EMC²
where information lives

Fibre Channel

Disk Storage System Interface Speed Roadmap

EMC²
where information lives

Recommendations for future drive interfaces

- Needs to co-exist with 4G Arbitrated Loops
  - Newer generation disks are frequently installed into previous generation systems.
  - Drives may be intermixed with previous generation disks.
- Needs to enable a complete coherent infrastructure from the processor bus to the drive media interface.
  - Speed agility and protocol conversion impact all aspects of the storage system infrastructure, not just the disk interface.
- Needs to remain cost-competitive with other enterprise-class infrastructures.
  - Will be measured against value propositions of other disk interface technologies.

EMC²
where information lives

Example Storage System Infrastructure

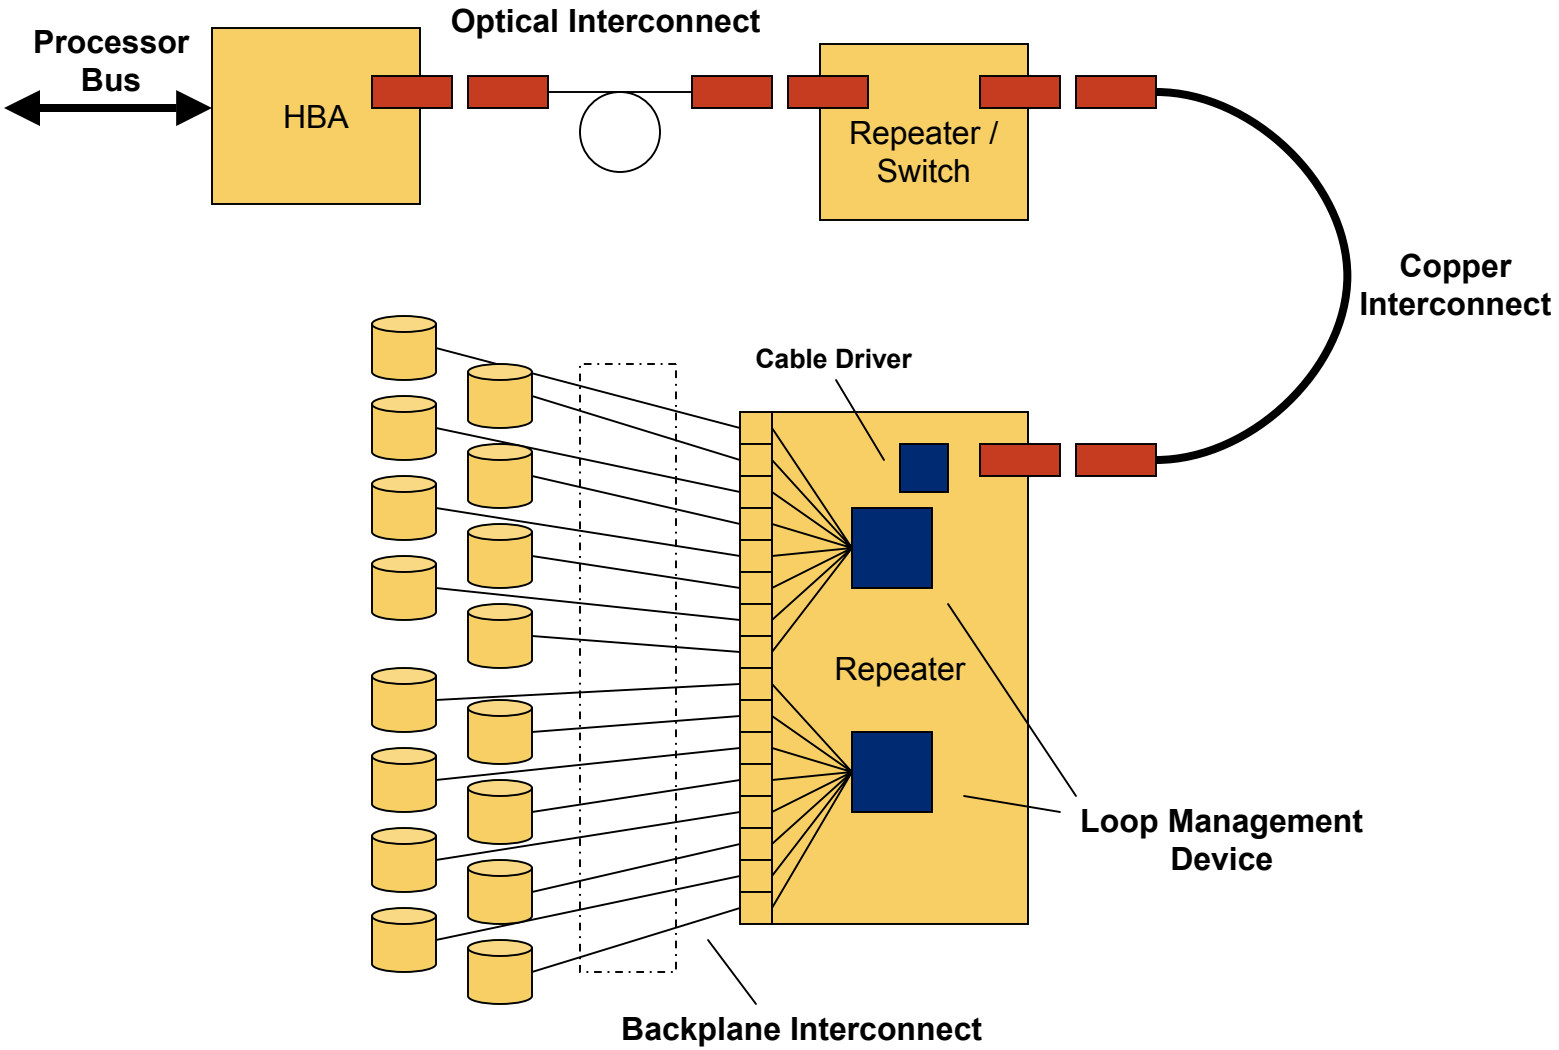

EMC²
where information lives

Proposals before us

- 4G migrates to 6G
  - Converge on SAS PHY
- 4G migrates to 8G
  - Line doubling of 4G
- 4G migrates to 10GFC
  - Use existing PHY roadmap but requires adoption of a different encoding scheme from previous AL generations.
- 4G migrates to 12G
  - Use 8B/10B encoding, data rate compatible with 10GFC.

EMC²
where information lives

Speeds & Feeds

Disk Interface Speed	Half-Duplex Data Rate	Processor Bus Requirement	Full-Duplex Data Rate	Processor Bus Requirement	Shared Media Rate† x16	x32	x64	Derived‡ 64K IOPS
