Topics

- Bit of Background (me and storage)
- Enterprise Storage Features, expected and not
- What Enterprise Storage Teams Think is Sexy
- Anecdotal evidence of enterprise storage thinking
- Q & A

My experience

- Assembly coding on core machine in HS
- Deep exp w Bubble Memory, UV-EEPROM
- Medical & Test tools developer: get it right
- ASICs @BayNet (packet drops blew my mind)
- @Rhapsody (cool buffer credits, hints of iSCSI)
- 6 yr side trip into Software Defined Networking
- Back to storage… wow, look at all the flash ☺

Storage Types

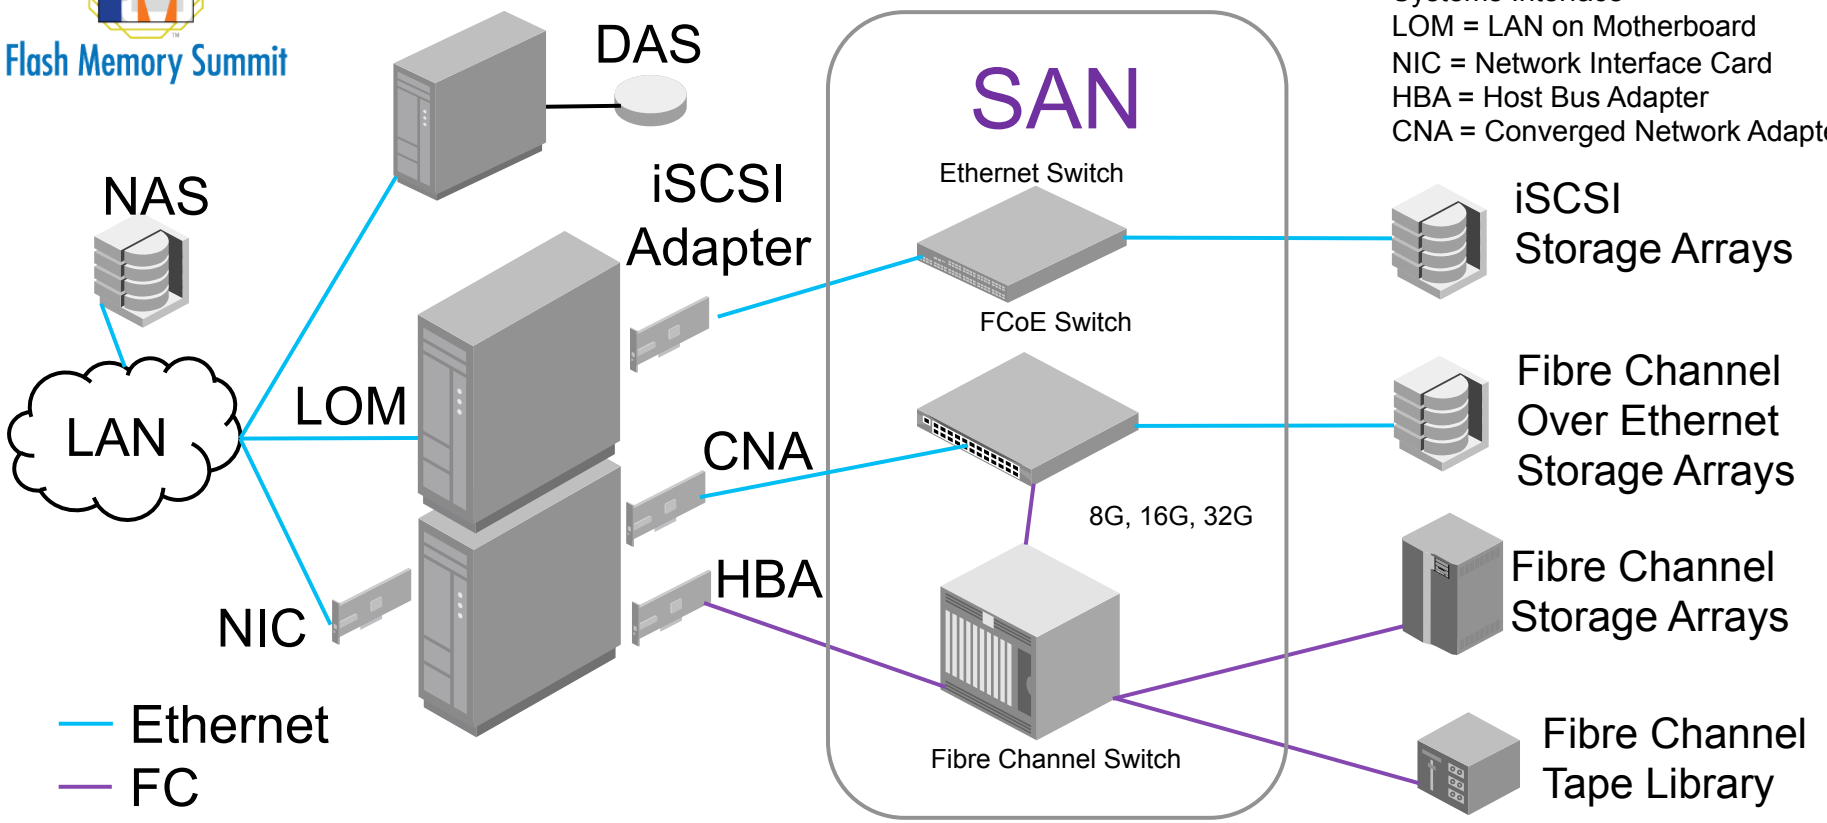

DAS = Direct Attached Storage
NAS = Network Attached Storage
iSCSI – Internet Small Computer Systems Interface
LOM = LAN on Motherboard
NIC = Network Interface Card
HBA = Host Bus Adapter
CNA = Converged Network Adapter

Source: http://www.ieee802.org/3/ad_hoc/bwa/public/sep11/kipp_01a_0911.pdf

Categorization (storage-oriented)

	Cost*	Performance	Reliability	Maturity
Fibre Channel	1.00	High	High	High
NAS (NFS, etc, over IP)	0.68	Low-Medium	Medium	High
iSCSI	0.59	Medium-High	Medium	High
DAS	0.46	High	High	High
Mainframe (FICON)	1.63	High	High	High
InfiniBand	1.43	High	High	Low
SAS SAN	0.70	Medium	Medium	Low
FCoE	0.79	High	Medium	Medium
NVMe over Fabrics	n/a	High**	High**	Low

*Normalized to FC Cost/GB in 2016 prices (Ref: IDC)

