

All-Flash Storage System

June 2016

Jungsoo Kim

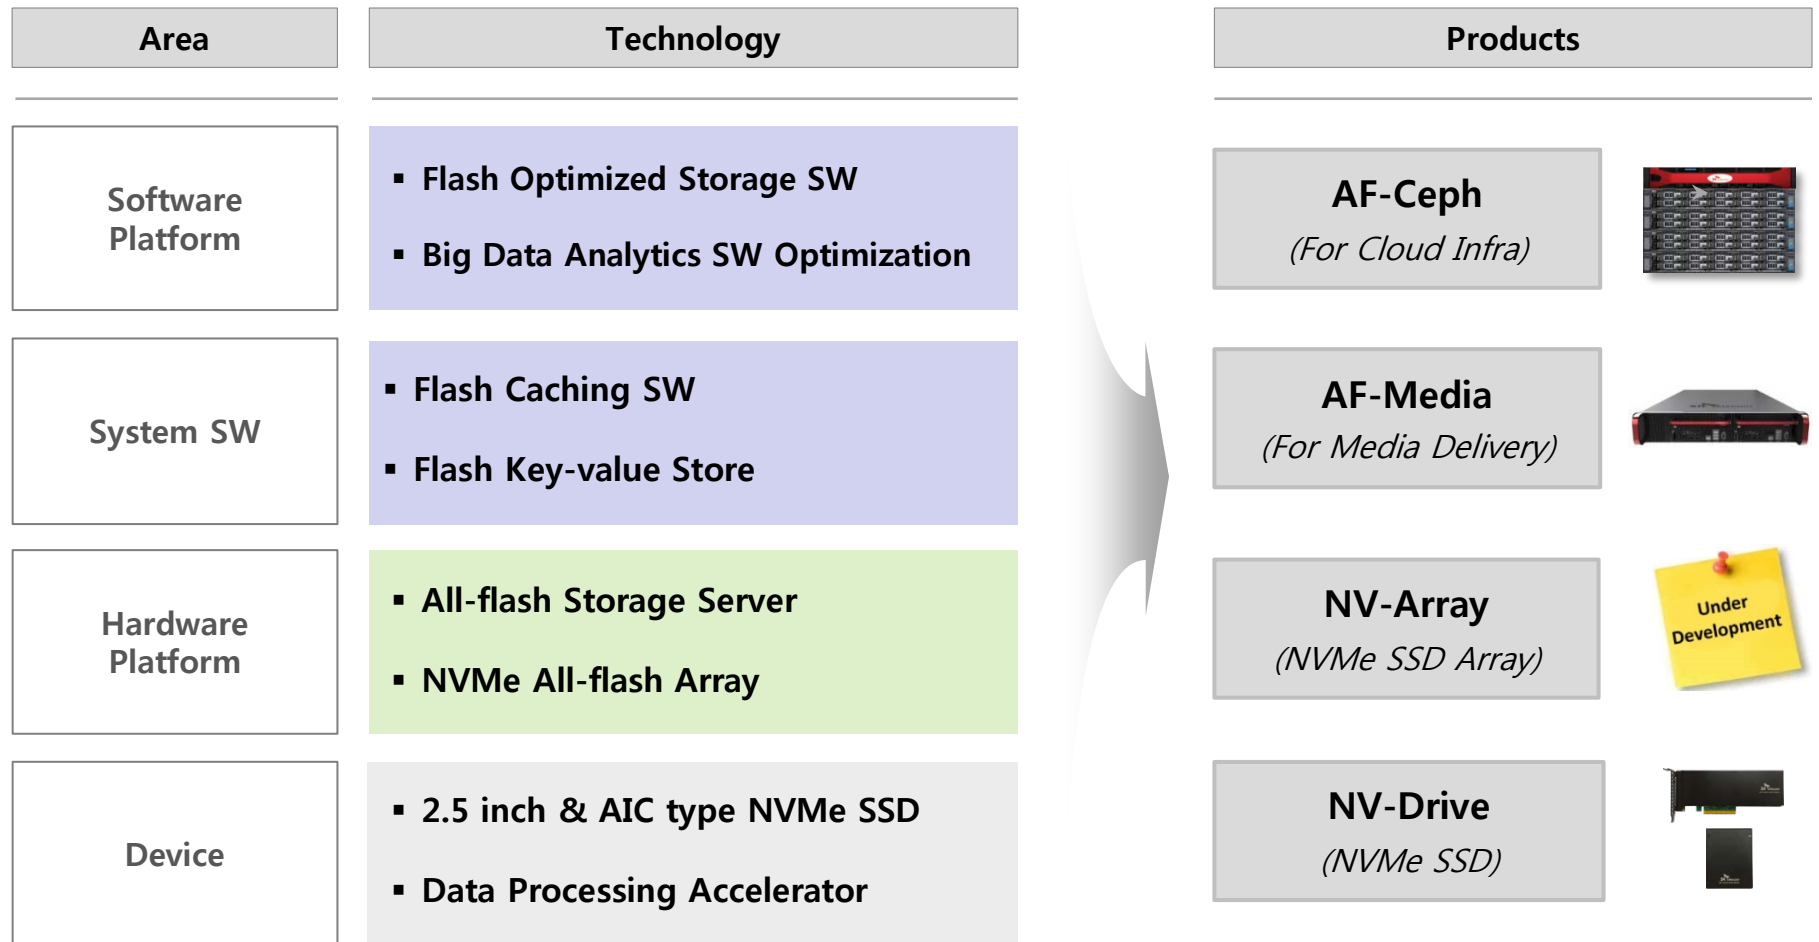
Manager, SK Telecom

Agenda

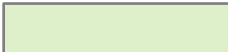
- SKT Storage Solution R&D Introduction
- Our approaches in developing storage system
- AF-Media details
 - Computing Board
 - Storage Module
 - System Reliability & Availability
 - Performance
 - Comparison
- AF-Media : Future Work
- NV-Array Development Collaboration
- Timeline

SKT Flash Storage Solution R&D Intro


From Device to Software Platform, SKT Storage R&D vertically aligned and focused on development of Flash-based technologies



 : Open Source S/W Project

 : Open Source H/W Project

In the middle of nowhere ...

