OCP SUMMIT

March 20-21
2018
San Jose, CA

OPEN
Compute Project

Facebook Flexible GPU Expander Big Basin Refresh

Whitney Zhao/HW Eng/Facebook Inc.

Xiaodong Wang/SW Eng/Facebook Inc.

Agenda

- Introduction ◆
- Architecture ◇
- Performance ◇
- Questions ◇

Agenda

Introduction ◆

Architecture ◇

Performance ◇

Questions ◇

Impact

Facebook’s commitment to developing AI & advancing ML

- LANGUAGE TRANSLATION
- FACE RECOGNITION
- SEARCH
- ADS
- NEWS FEED
- SIGMA
- LUMOS

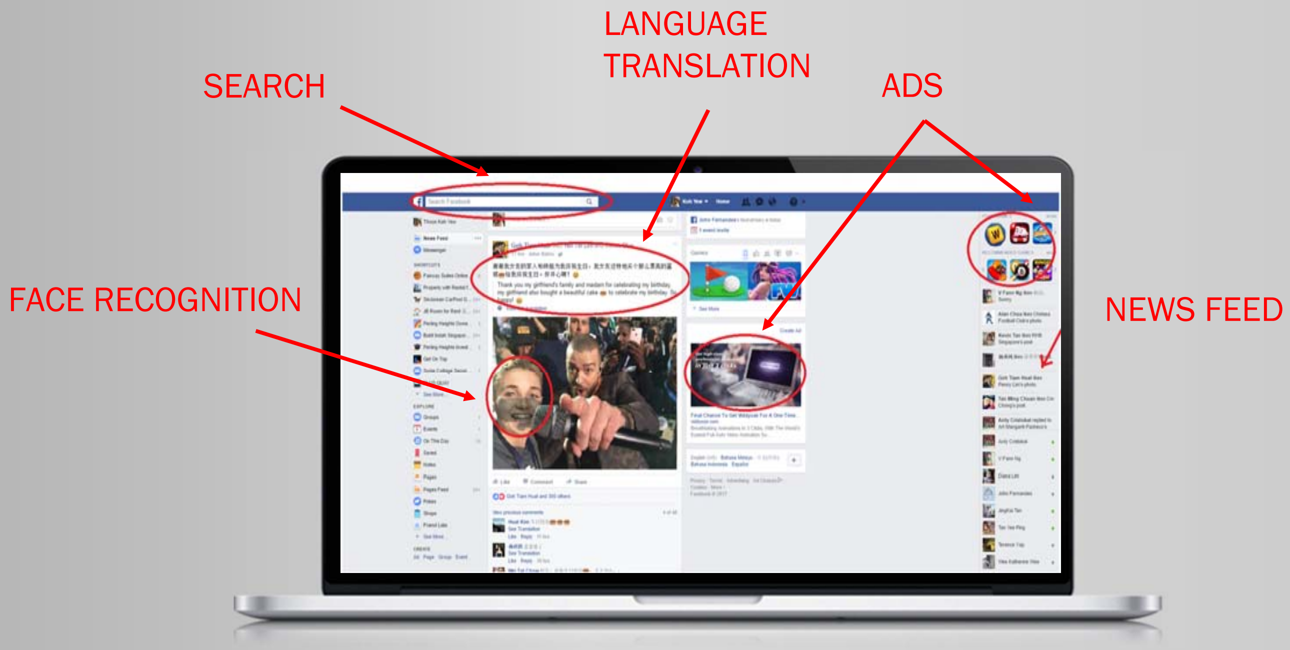

SEARCH

LANGUAGE TRANSLATION

ADS

FACE RECOGNITION

NEWS FEED

Goal

- Open, full contribution to OCP
- Disaggregation/Modularity
- Serviceability

2016: Big Sur

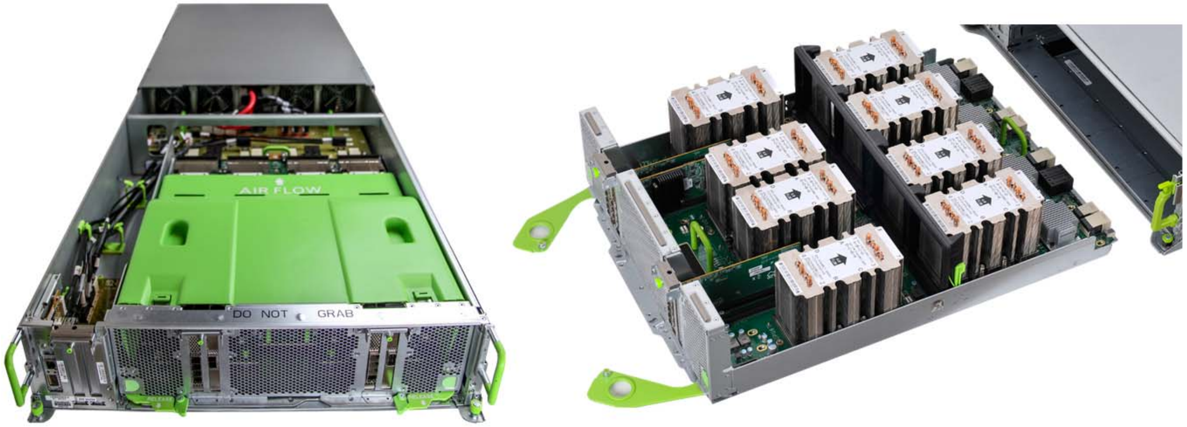

2017: Leopard + Big Basin
2018: Tioga Pass + Big Basin V2

Big Basin V2 Overview

- 3 OU chassis
- Open Rack v2 compatible
- 8x Nvidia Tesla V100 GPUs; NVLink capable
- 300W TDP for each Tesla V100 GPU
- Facebook 2S Server Tioga Pass as Head node

