

**OCP U.S. SUMMIT 2016** March 9-10 | San Jose, CA

OCP U.S. SUMMIT 2016

# SCube AF-Media



High Density/High Efficient All-Flash Storage Server

Jungsoo Kim  
Project Manager

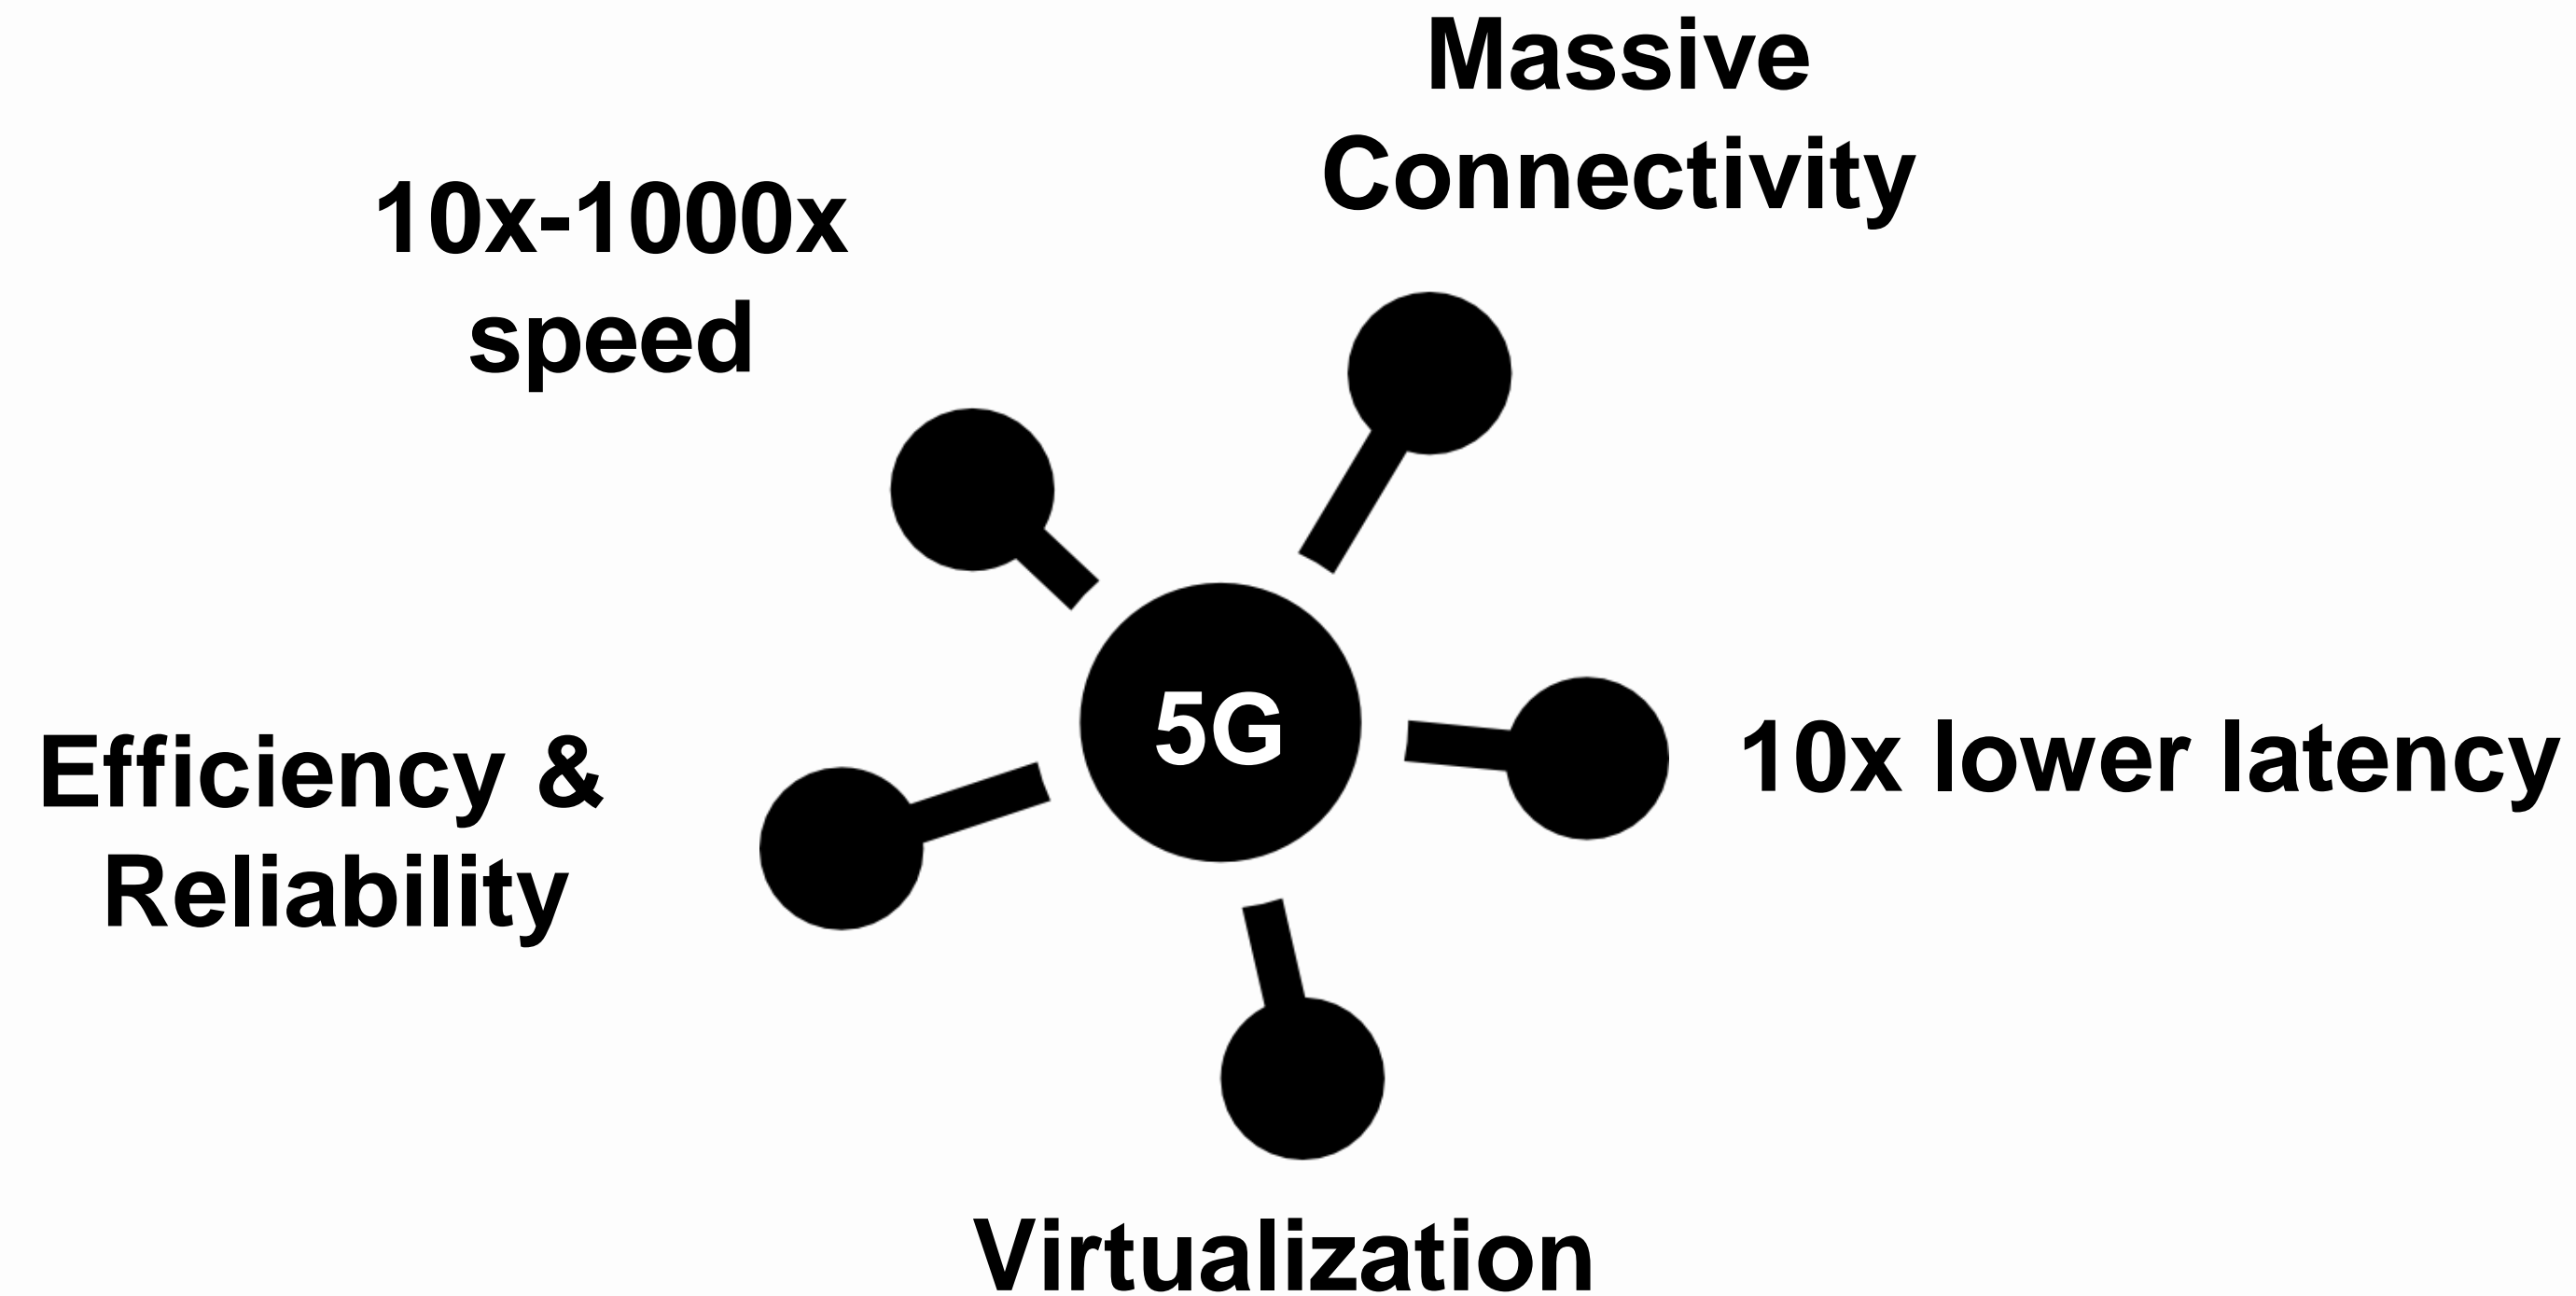
# Agenda

- Why we care about All-Flash Storage ...
- Transforming to 5G Network
- Open HW & SW Projects @ SKT
- Our approaches in developing storage system
- AF-Media details
- Future Work for AF-Media
- Future Plan : NV-Array