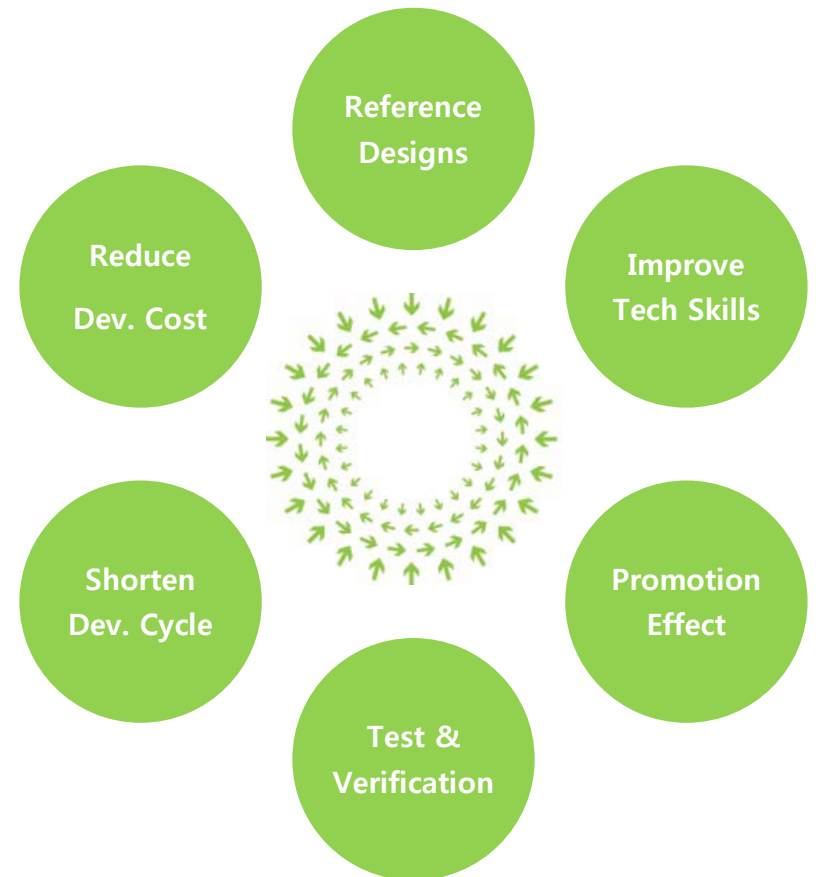


Our Approaches

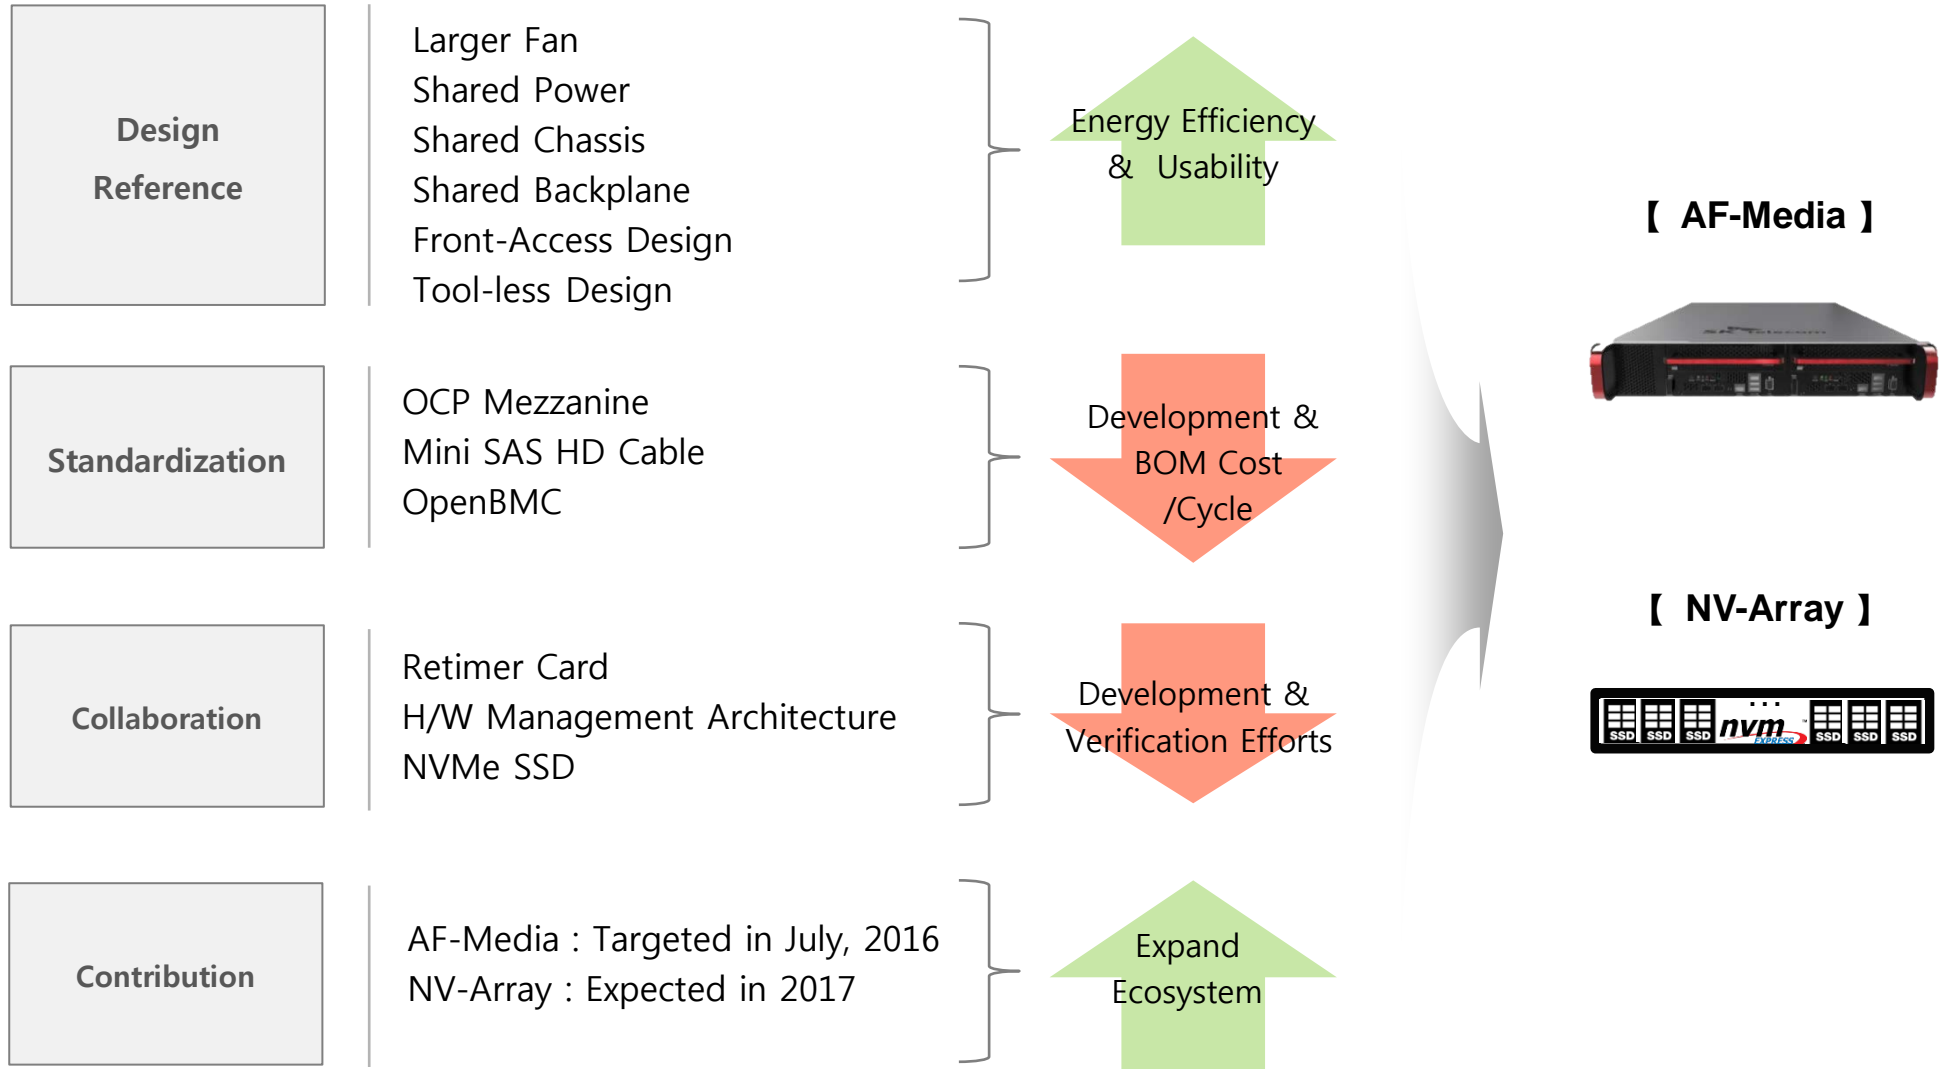
Building Ecosystem



Open Hardware Community



How we utilize Open Hardware ...