A deeper look into Big Basin

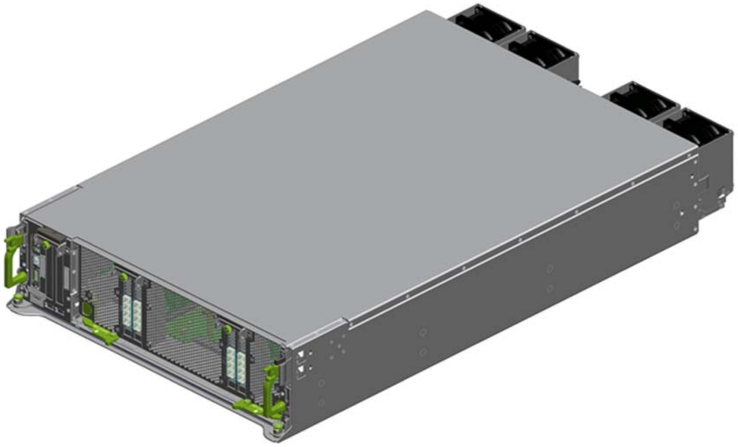

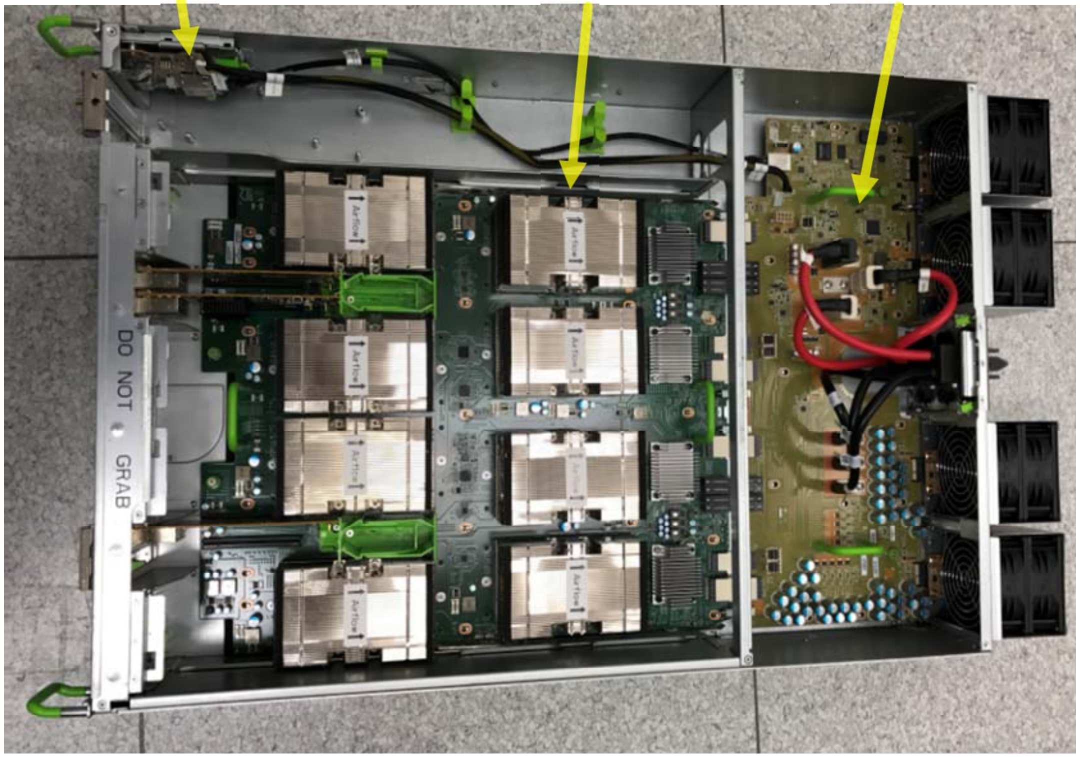

IO board

Baseboard

Midplane board

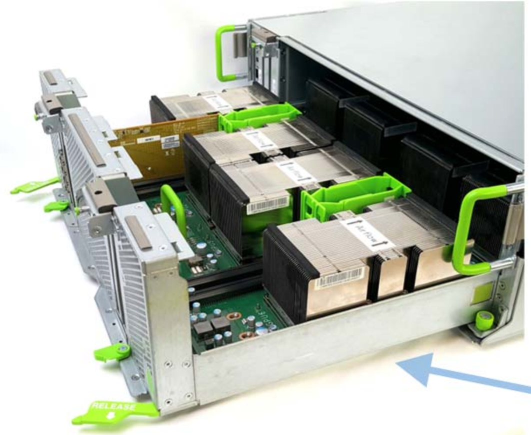

Baseboard on sliding tray

Serviceability

- Quick repairs at data center
- Telemetries accessible from head node
- Provisioning Big Basin with its head node is not much different from provisioning existing servers; these servers come with additional GPUs.