# Why we care about All-Flash Storage ...



# Transforming to 5G Network



New infrastructure should be Programmable, Scalable, Flexible, and Cost Effective

Software Defined Technologies based on Open Software & Open Hardware

# Open HW & SW Projects @ SKT

## Open Software



*OpenStack, ONOS, Ceph, Cloud Foundry, Hadoop ...*

## Software-Defined Technologies

## Open Hardware


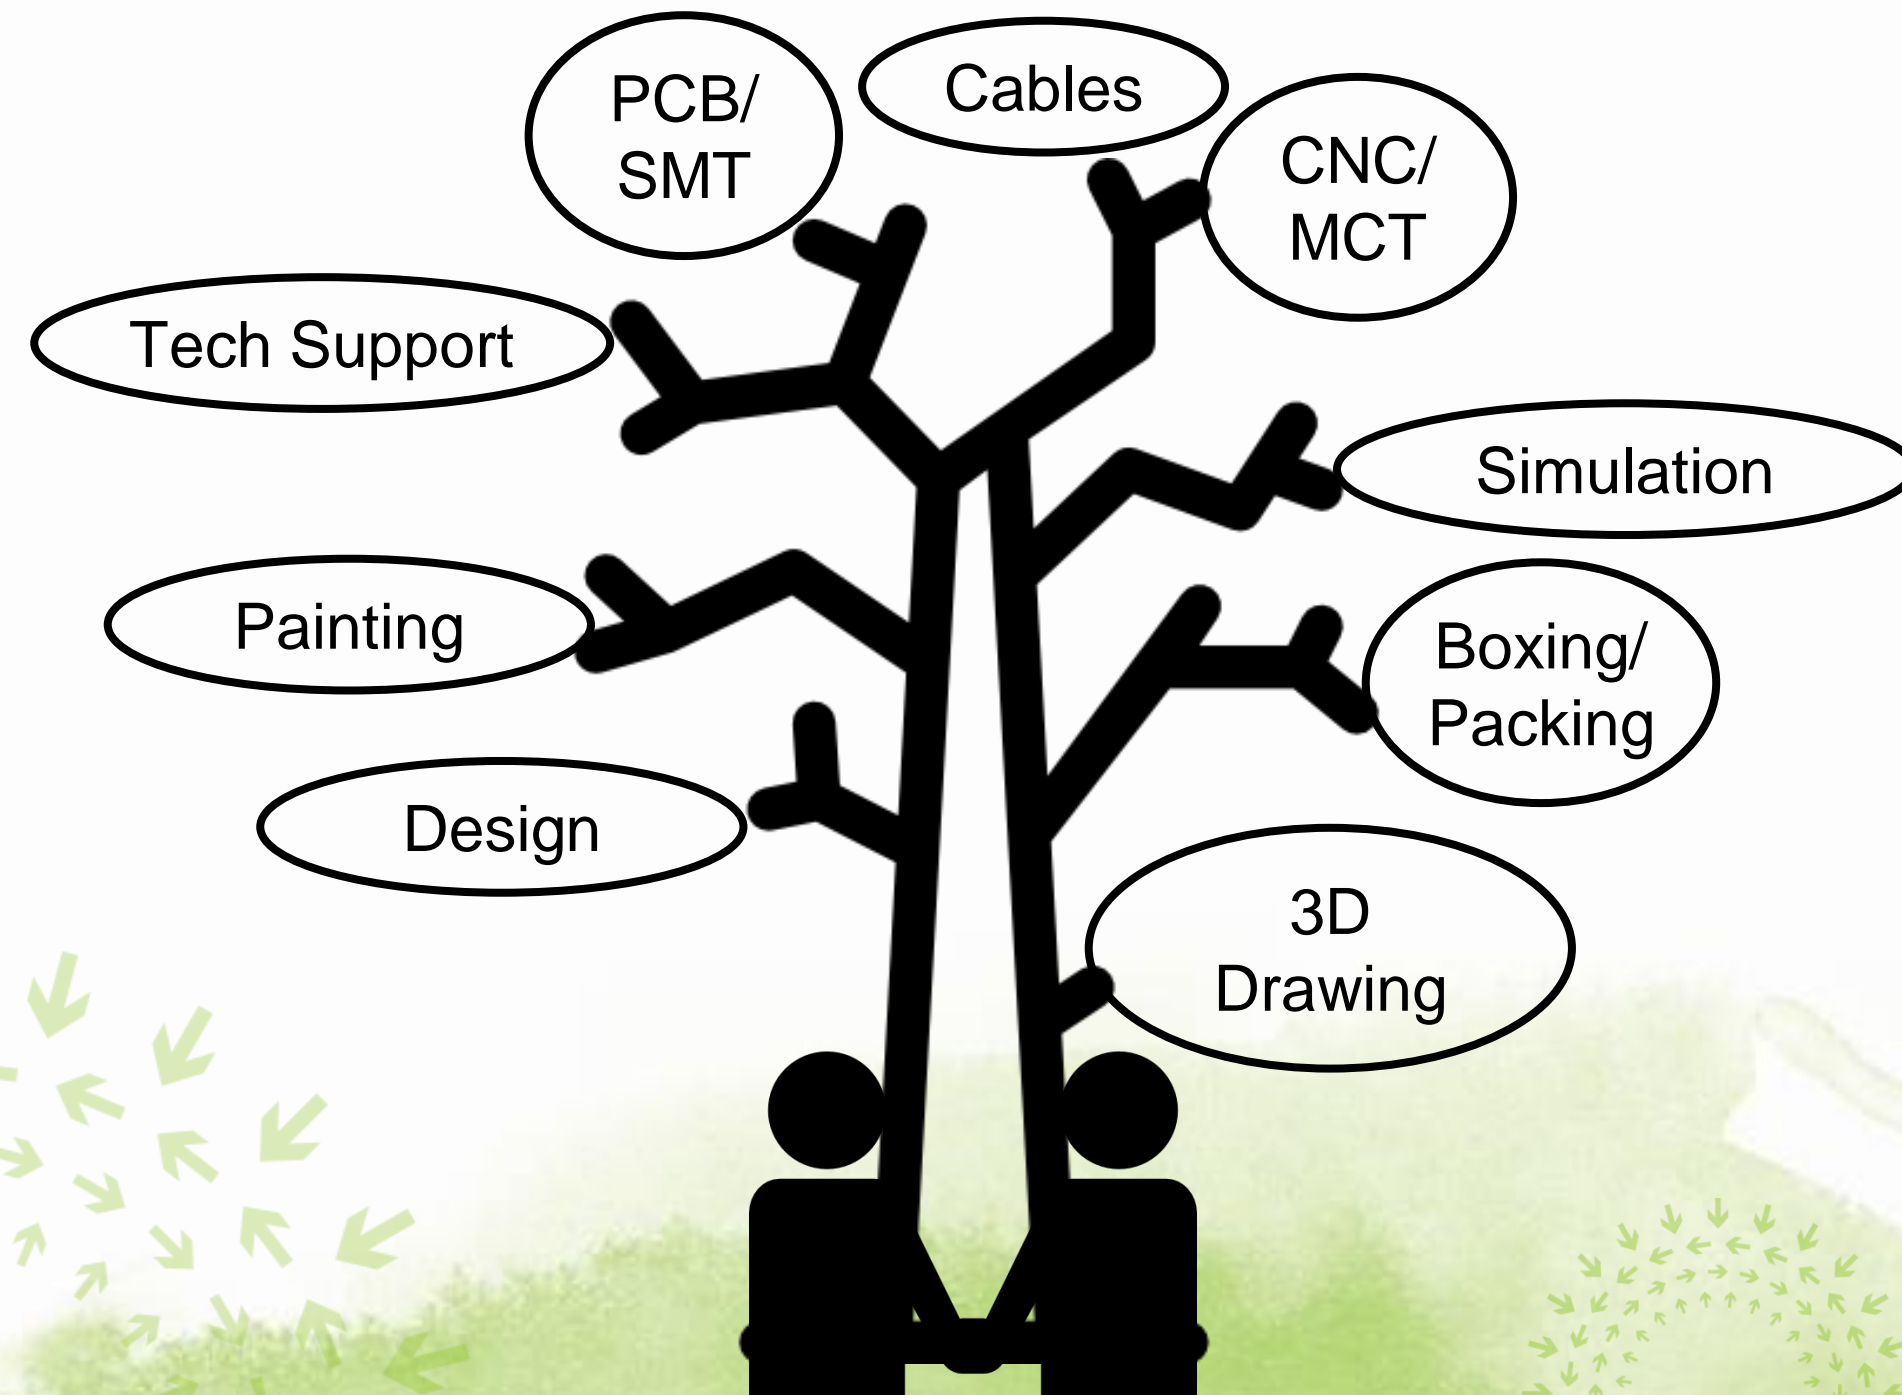
*Open Compute Project (OCP), Telecom Infra Project (TIP)  
All-Flash Storage, Server Switch, Telco Specific H/W...*