AF-Media : All-Flash Storage Server



High Density

SATA SSD 16ea in 1U size
High Capacity



High Bandwidth

High IO Performance
40GbE Network Support



Availability

Node Hot-Plug &
SSD Hot-Swap Support



Low Power

Under 150W/Node

【 Application 】

Media
Streaming

High Perf.
File Server

CDN
Edge Server

Proxy
Cache Server

【 AF-Media S16 】

2U



Node 1

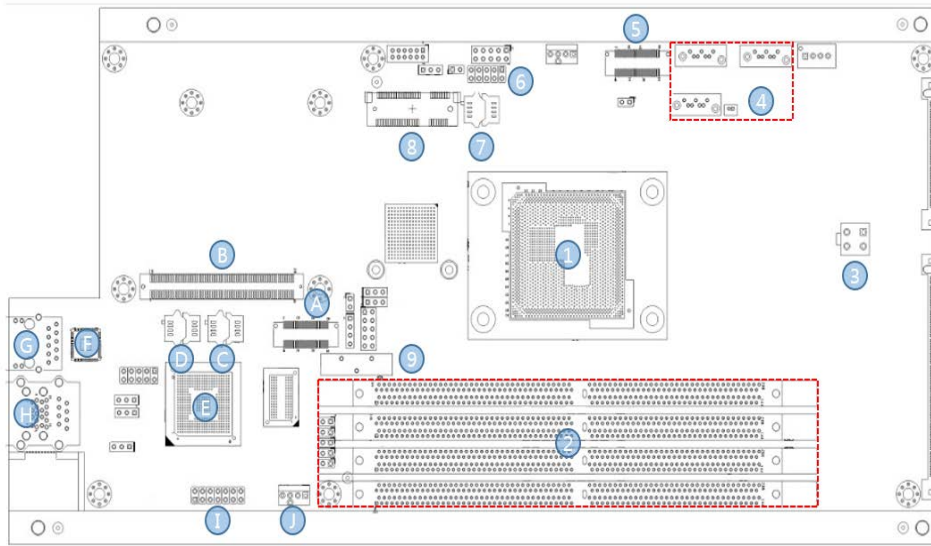
Node 2

For Each Server,

- Network Throughput : ~ 40Gbps
- Storage Capacity : 32TB/Node (w/ 2TB SSD)

AF-Media : Computing Board


【 Specification 】



No.	Item	Description	No.	Item	Description
1	CPU	D1540	B	NIC Slot	OCP Mezz. V2
2	DIMM Slot	DDR4 RDIMM	C	BMC FW	BMC Flash
3	JPWR	12V Conn.	D	vBIOS FW	Video BIOS Flash
4	SATA	SATA3 Conn.	E	BMC	AST2400
5	XDP (CPU)	Debug Port	F	Ethernet Controller	Intel I210
6	JSPI	Front Panel	G	MGMT	RJ45
7	BIOS FW	BIOS Flash	H	Console/USB	RJ45, USB2.0
8	mSATA	mSATA Slot	I	VGA	VGA HDR
9	BAT	RTC Battery	J	FAN	FAN HDR
A	XDP (PCH)	Debug Port	-		

Items	Description
Processor	Intel Xeon D-1548 <ul style="list-style-type: none"> • 8 Cores/16 Threads • 2.0GHz/2.6GHz • TDP < 45W
Memory	<ul style="list-style-type: none"> • RDIMM DDR4-2133, 2Ch 4 Slot • up to 128GB
Storage	<ul style="list-style-type: none"> • OS: mSATA 64GB SATA3
Network Card Interface	<ul style="list-style-type: none"> • OCP Mezzanine V2.0 • PCIe 3.0 x8
Interface for Storage	<ul style="list-style-type: none"> • PCIe 3.0 x16
System Management	<ul style="list-style-type: none"> • AST2400


AF-Media : Storage Module



【 Storage Sled 】



【 Drive Attention LED 】



【 Cable Chain 】



AF-Media : System Reliability & Availability

[Node Hot-Plug]



[SSD Hot-Swap]




[Redundant PSU]



AF-Media : Performance

[Video Streaming Throughput] (HLS: Http Live Streaming Protocol)





1 Mbps (SD)		2Mbps (HD)		7Mbps (Full-HD)		35Mbps (UHD)	
User	Throughput	User	Throughput	User	Throughput	User	Throughput
26,800	33.21 Gbps	16,000	36.23 Gbps	4,900	37.37 Gbps	925	37.94 Gbps

AF-Media : Comparison

【 AS-IS 】

【 AF-Media 】

Form Factor			
		S16	P2
CPU	Intel E5 2-Socket	Intel Xeon-D	Intel Xeon-D
Max. Throughput (FTP / RTSP)	~12Gbps / ~10Gbps	~38Gbps / ~25 Gbps	~38Gbps / ~25 Gbps
Storage	450GB SAS HDD * 48EA (for 15TB)	1TB SATA SSD * 16ea	NVMe PCIe Card 3.2TB * 2ea
Power Consumption	880W	150W (↓ 82%)	130W (↓ 85%)