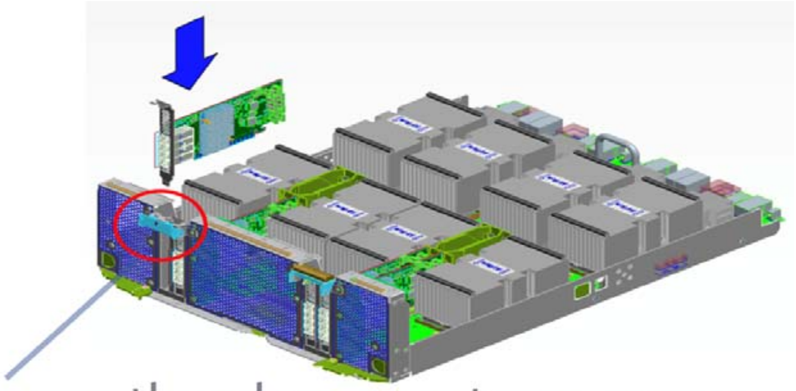

Loosen thumb screws to open cap

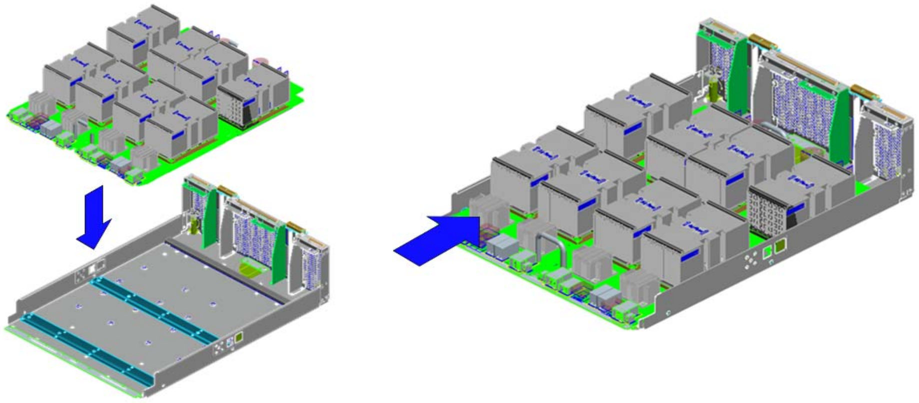

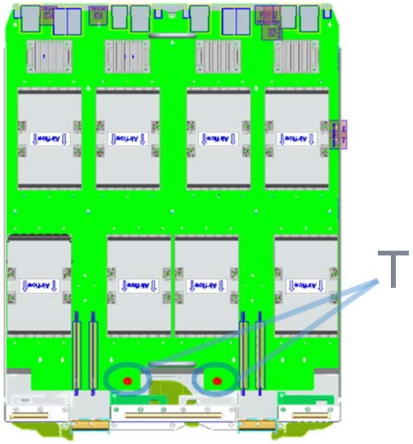

Thumb screws

Agenda

Introduction ◇

Architecture ◆

Performance ◇

Questions ◇

Architecture (Headnode to Big Basin)

- MiniSAS HD cable(2 for each x16)
  - Standard PCIe x16
  - Present Pin
  - USB2.0
  - IPMB/I2C

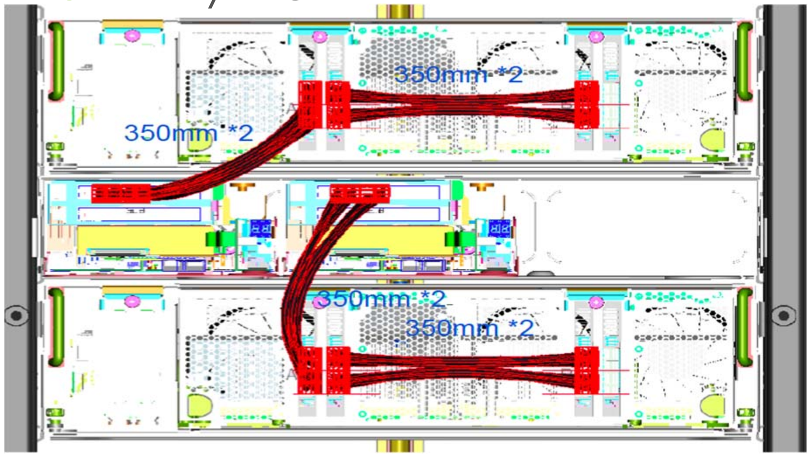

Leopard + Big Basin(Tesla P100)

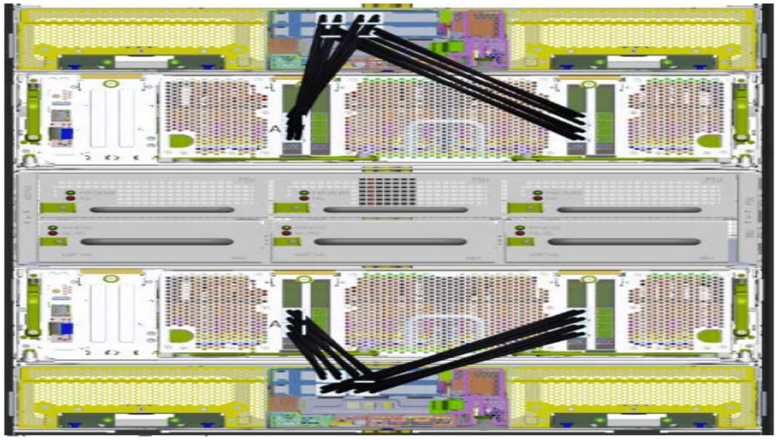

Tioga Pass + Big Basin V2(Tesla V100)

Architecture (PCIe)