In the middle of nowhere ...



# Our Approaches

## Building Ecosystem



- Reference Designs
- Reduce Development Cost
- Shorten Development Cycle
- Test & Verification
- Improve Technical Skills
- Increase Promotion Effects




# 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

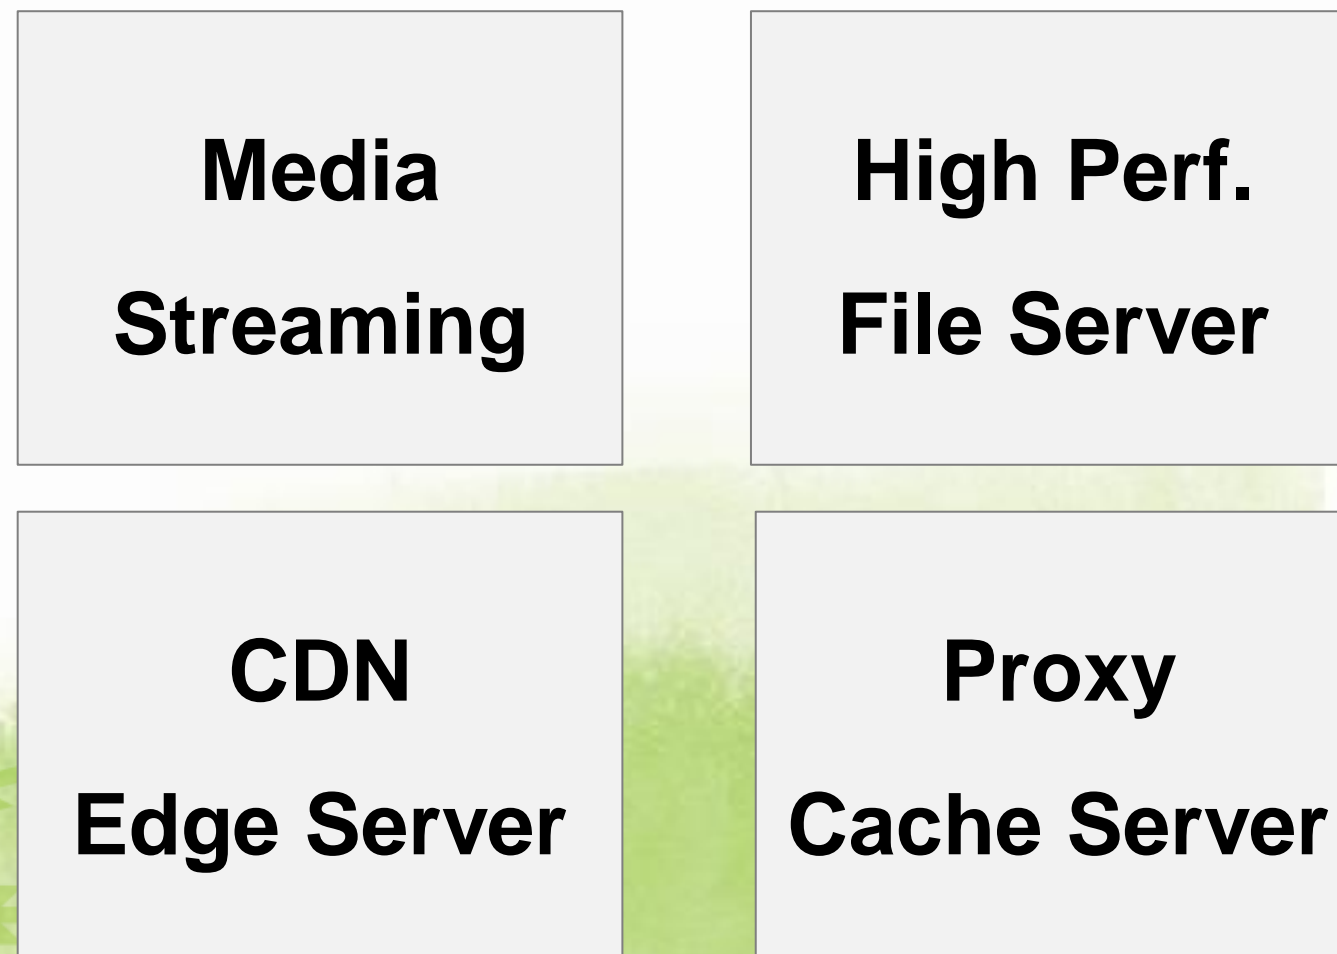
Node Hot-Plug &  
SSD Hot-Swap Support




## Low Power

Under 150W/Node

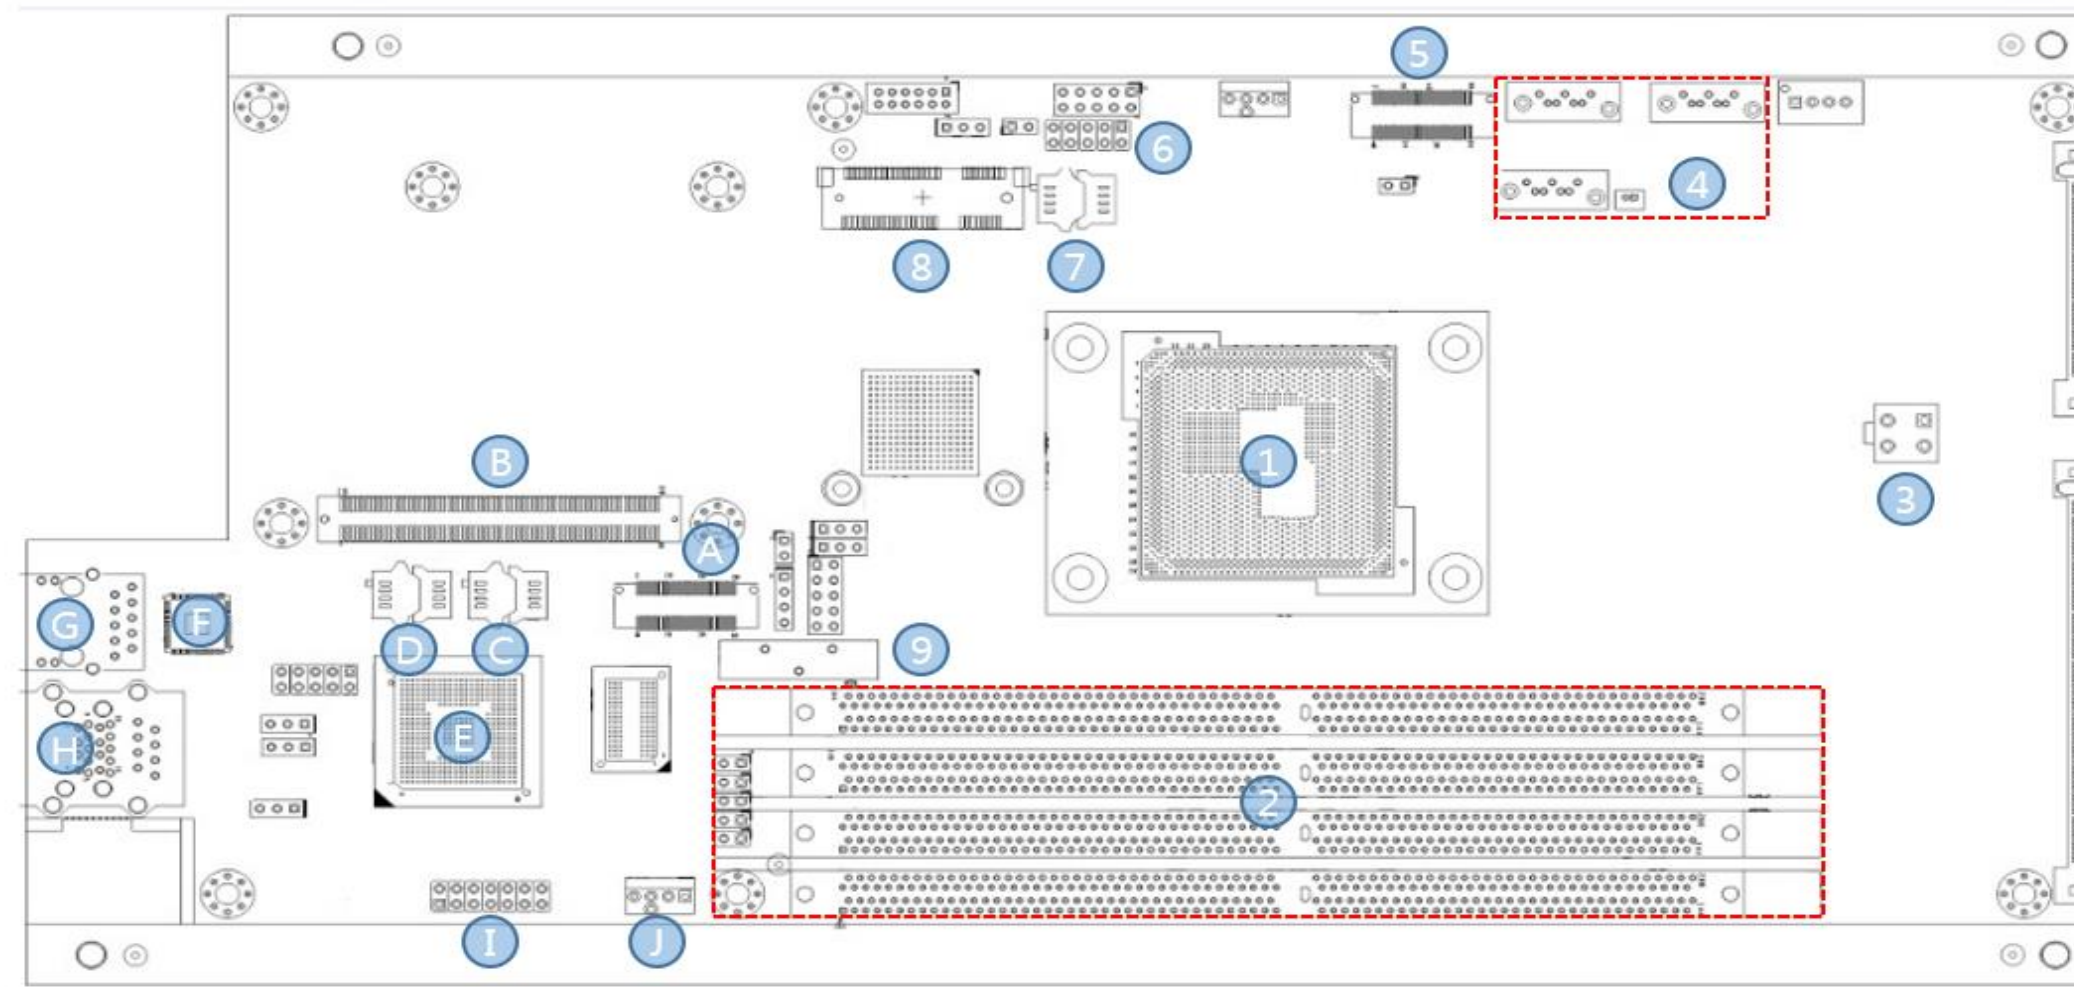
## 【 Application 】



## 【 AF-Media S16 】



# AF-Media : Computing Board



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	-		