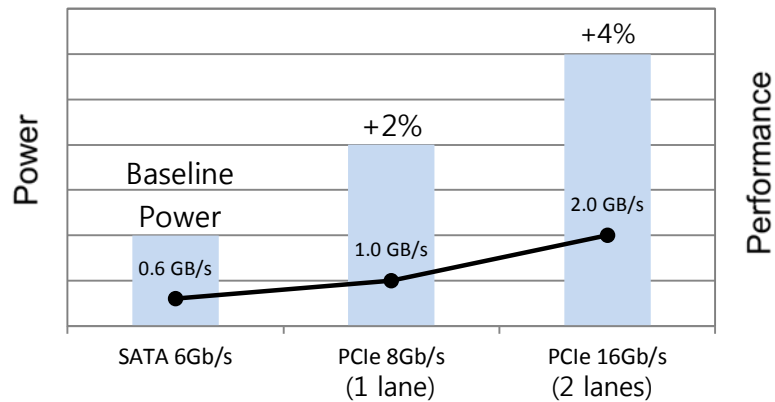
AF-Media : Future Work

1. **BOM Reduction : Inexpensive front panel, lower capacity PSU, etc.**
2. **OS storage : mSATA -> M.2**
3. **New SoC : The latest Intel Xeon D SoC (16 Core ?)**
4. **More capacity : 32TB/node with 2TB SSD**
5. **Complete Tool-less design**
6. **OpenBMC support**
7. **Design contribution to the open source community**


NV-Array : NVMe SSD Array

【 All-Flash & All-PCIe 】

【 NVMe SSD Array 】



- Two lanes of PCIe 3.0 offers 3.3x the performance of SATA 6Gb/s with only 4% increase in power
- NVMe SSD Array with PCIe host connection maximizes throughput and latency

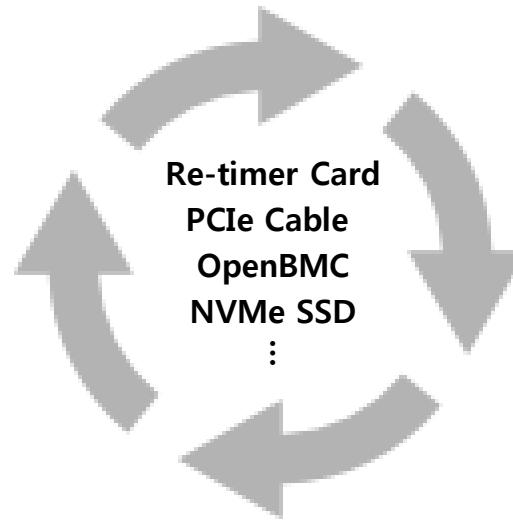


- NVMe SSD moves IO bottleneck from storage to other components in the system
- The ratio between NVMe SSD and CPU has become very important for efficiency

NV-Array : Development Collaboration

【 NV-Array 】

Form Factor	19 inch 1RU
Drive Type	2.5'' NVMe SSD
Number of SSD	20ea
Raw Capacity	80TB (w/ 4TB NVMe SSD)
Density	80TB/1RU
SSD Connector	U.2 (SFF-8639)
Disk Hot-Plug	Support
External Cable	Mini SAS HD (SFF-8644)
Number of Hosts	Up to 4 Hosts
Host Connection	PCIe 3.0 x8
Management	OpenBMC