Leopard + Big Basin

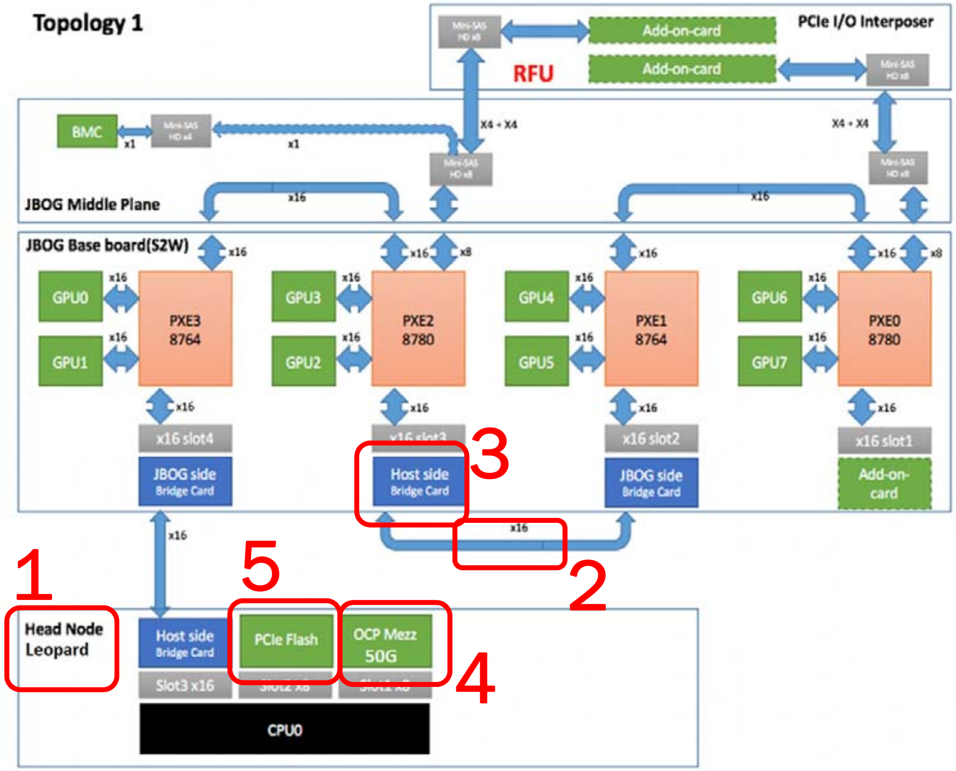

Tioga Pass + Big Basin V2

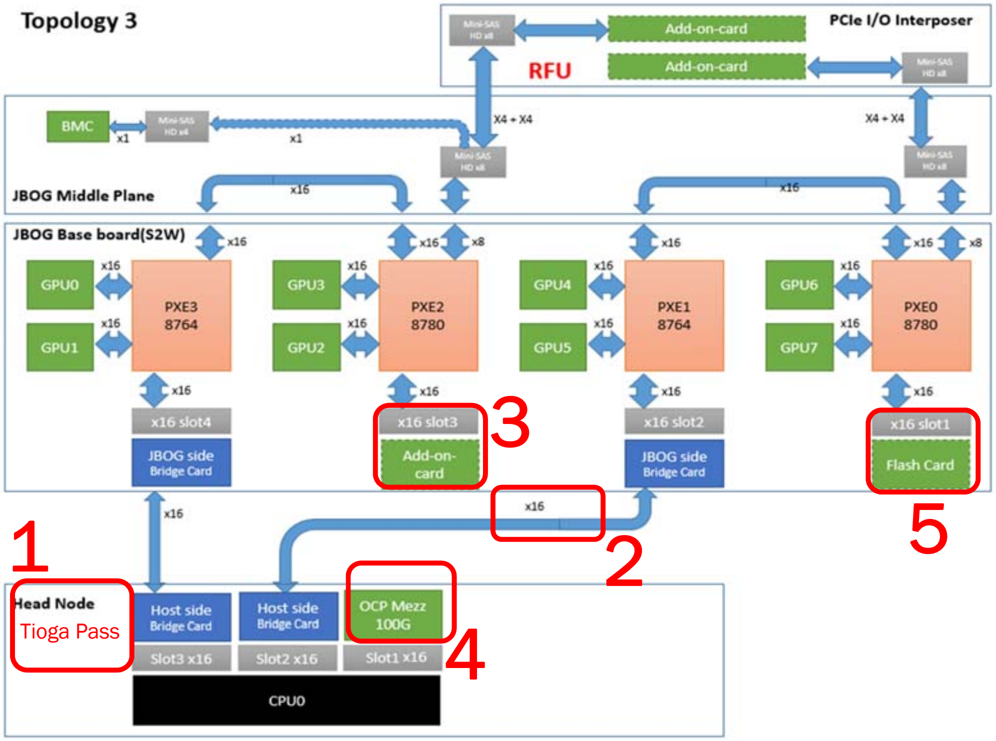

Architecture (NVLINK)

NVHS 8 lanes
20 Gbps for Pascal/25.78125Gbps for Volta

NVHS 8 lanes Volta Only
25.78125Gbps for Volta

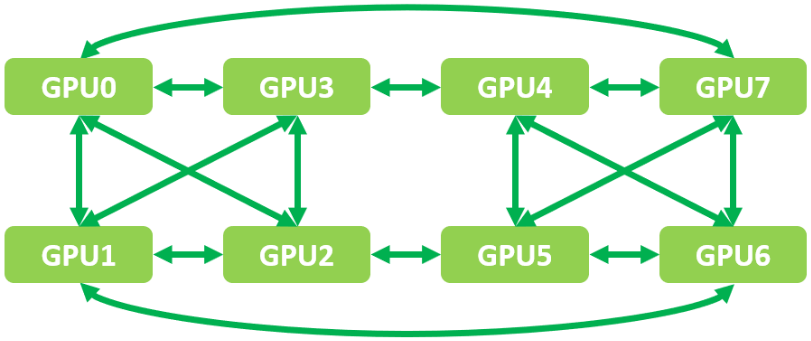

Big Basin W/Nvidia Tesla P100

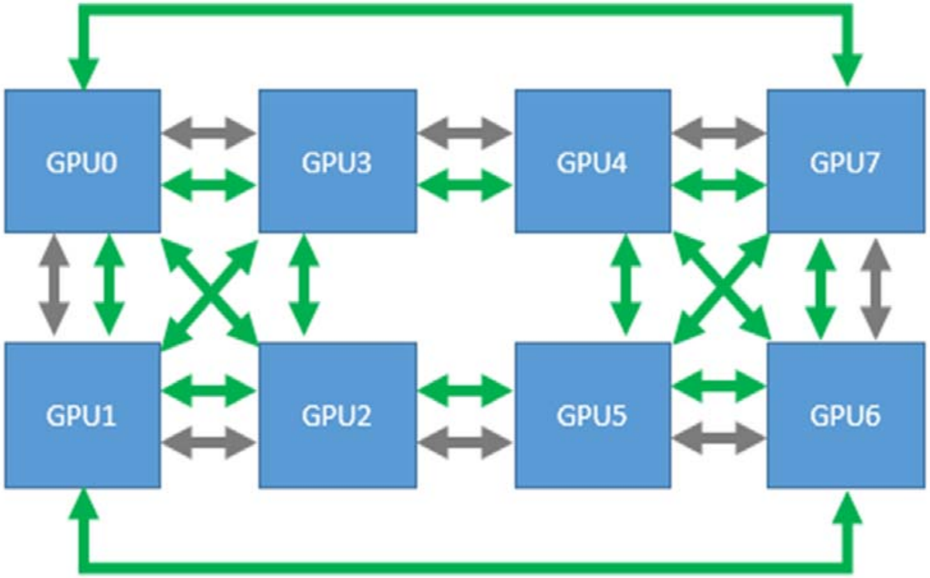

Big Basin V2 W/ Nvidia Tesla V100

Architecture (IPMB/I2C/PMBUS)