## 【 Specification 】

### Intel Xeon D-1540

- 8 Cores/16 Threads
- 2.0GHz/2.6GHz
- TDP < 45W

### DDR4 RDIMM ECC

- 2 Channel 4 DIMMs
- Up to 128GB
- 2133MHz

### mSATA SSD

- SATA3
- for OS

### Network Card Interface

- OCP Mezzanine V2.0
- PCIe 3.0 x8

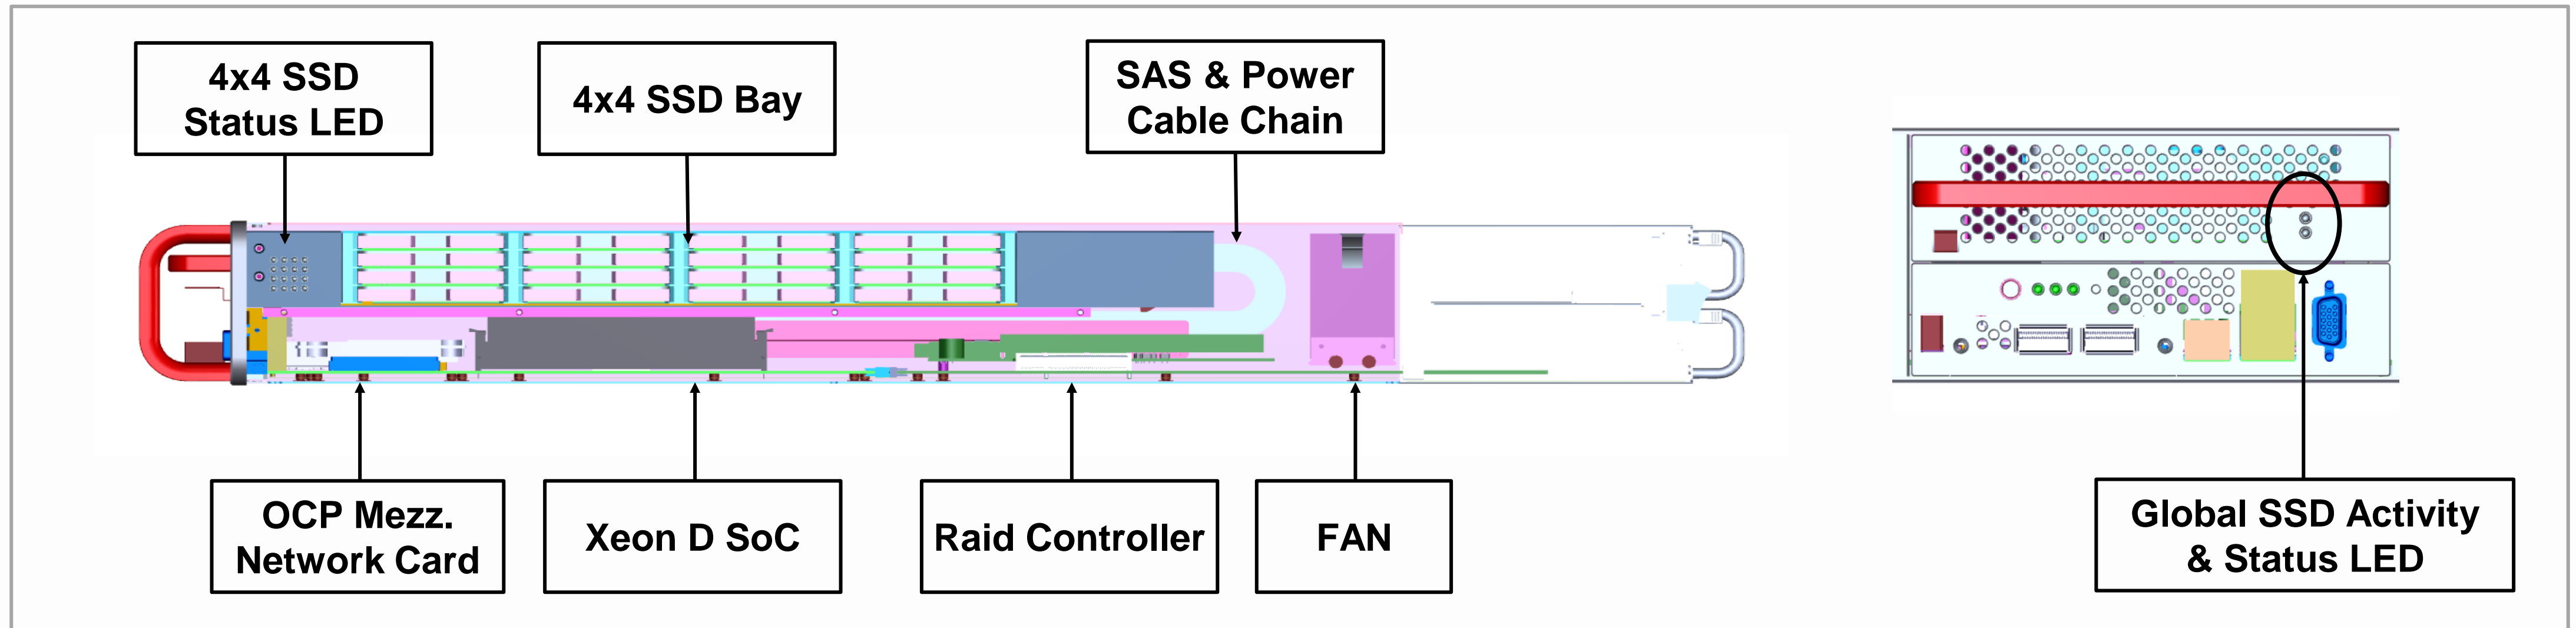
### Storage Side Connection

- PCIe 3.0 x16


### System Management

- AST2400


# AF-Media : Storage Module




**【 Storage Sled 】**



**【 SSD Status 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)





1Mbps (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%)


# Future Work for AF-Media

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




# Future Plan : NV-Array (NVMe SSD Array)

[ All-Flash & All-PCIe]



- 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 Array]



- 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




# Future Plan : NV-Array

## [ Target Specification ]

Form Factor	19 inch 1RU
SSD Type	2.5" NVMe PCIe SSD
Number of SSD	Max 24ea
Raw Capacity	96TB (w/ 4TB NVMe SSD)
SSD Interface	SFF-8639 (U.2)
External Port Type	SFF-8644 (Mini SAS HD)
Number of Hosts	Max 6
Total Bandwidth	36GB/s (48 lanes)