**Projected

Enterprise Flash Growing Well

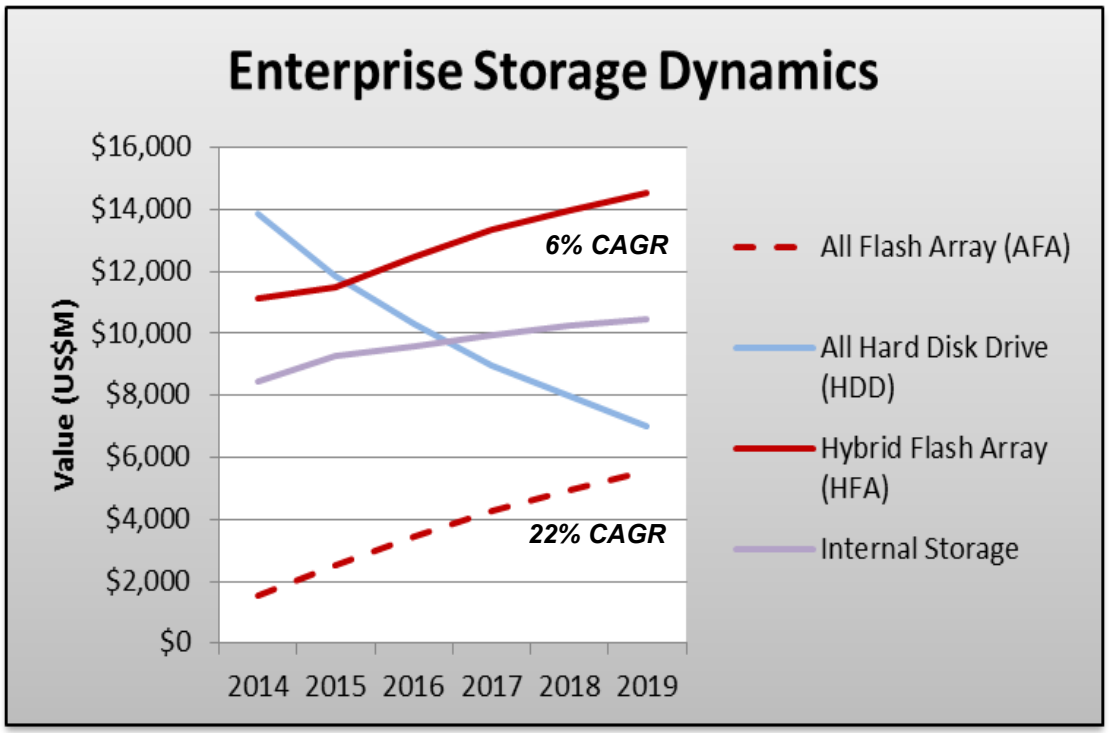

Source: IDC September 2015 WW Quarterly Disk Storage Systems Forecast

Enterprise Features, expected (duh)

- Able to get Really big
- Able to go Really fast
- “Standard”
- Reliable / available

Enterprise Features, more nuanced

- More worried about OpEx than CapEx
- Low risk across many forms of risk
- Familiar / comfortable (surprise-free)
- Predictable, structured
- Highly interoperable (not just “certified std”)
- Pre-tested, fingerpoint-free, 24x7 support
- Backward compatibility / long-term support

What ES teams think is sexy

Image credit: //www.pinterest.com/pin/223069606256631992/

- Low risk, repeatability is exciting
  - Belt-and-suspenders is okay fashion-wise
  - "hot" upgrades are dreamy (but test in lab 1st!)
- Staying home Sundays is a thrill
  - Ditto for sleeping the whole night through
  - Incremental adoption paths simplify life
- Happiest if "customer" never thinks about them
  - Have nightmares that some mistake will make the headlines

ES Team Mindset differences

- DRAM ECC detects multi-bit, corrects 1-bit errors
  - Multi-bit errors happen sometimes… to recover, rerun the app
  - Makes sense: computation has many exposures besides ECC
- Good idea to deploy this level of ECC on servers
  - Not required to deploy such ECC on laptops
- E-Storage always protected… RAID / Erasure-coding
  - Makes sense: compute / memory recovery leverages storage
- E-Storage has an SLA; Memory is part of compute
  - “do over” is “normal” in compute (as well as TCP)

Mindset differences mean… #1

- No single point of failure
  - Dual fully-redundant SANs
- Bias toward optics over copper
- Reviewing vendor test matrices
  - Buy SAN kit from storage vendor (ensure unified support)

Mindset differences mean… #2

- Careful, methodical planning and processes
  - Writing, testing scripts that may only be run once
  - Migration plans include clear rollback plans (because “duh”)
- Write caching, etc, needs battery to handle power fail
- Asking Root Cause Analysis after any device failure
  - Then asking why all other similar devices won’t also fail

Anecdotal evidence

- Super performant “Enterprise” flash product fizzled
- Firm buys flash array, cuz vendor priced it best
  - Some must have flash, others buy it too, if robust, full-featured
- Compelling storage start-ups can only go so far solo
  - “Successful” ones (e.g. Rhapsody) gain cred from acquisition
- FC born / grew as Ethernet crushed ATM, Token Ring
  - By contrast, iSCSI adoption is modest… Hmm.
- Firms have mixed storage needs, but buy for top SLA

Summary

- Classic enterprise storage remains strong
  - Speed is good, but robustness, risk-avoidance non-negotiable
- For differentiated market traction, products should…
  - Include sync replication / failover features
    - Battery protection, etc, to write back any caches on power fail
  - Have strong redundancy (e.g. dual port NVMe modules)
  - Support non-disruptive upgrades to allow five 9's availability
  - Be fully tested / 24x7 support / interoperable with key vendors
  - Offer low risk migration path (some backward compatibility)