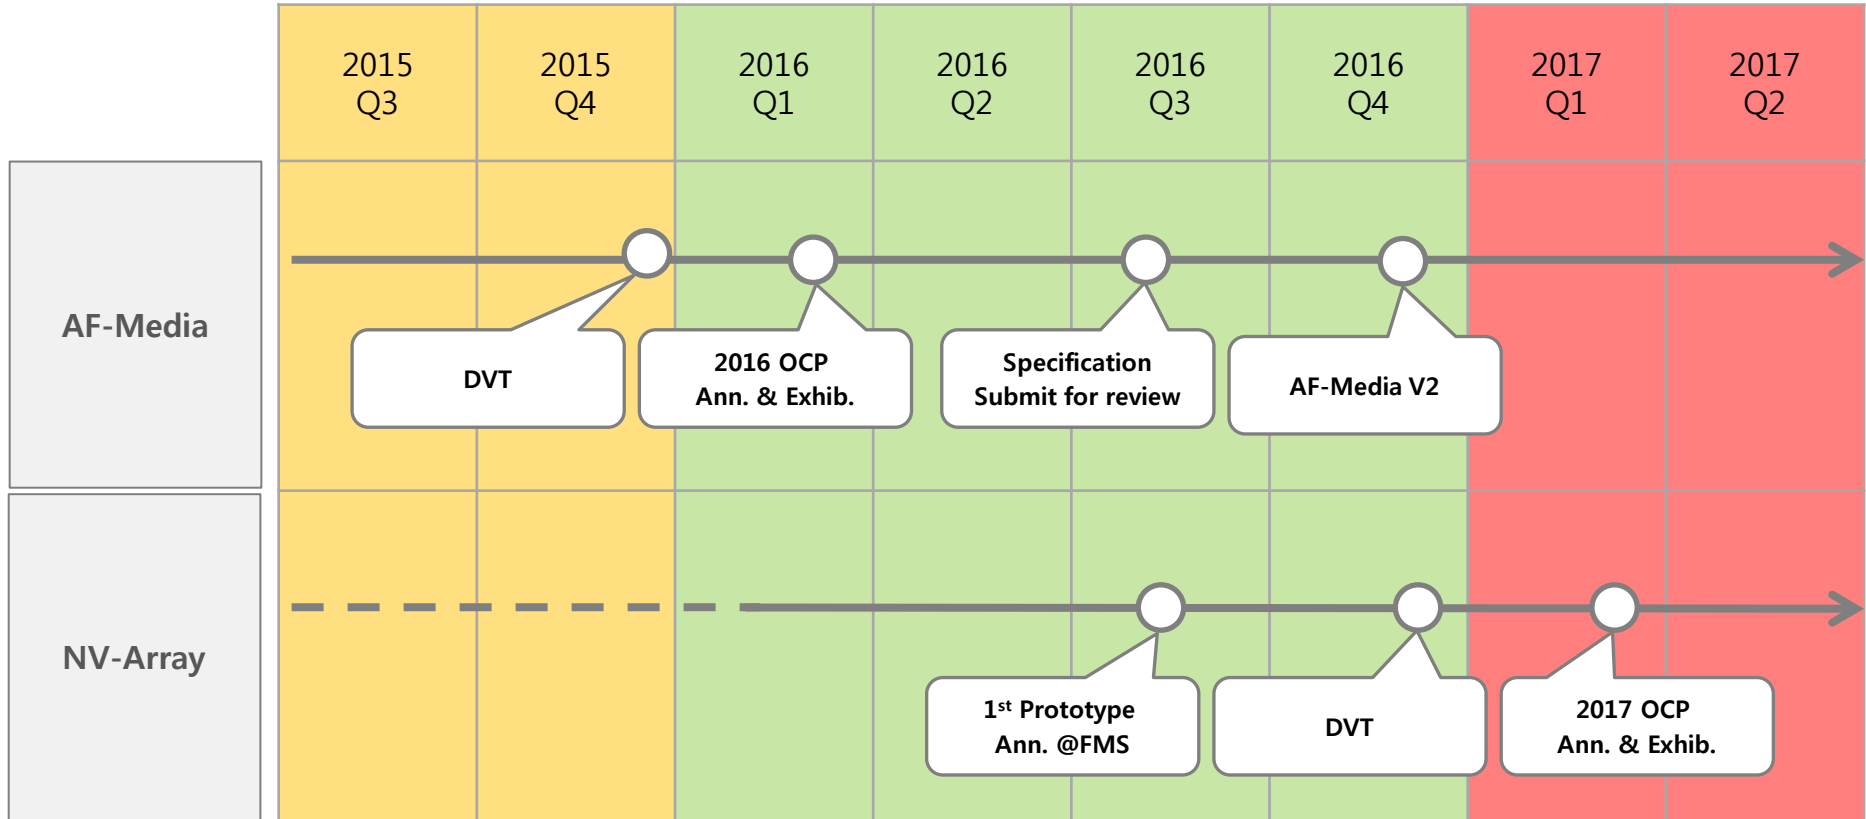


【 Lightning 】

Form Factor	21 inch 20U
Drive Type	2.5'' & M.2 NVMe SSD
Number of SSD	30ea
Raw Capacity	120TB (w/ 4TB NVMe SSD)
Density	60TB/10U
SSD Connector	U.2 & M.2 (SFF-8639)
Disk Hot-Plug	Support
External Cable	Mini SAS HD (SFF-8644)
Number of Hosts	Up to 4 Hosts
Host Connection	PCIe 3.0 x16, x8, x4
Management	OpenBMC

*This is preliminary specification and subject to change

Timeline



Thank you!