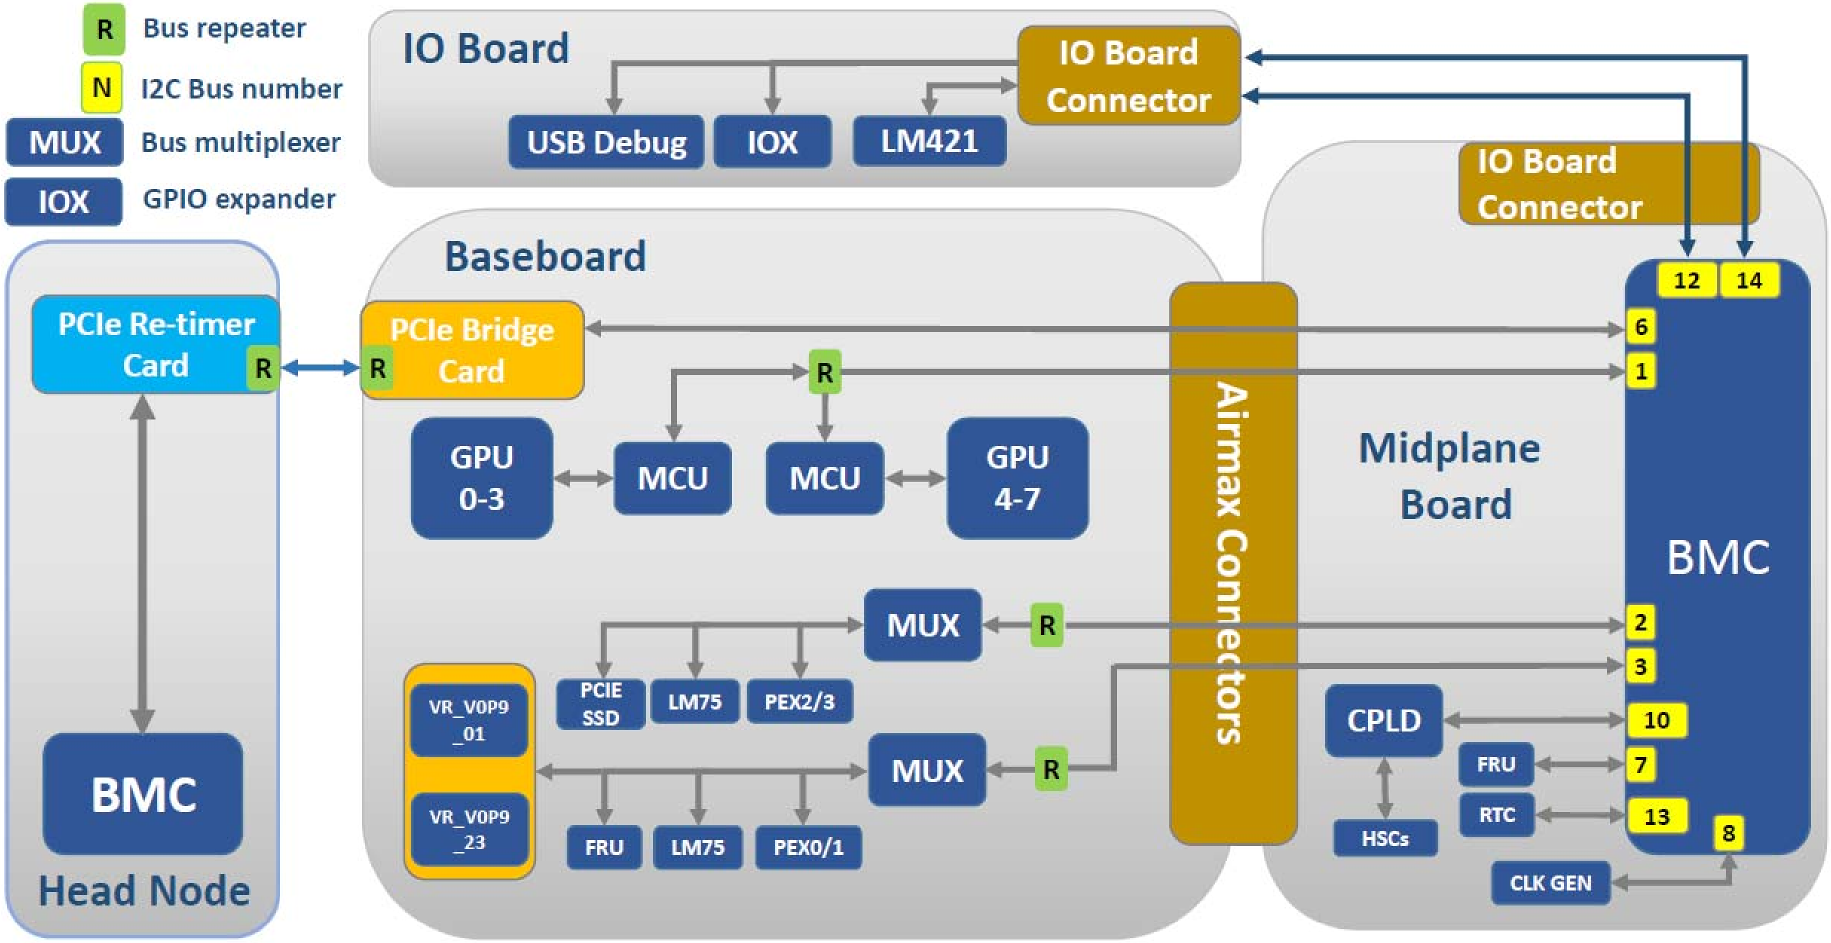

Agenda

- Introduction ◇
- Architecture ◇
- **Performance** ◆
- Questions ◇

Performance

- Hardware Spec Improvement

- Application performance
  - Computer vision
    - Single-GPU
    - Multi-GPU scalability
    - TensorCore
  - Neural machine translation

Performance

- Comparisons of GPU Hardware

	Metrics	NVIDIA V100	NVIIDA P100	Improvement
Performance	FP-32	15 TFLOPS	10.6 TFLOPS	1.42x
	FP-16	30 TFLOPS	21.2 TFLOPS	
	TensorCore	125 TFLOPS	NA	Up to 5x
	Mem Bandwidth	900 GB/s	720 GB/s	1.25x
	NVLink	300 GB/s	160 GB/s	1.88x
Power		300 W	300 W	

Performance

- Comparisons of GPU Hardware

- Head-node upgrade: Tioga Pass
  - New CPU architecture: Broadwell to Skylake
  - Double PCIe bandwidth
  - Upgraded 100G NIC

- CUDA 9 + cudnn 7: faster libraries, etc.

