



OPEN
Compute Project

OCP U.S. SUMMIT 2017

Santa Clara, CA



Bryce Canyon

Facebook's Next Generation Storage Server

Jason Adrian / Hardware Systems Engineer / Facebook

Dominic Cheng / Hardware Engineer / Facebook

Madhavan Ravi / Hardware Engineer / Facebook

OPEN HARDWARE.



OPEN SOFTWARE.



OPEN FUTURE.



What is Bryce Canyon?





Why did we build Bryce Canyon?


- Higher CPU power per slot
- Thermal efficiency
- Power efficiency
- Flexible platform for future growth
- Hot swap drives

Component Details


Thermal	4x dual rotor 92 mm fans
Power	ORv2 bus bar clip
Storage Protocol	2x Storage Controller Card
Serviceability	Cable track for hot swap
Compute	2x 1P Server Slot (Mono Lake)
Slide Rails	
Network + extensibility	IO Module



Dual Storage Server




- Independent power distribution to each storage server and its HDDs
- Top View: Compute card, Storage controller card, HDDs
- Front View: IO Module



- Server A components
- Server B components

Drive Plane Board and Mono Lake

Top View



Rear Drive
Plane Board

Air Baffle
/Stiffener
On the bottom


Front Drive
Plane Board

Top View

Up to 4x 32GB
DDR4 DIMMs


CPU

M.2 SATA
boot drive



Storage Controller Card variants


Top View




12 Gb SAS
Expander

SAS IO
Controller

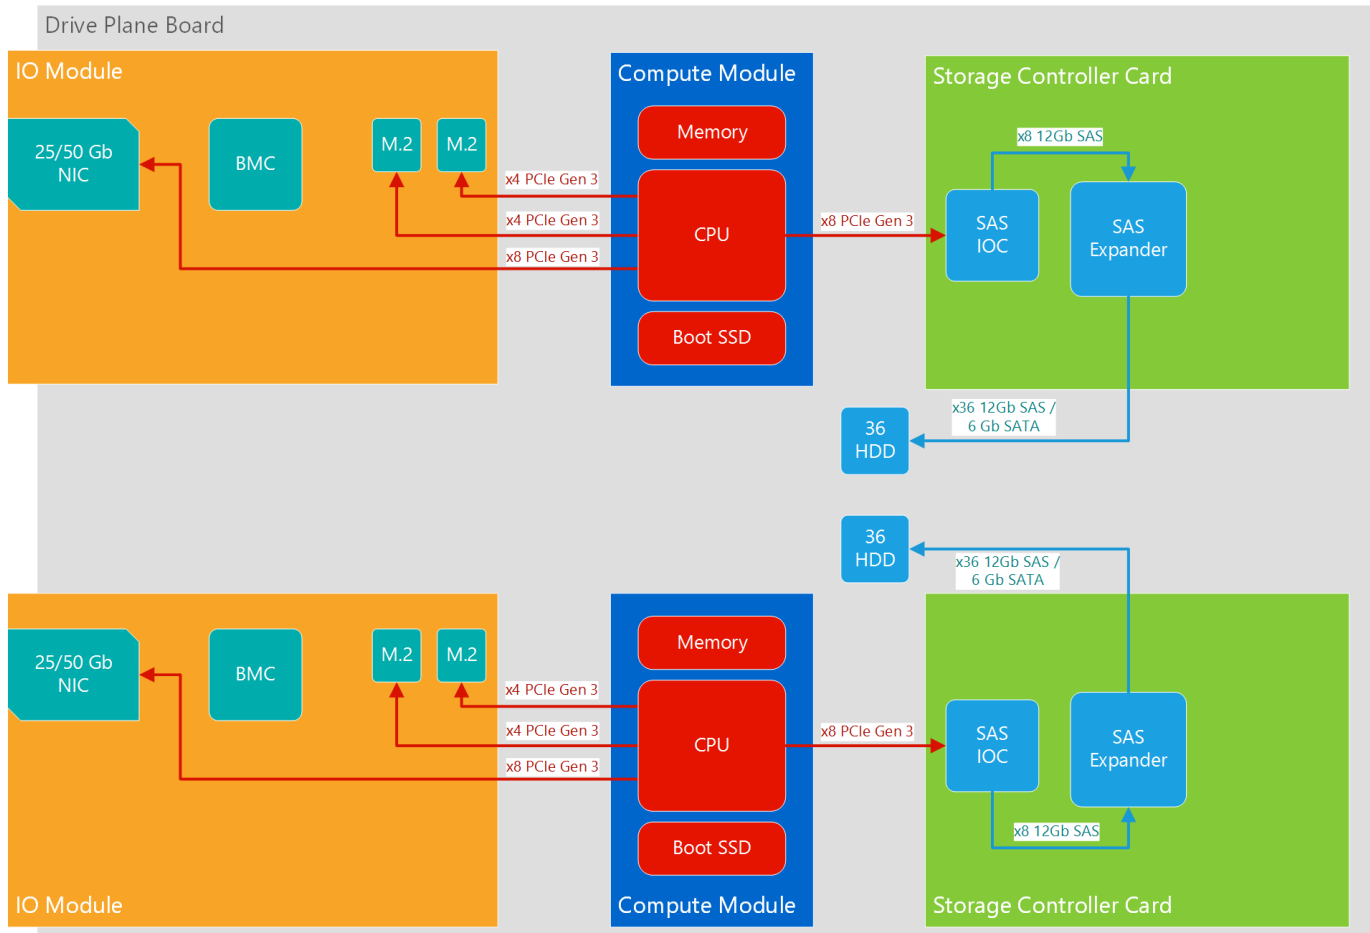
Top View




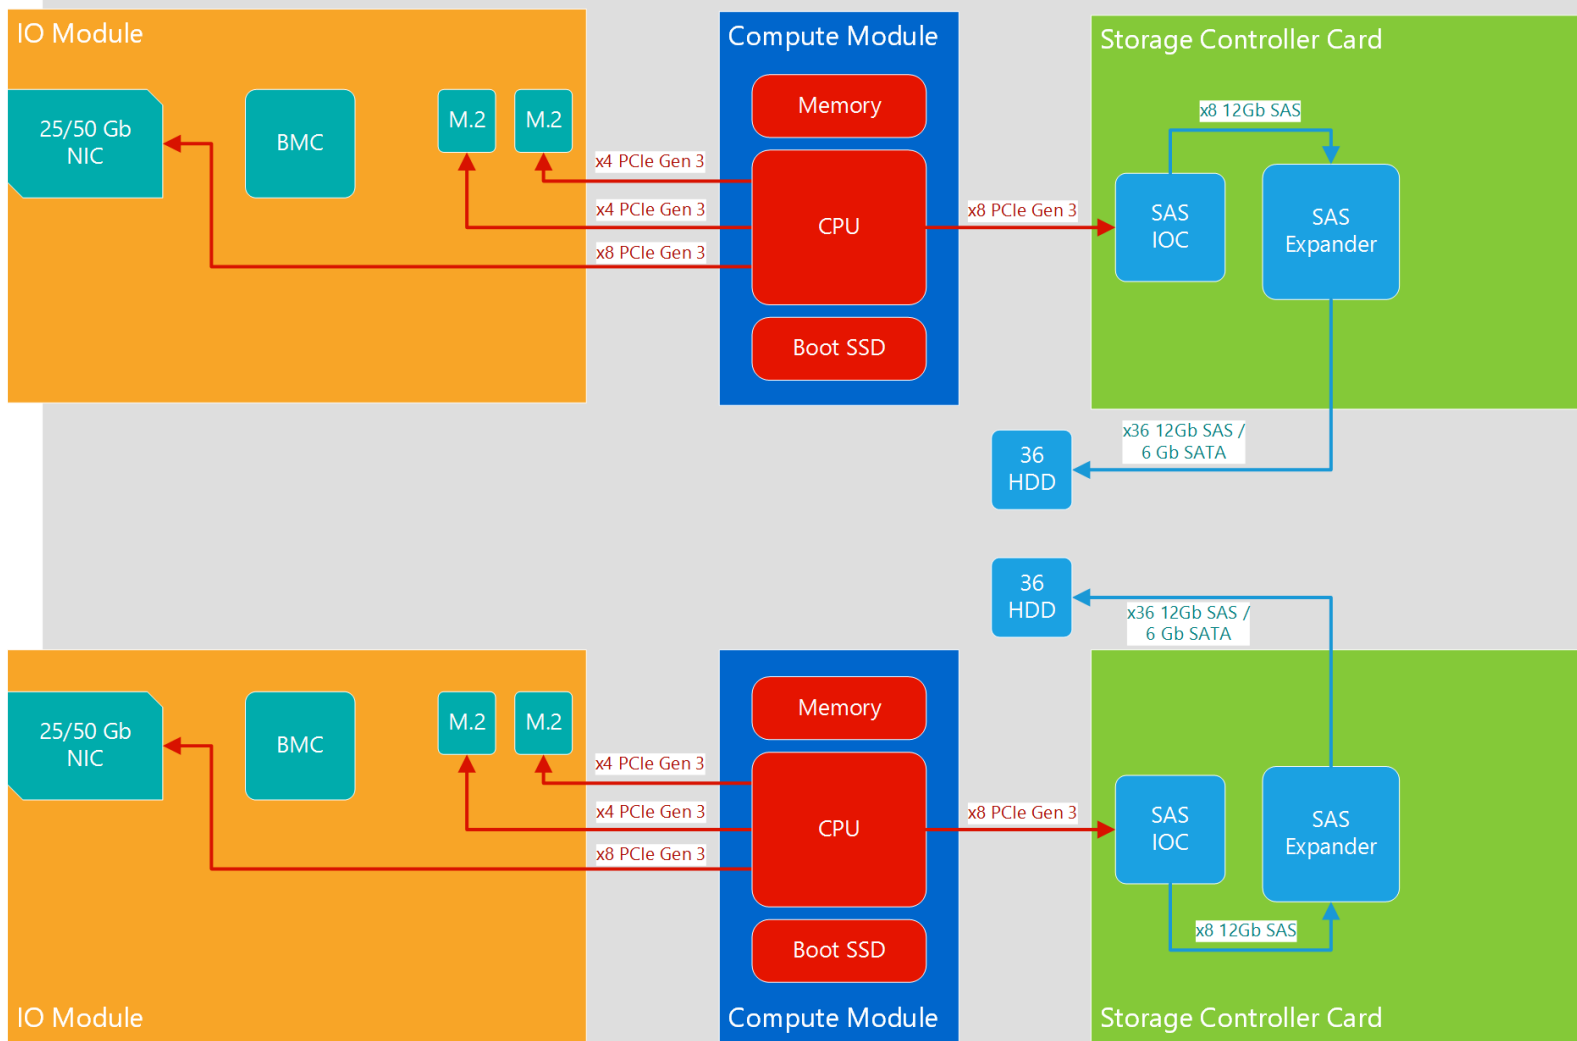
IO Module (M.2 and IOC variants)