1G	100MB/Sec	PCI 32/33	200MB/Sec	PCI 64/33	6.25	3.125	1.563	95 / 47 / 23
2G	200MB/Sec	PCI 64/33	400MB/Sec	PCI(X) 64/66	12.50	6.25	3.125	190 / 95 / 47
4G	400MB/Sec	PCI(X) 64/66	800MB/Sec	PCI-X 1.0-133 /	25.00	12.50	6.25	381 / 190 / 95
6G	600MB/Sec	PCI-X 1.0-133	1200MB/Sec	PCI-X 2.0 266, PCI-e 4X	37.5	18.75	9.375	571 / 285 / 99
8G	800MB/Sec	PCI -X 1.0-133	1600MB/Sec	PCI-X 2.0-266 PCI-e 4X	50.00	25.00	12.50	762 / 381 / 190
10GFC 12G	1200MB/Sec	PCI -X 2.0-266/ PCI-e 4X	2400MB/Sec	PCI-X 2.0-533, PCI-e 8X	75.00	37.5	18.75	1143 / 571 / 285

† Apportioned bandwidth available to each drive of a 16, 32, or 64-drive loop, MB/sec.

‡ Average number of 64K I/o's that apportioned bandwidth could support.

EMC²
where information lives

Speeds & Feeds

Disk Interface Speed	Half-Duplex Data Rate	Processor Bus Requirement	Full-Duplex Data Rate	Processor Bus Requirement	Shared Media Rate† x16	x32	x64	Derived‡ 64K IOPS
1G	100MB/Sec	PCI 32/33	200MB/Sec	PCI 64/33	6.25	3.125	1.563	95 / 47 / 23
2G	200MB/Sec	PCI 64/33	400MB/Sec	PCI(X) 64/66	12.50	6.25	3.125	190 / 95 / 47
4G	400MB/Sec	PCI(X) 64/66	800MB/Sec	PCI-X 1.0-133 /	25.00	12.50	6.25	381 / 190 / 95
6G	600MB/Sec	PCI-X 1.0-133	1200MB/Sec	PCI-X 2.0 266, PCI-e 4X	37.5	18.75	9.375	571 / 285 / 99
8G	800MB/Sec	PCI -X 1.0-133	1600MB/Sec	PCI-X 2.0-266 PCI-e 4X	50.00	25.00	12.50	762 / 381 / 190
10GFC 12G	1200MB/Sec	PCI-X 2.0-266/ PCI-e 4X	2400MB/Sec	PCI-X 2.0-533, PCI-e 8X	75.00	37.5	18.75	1143 / 571 / 285