Impact - Computer Vision

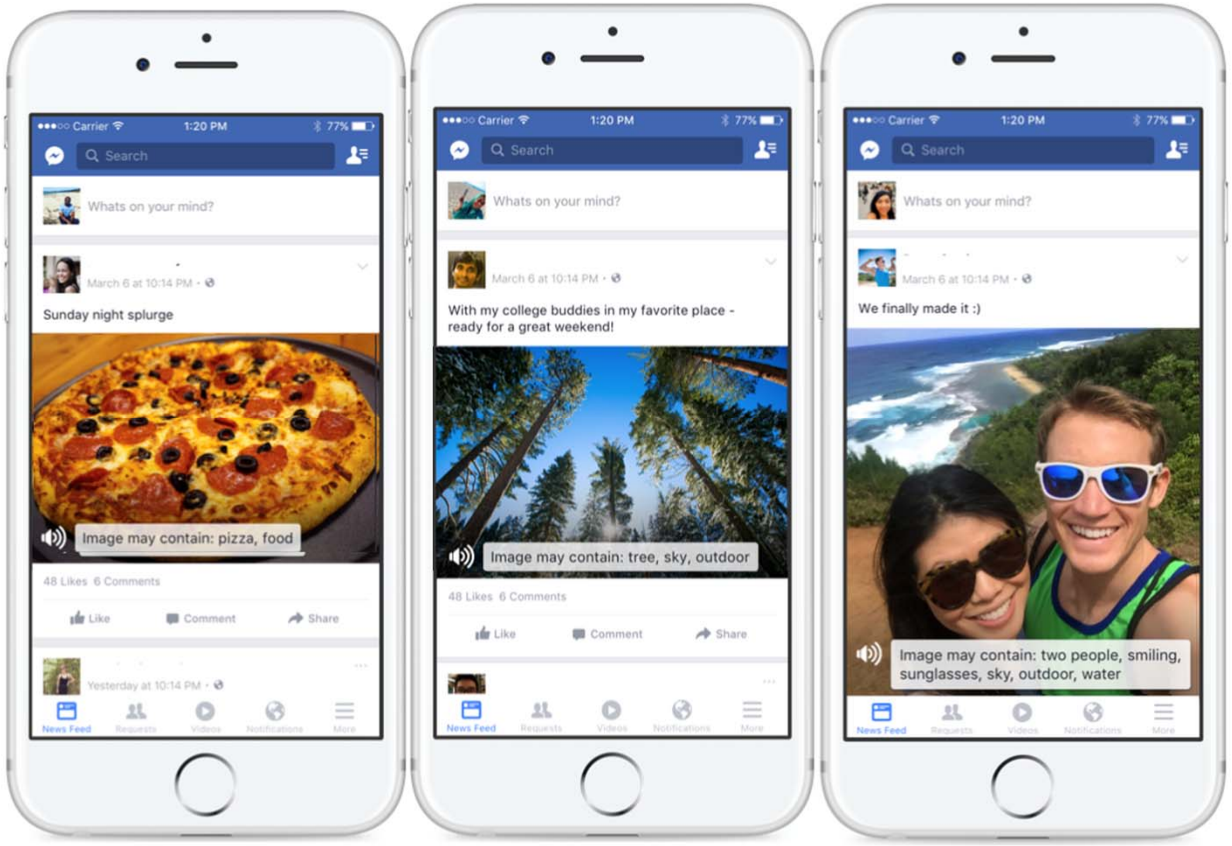

Whats on your mind?

March 6 at 10:14 PM ·

Sunday night splurge

Image may contain: pizza, food

48 Likes 6 Comments

Yesterday at 10:14 PM ·

Whats on your mind?

March 6 at 10:14 PM ·

With my college buddies in my favorite place - ready for a great weekend!

Image may contain: tree, sky, outdoor

48 Likes 6 Comments

Whats on your mind?

March 6 at 10:14 PM ·

We finally made it :)

Image may contain: two people, smiling, sunglasses, sky, outdoor, water

Performance metrics in Computer Vision

- Computer Vision: resnet-50
  - 1-GPU training speed: use P100 + CUDA 8 as baseline