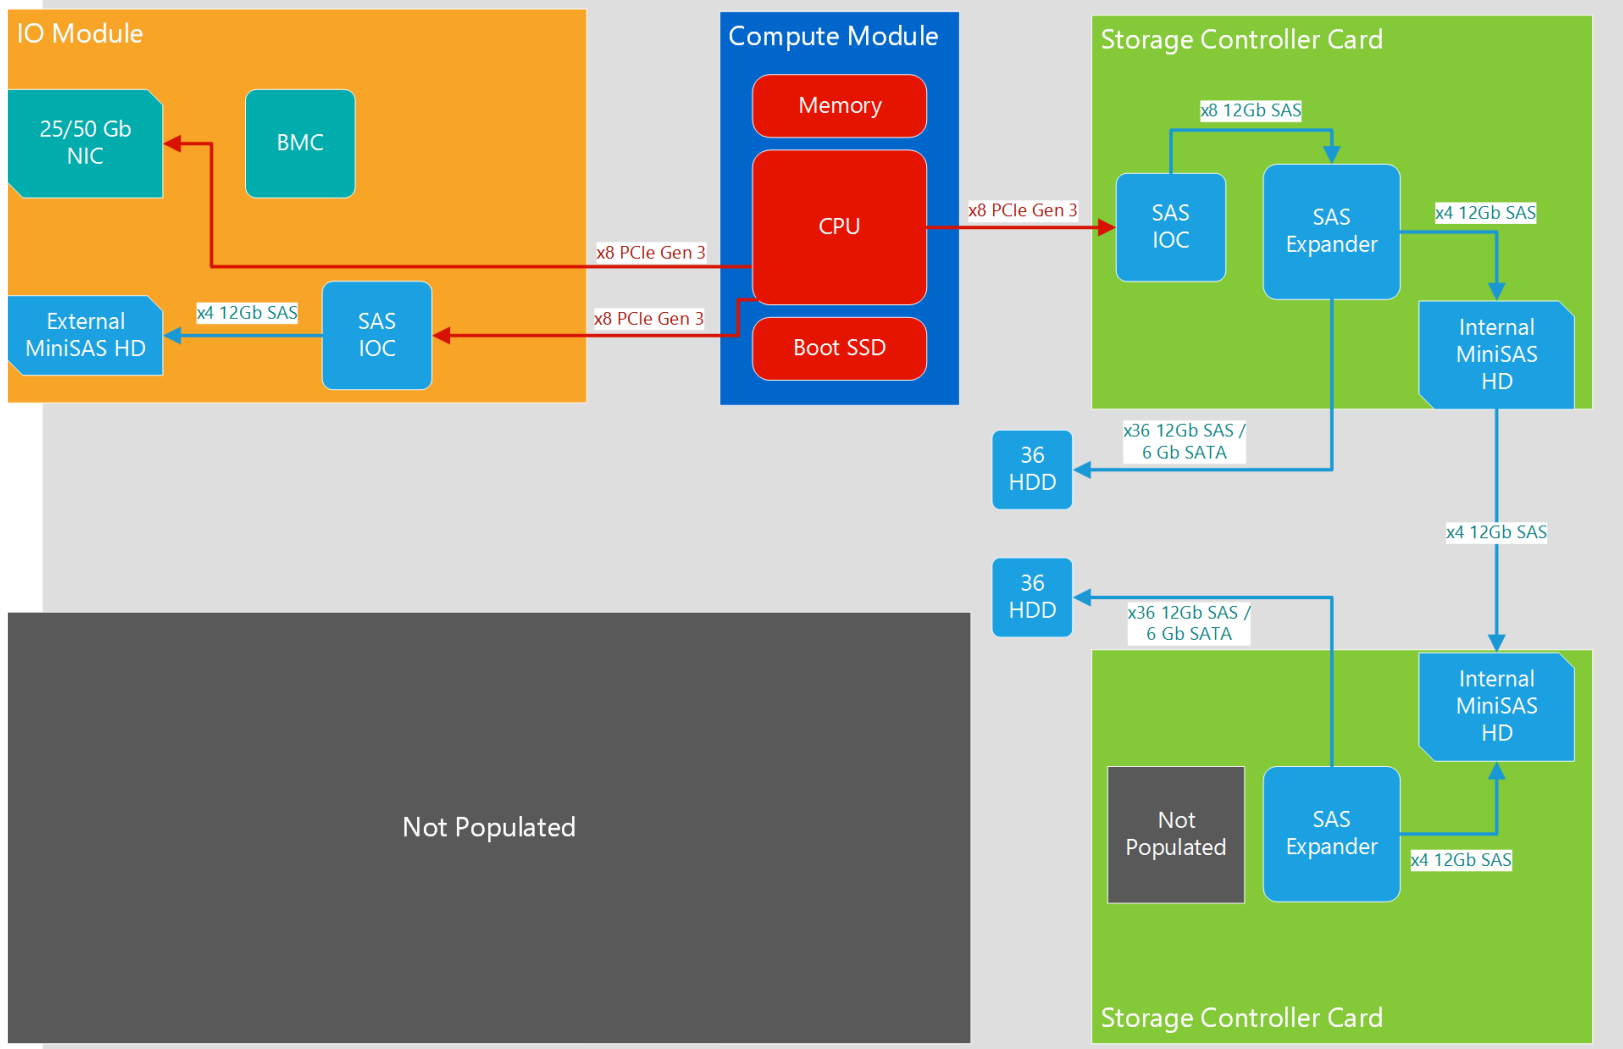
Block diagram


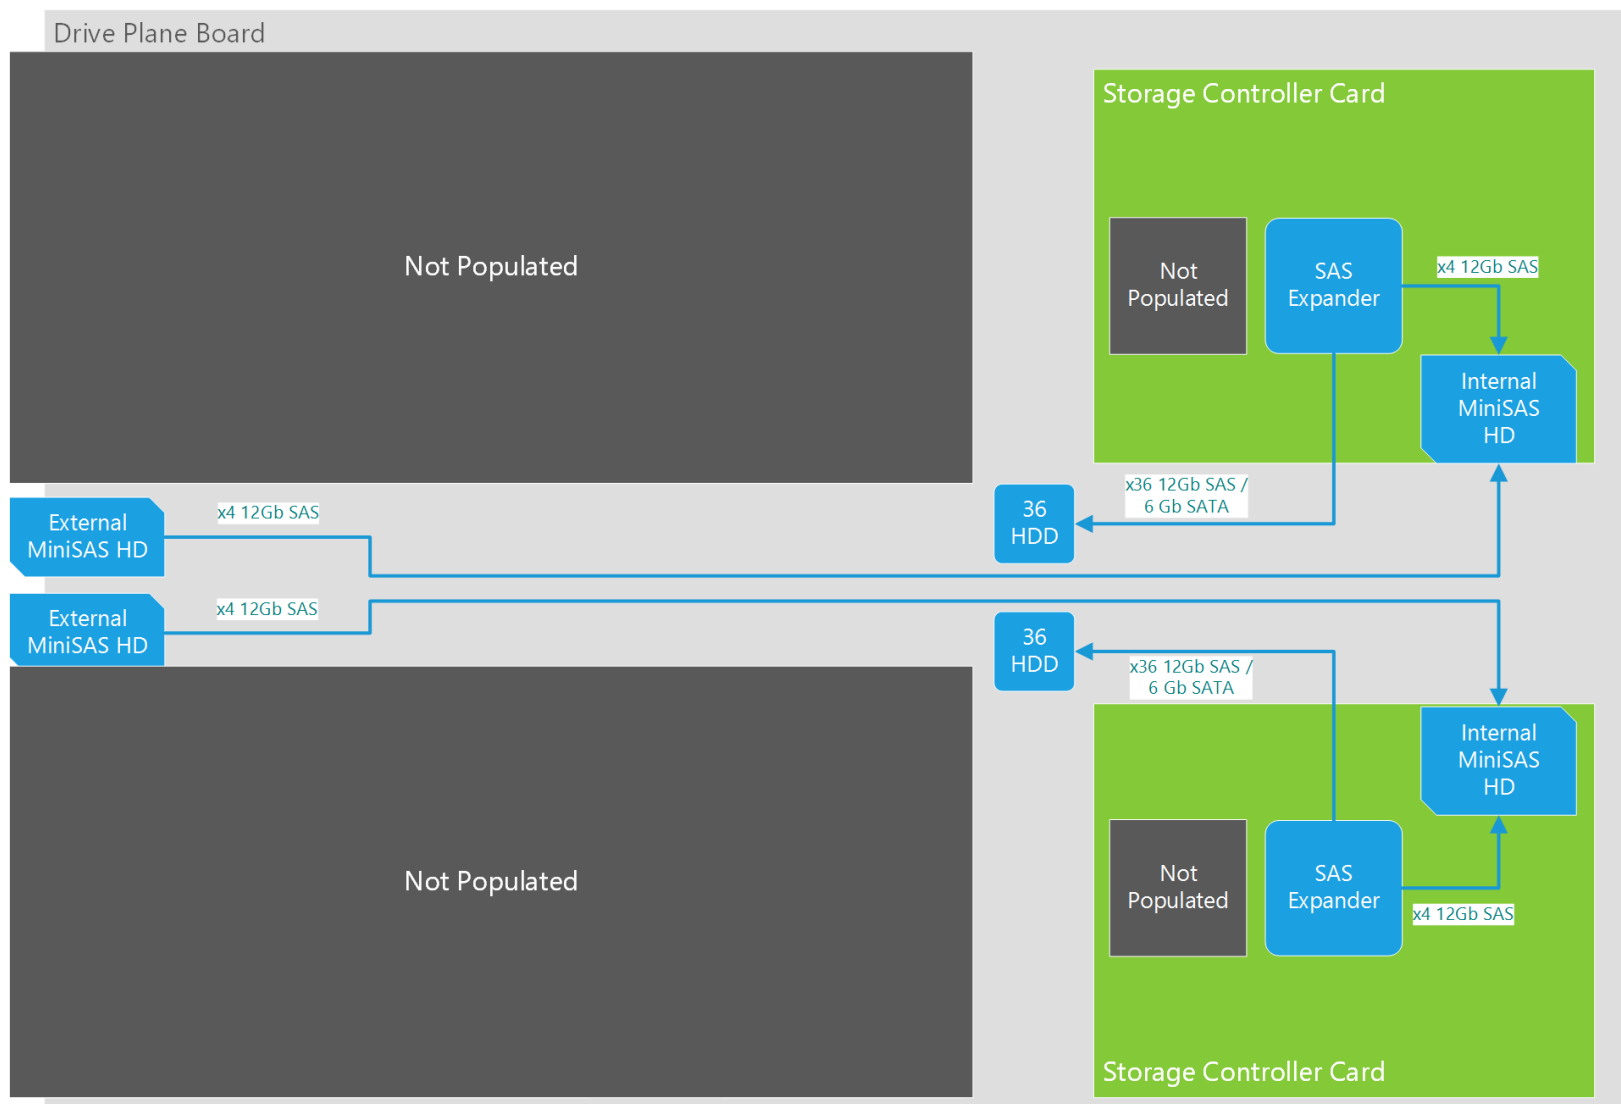


Drive Plane Board



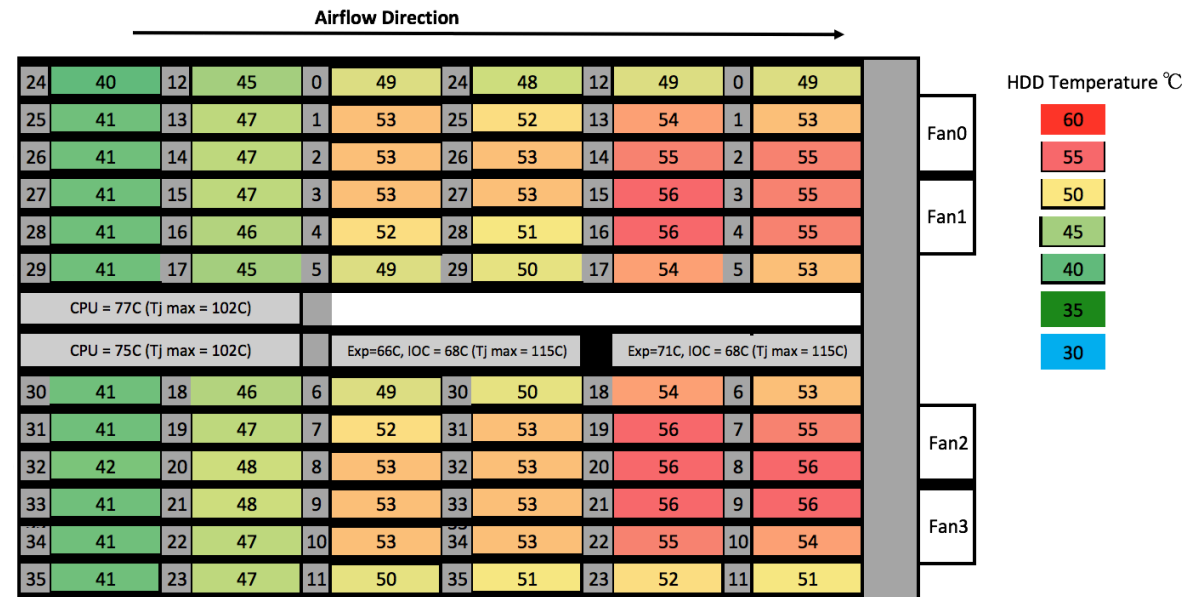
Drive Plane Board






Thermal Performance

- Target max HDD temp=55C
- Measured data for a dual storage server configuration
- CPU, DIMM, M.2, NIC and all HDD stress scripts executed in parallel



- Ambient Temperature = 30C
- Fan Duty Cycle = 30%
- Measured airflow = 125.9 CFM
- System power = 952 W
- CFM/W = 0.132



Carrier-less HDD



Handles to move the drawer



Fan Module



Evolution of storage at Facebook

	Warm Storage		Cold Storage	
	Prior Gen	Bryce Canyon	Prior Gen	Bryce Canyon
HDD per U	15	18	15	18
HDDs per Rack	450	576	480	648
Performance/ Compute Module	1 (Baseline)	6x	1 (Baseline)	6x
Max Memory	32 GB	128 GB	32 GB	128 GB
NVMe Slots (M.2)	0	2	0	0
SAS/SATA Support	6 Gbps / 6 Gbps	12 Gbps / 6 Gbps	6 Gbps / 6 Gbps	12 Gbps / 6 Gbps
Max Network Bandwidth	25 Gbps *	50 Gbps	25 Gbps *	50 Gbps

* Limited by x4 Gen 2 PCIe



Open BMC

- Bryce Canyon supports OpenBMC
- Supports SSH
- One interface to access and control everything
 - System thermal sensors across all PCBAs & HDDs
 - System fan speed algorithm
 - Record system power consumption
 - Supports fru-util, sensor-util, power-util and sol-util
- For more information on Open BMC, please visit <https://github.com/facebook/openbmc>



OPEN

Compute Project