System Bus requirements for 6G and 8G are the same

EMC²
where information lives

Speeds & Feeds

Disk Interface Speed	Half-Duplex Data Rate	Processor Bus Requirement	Full-Duplex Data Rate	Processor Bus Requirement	Shared Media Rate† x16	x32	x64	Derived‡ 64K IOPS
1G	100MB/Sec	PCI 32/33	200MB/Sec	PCI 64/33	6.25	3.125	1.563	95 / 47 / 23
2G	200MB/Sec	PCI 64/33	400MB/Sec	PCI(X) 64/66	12.50	6.25	3.125	190 / 95 / 47
4G	400MB/Sec	PCI(X) 64/66	800MB/Sec	PCI-X 1.0-133 /	25.00	12.50	6.25	381 / 190 / 95
6G	600MB/Sec	PCI-X 1.0-133	1200MB/Sec	PCI-X 2.0 266, PCI-e 4X	37.5	18.75	9.375	571 / 285 / 99
8G	800MB/Sec	PCI -X 1.0-133	1600MB/Sec	PCI-X 2.0-266 PCI-e 4X	50.00	25.00	12.50	762 / 381 / 190
10GFC 12G	1200MB/Sec	PCI -X 2.0-266/ PCI-e 4X	2400MB/Sec	PCI-X 2.0-533, PCI-e 8X	75.00	37.5	18.75	1143 / 571 / 285

Architecting for 10GFC eliminates processor bus b/w as a factor

EMC²
where information lives

Speed Agility - A Little History

- 1G Fibre Introduced.
  - No speed agility supported.
- 2G Fibre Introduced
  - SAN customers required Auto-Negotiation in order to adopt.
  - AL adopted backplane hardstrap method.
- 4G Fibre TBD
  - SAN will continue to use AN.
  - AL will continue to use hardstrap.
- NexGen Fibre
  - Will this be specified in the SAN? Can this work inside the scope of the present AN algorithm (i.e., 4 levels)?
  - 4G backplane compatibility dictates the adoption of hardstrap method.

EMC²
where information lives

Technology Components

Disk Interface Speed	Optical Cable Form Factor	Copper Cable Form Factor	Disk Connector
1G	GLM/GBIC/SC	DB9, HSSDC	SCA
2G	SFF/SFP/LC	HSSDC(2), SFP	SCA
4G	SFF/SFP/LC	HSSDC(2), SFP	SCA
6G	TBD/LC (?)	CX4, IB / SFP (?)	SCA
8G	XFP,X2,XPAK/LC (?)	CX4, SFP (?)	SCA
10GFC 12G	XFP,X2,XPAK/LC	CX4, XFI_Cable	TBD

EMC²
where information lives

Physical Technology

Disk Interface Speed	Line Rate, Gbps	Optical Cable Distance Objective	Copper Cable Distance Objective	Backplane Distance Objective	SERDES (CMU)
4G/ 6G	4.250GHz/ 6.375GHz	150M/ 120M?	7M / 6M?	0.5M / TBD	40X, 120X/20X
4G/ 8G	4.250GHz/ 8.500GHz	150M / 100M?	7M / 5M?	0.5M / TBD	40X, 80X
4G/ 10GFC	4.250Gbps 10.51Gbps	150M / 80M	7M / 3M?	0.5M / TBD	Dual Oscillator
4G/ 12G	4.250GHz 12.750GHz	150M / 50M?	7M / 2M?	0.5M / TBD	40X / 120X

EMC²
where information lives

Building the Infrastructure

- 4G / 6G
  - O/E & Cable FF need to be compatible.
  - Link budget will be a challenge.
  - Interesting alignment to 10GFC (2x6G = 10GFC/12G).
- 4G / 8G
  - O/E & Cable FF need to be compatible.
  - Link budget will be a challenge.
  - Should we worry about the question of what next *now*?
- 4G / 10GFC
  - Incompatible Optical and cable Form Factors.
  - Incompatible Encoding formats.
  - 10GFC doesn't carry AL.
- 4G / 12GFC
  - Incompatible Optical and cable Form Factors
  - Link budgets extremely challenging.
  - Cost issues?

EMC²
where information lives

Conclusions

- 4G → 6G
  - Requires same system b/w upgrade as 8G.
  - Not clear if there are any system-wide advantages in cost (relative to 8G).
- **4G → 8G**
  - Winner by default?
  - Least impact in terms of compatibility.
  - Best “bang for buck”?
- 4G → 10GFC
  - Not clear what the advantage is of extending SAN to disk interface is.
  - Interoperability requirement imposes heavy burden.
- 4G → 12G
  - Requires additional link budget (over 10GFC)
  - Need to assess cost vs. value proposition