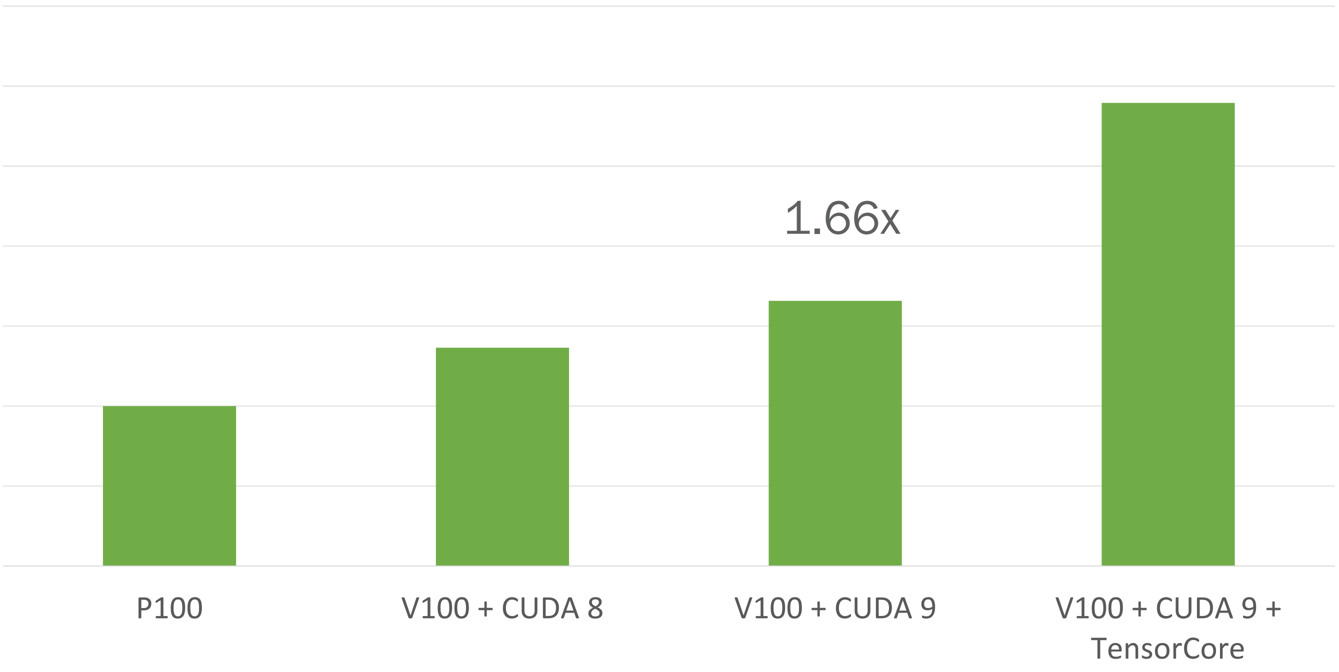

Computer Vision Performance

- Computer Vision
  - Multi-GPU speedup vs. 1 P100

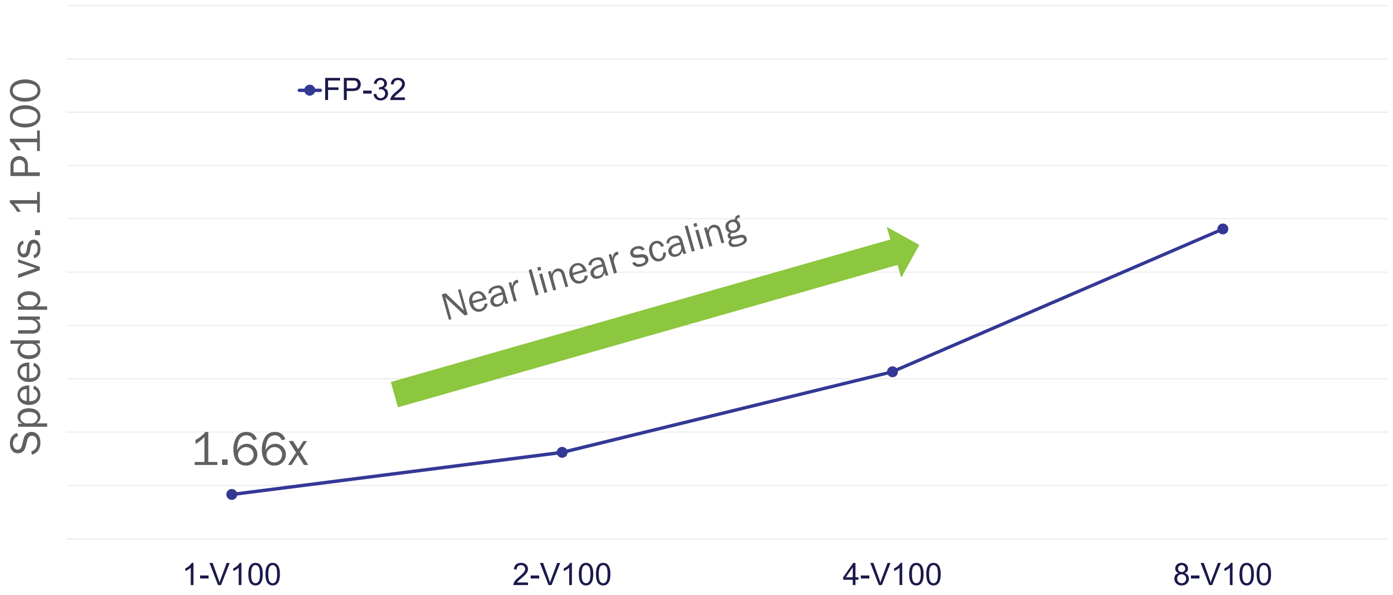

Computer Vision Performance

- Computer Vision
  - High-bandwidth FP-16 TensorCore (WIP)

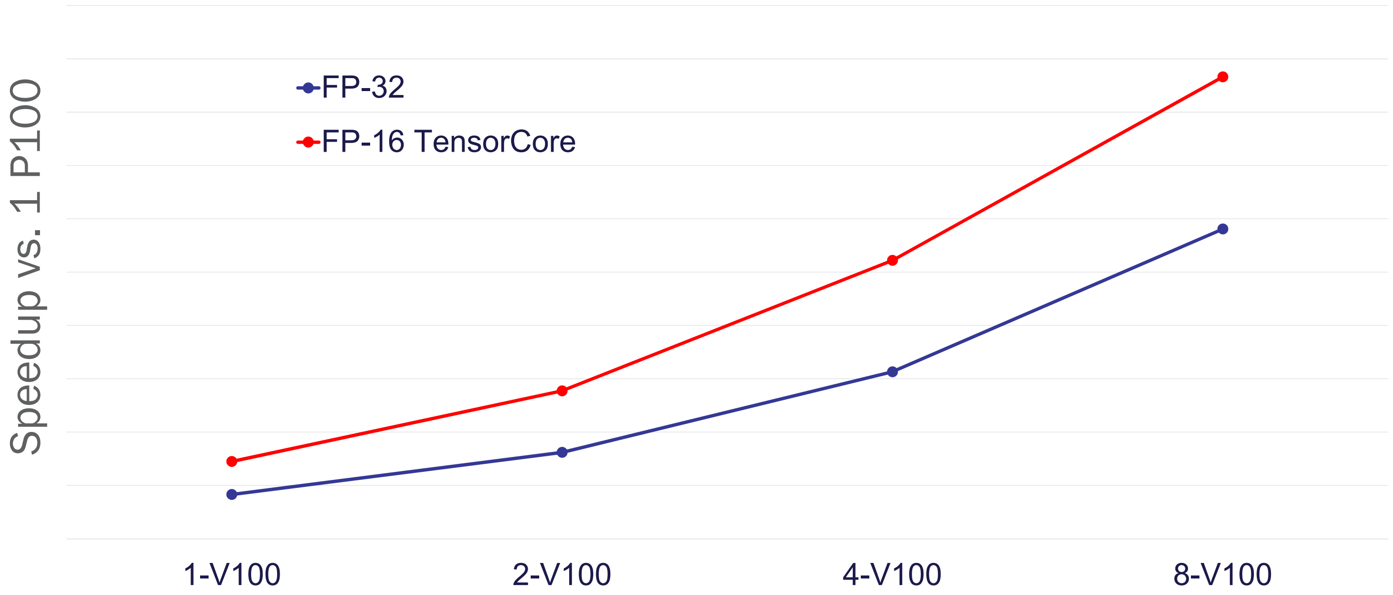

Machine Translation

Better Translation Quality

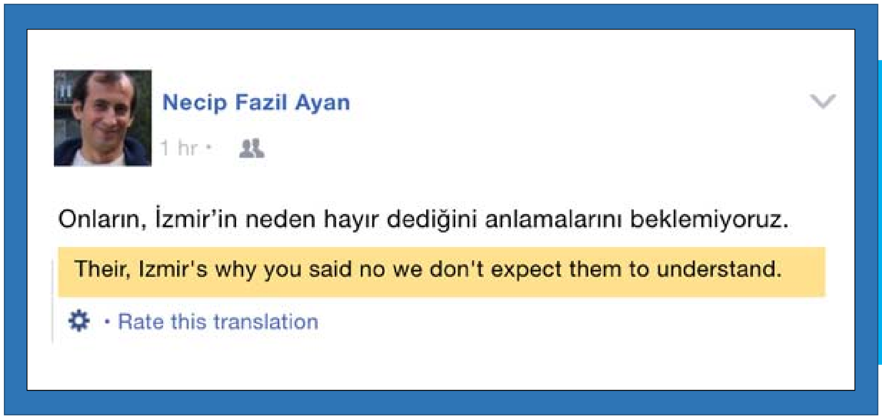

Necip Fazil Ayan

1 hr ·

Onların, İzmir'in neden hayır dediğini anlamalarını beklemiyoruz.