\*This is preliminary specification and subject to change



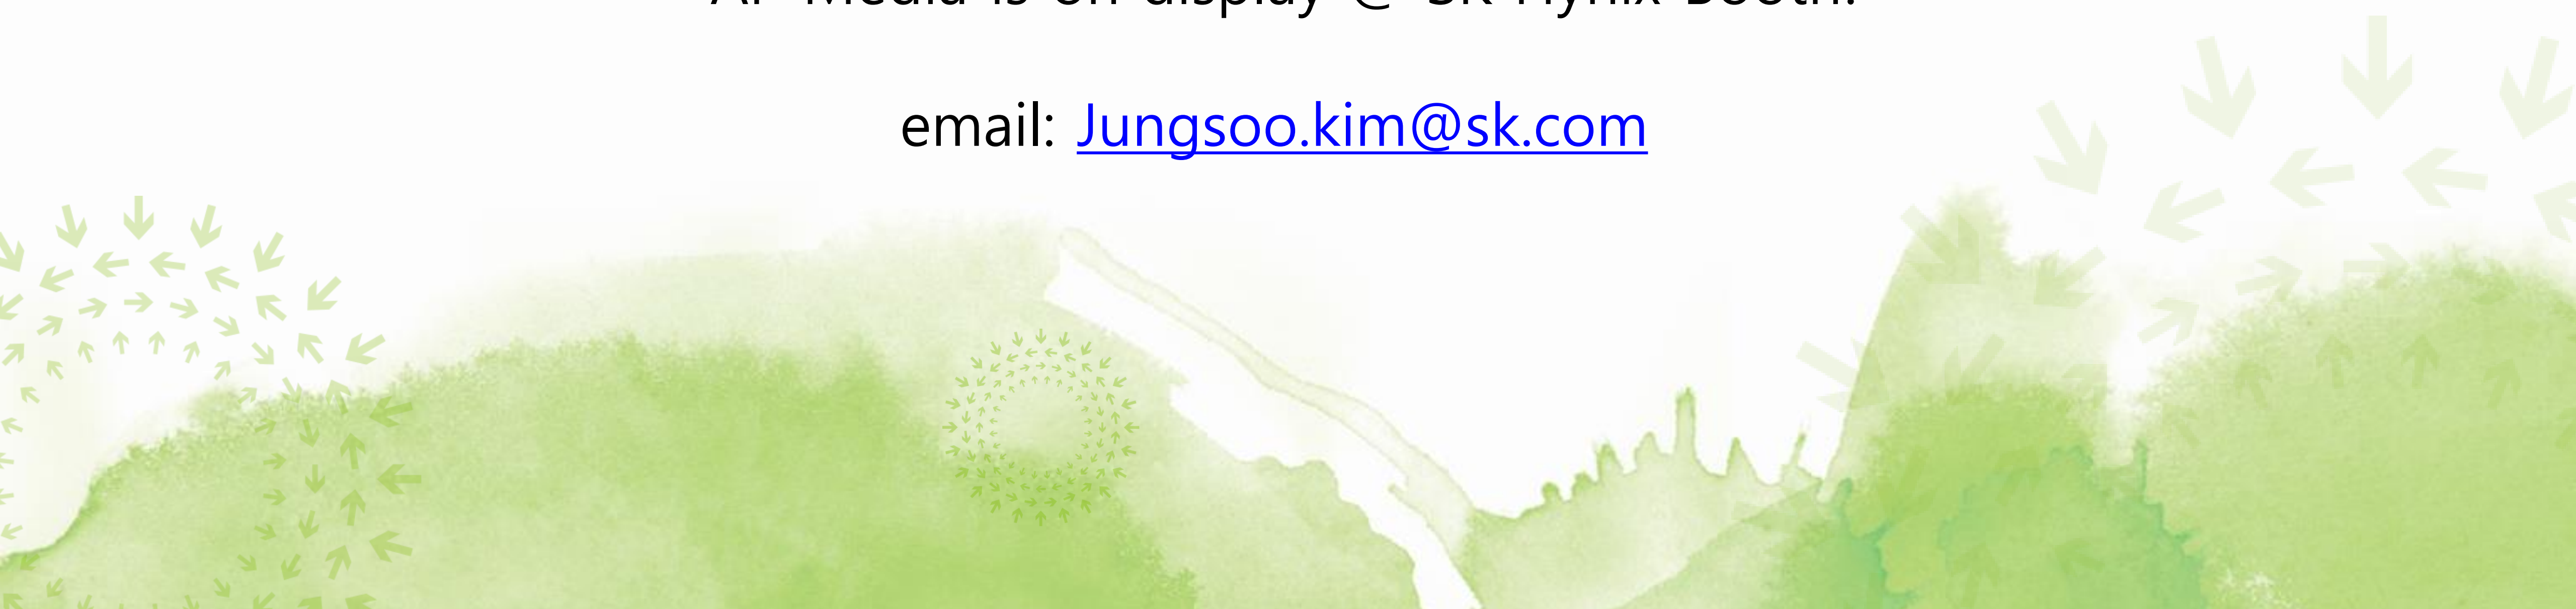
Collaboration with  
'Lightning'


Contribution planned for Open Compute Project

# Thank you

AF-Media is on display @ SK Hynix Booth!

email: [Jungsoo.kim@sk.com](mailto:Jungsoo.kim@sk.com)





# OPEN

Compute Project