Their, Izmir's why you said no we don't expect them to understand.

Phrase-based statistical approach

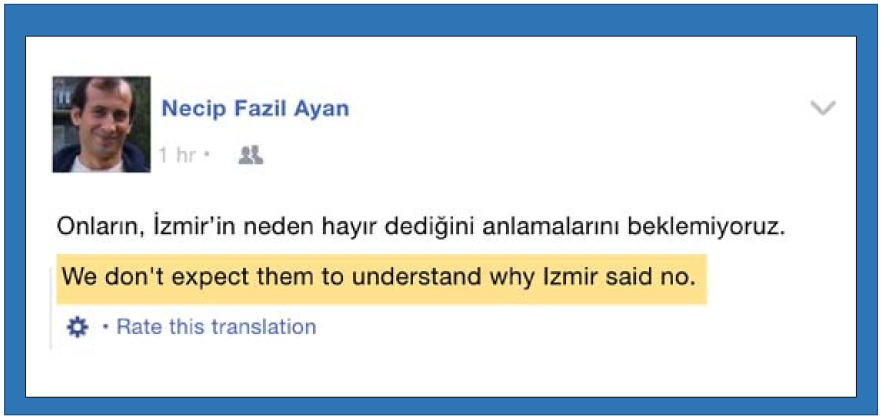

Necip Fazil Ayan

1 hr ·

Onların, İzmir'in neden hayır dediğini anlamalarını beklemiyoruz.

We don't expect them to understand why Izmir said no.

Neural network approach

Machine Translation Performance

- Neural Machine translation

P100 + CUDA 8
Training
Throughput as
Baseline

V100 + CUDA 8
1.45X

V100 + CUDA 9
2.2X

THIS JOURNEY
1% FINISHED

Questions?

OCP Marketplace

- http://www.opencompute.org/products/specsanddesign?keyword=Big+basin

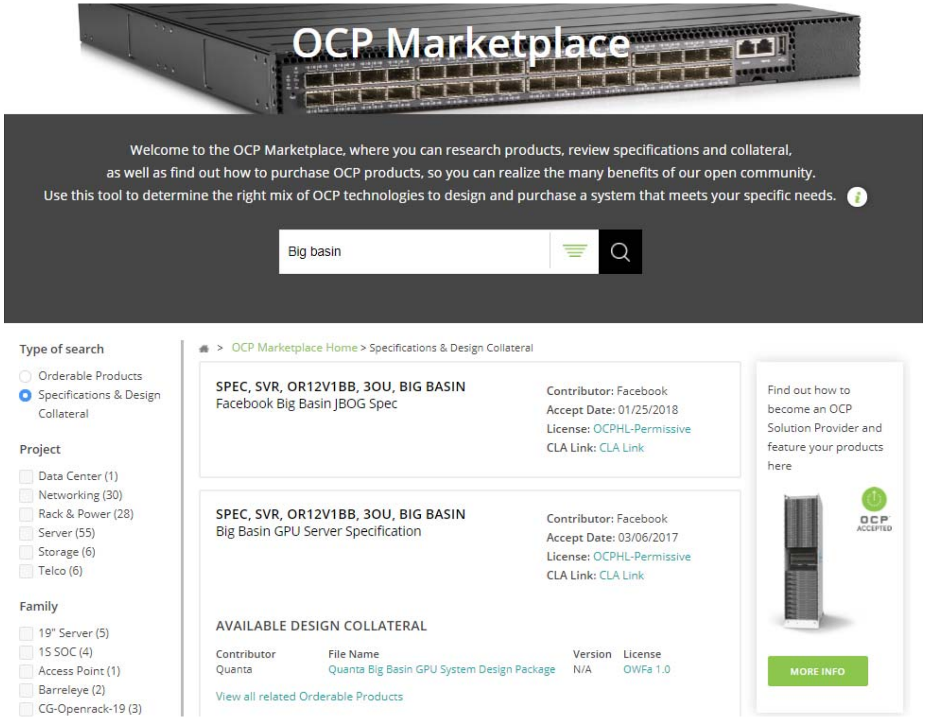

Welcome to the OCP Marketplace, where you can research products, review specifications and collateral, as well as find out how to purchase OCP products, so you can realize the many benefits of our open community. Use this tool to determine the right mix of OCP technologies to design and purchase a system that meets your specific needs.

Big basin

Type of search

○ Orderable Products
● Specifications & Design Collateral

Project

□ Data Center (1)
□ Networking (30)
□ Rack & Power (28)
□ Server (55)
□ Storage (6)
□ Telco (6)

Family

□ 19" Server (5)
□ 1S SOC (4)
□ Access Point (1)
□ Barreleye (2)
□ CG-Openrack-19 (3)

> OCP Marketplace Home > Specifications & Design Collateral

SPEC, SVR, OR12V1BB, 3OU, BIG BASIN
Facebook Big Basin JBOG Spec

Contributor: Facebook
Accept Date: 01/25/2018
License: OCPHL-Permissive
CLA Link: CLA Link

SPEC, SVR, OR12V1BB, 3OU, BIG BASIN
Big Basin GPU Server Specification

Contributor: Facebook
Accept Date: 03/06/2017
License: OCPHL-Permissive
CLA Link: CLA Link

AVAILABLE DESIGN COLLATERAL

Contributor	File Name	Version	License
Quanta	Quanta Big Basin GPU System Design Package	N/A	OWFa 1.0

View all related Orderable Products

Find out how to become an OCP Solution Provider and feature your products here

OCP
SUMMIT