



OPEN
Compute Project



OCP U.S. SUMMIT 2017

Santa Clara, CA



Wedge 100S Hardware Overview

Xu Wang
Hardware Engineer
Facebook, Inc.

OPEN HARDWARE.

OPEN SOFTWARE.

OPEN FUTURE.



Wedge 100S

An Improved Wedge 100

19-in SKU

For regular rack



21-in SKU

for Open Rack



Wedge 100S

Open 32x100GE TOR Switch

Facebook's second generation
Open TOR Switch based on
Tomahawk Switch ASIC

Standard COM-Express
module as control plane
CPU

BMC: Aspeed AST1250

Supports Open BMC and
FBOSS

Wedge 100S

Benefits and Applications



Connecting next generation servers

- Increased link speed with 50G and 25G NIC
- Increased compute density
- Lower oversubscription ratio



Improved system thermal design

- Enable adoption of 55°C 100G CWDM4 optic transceivers



Support for Open Rack V2

- Powered by Rack Bus Bar
- Rack Monitor



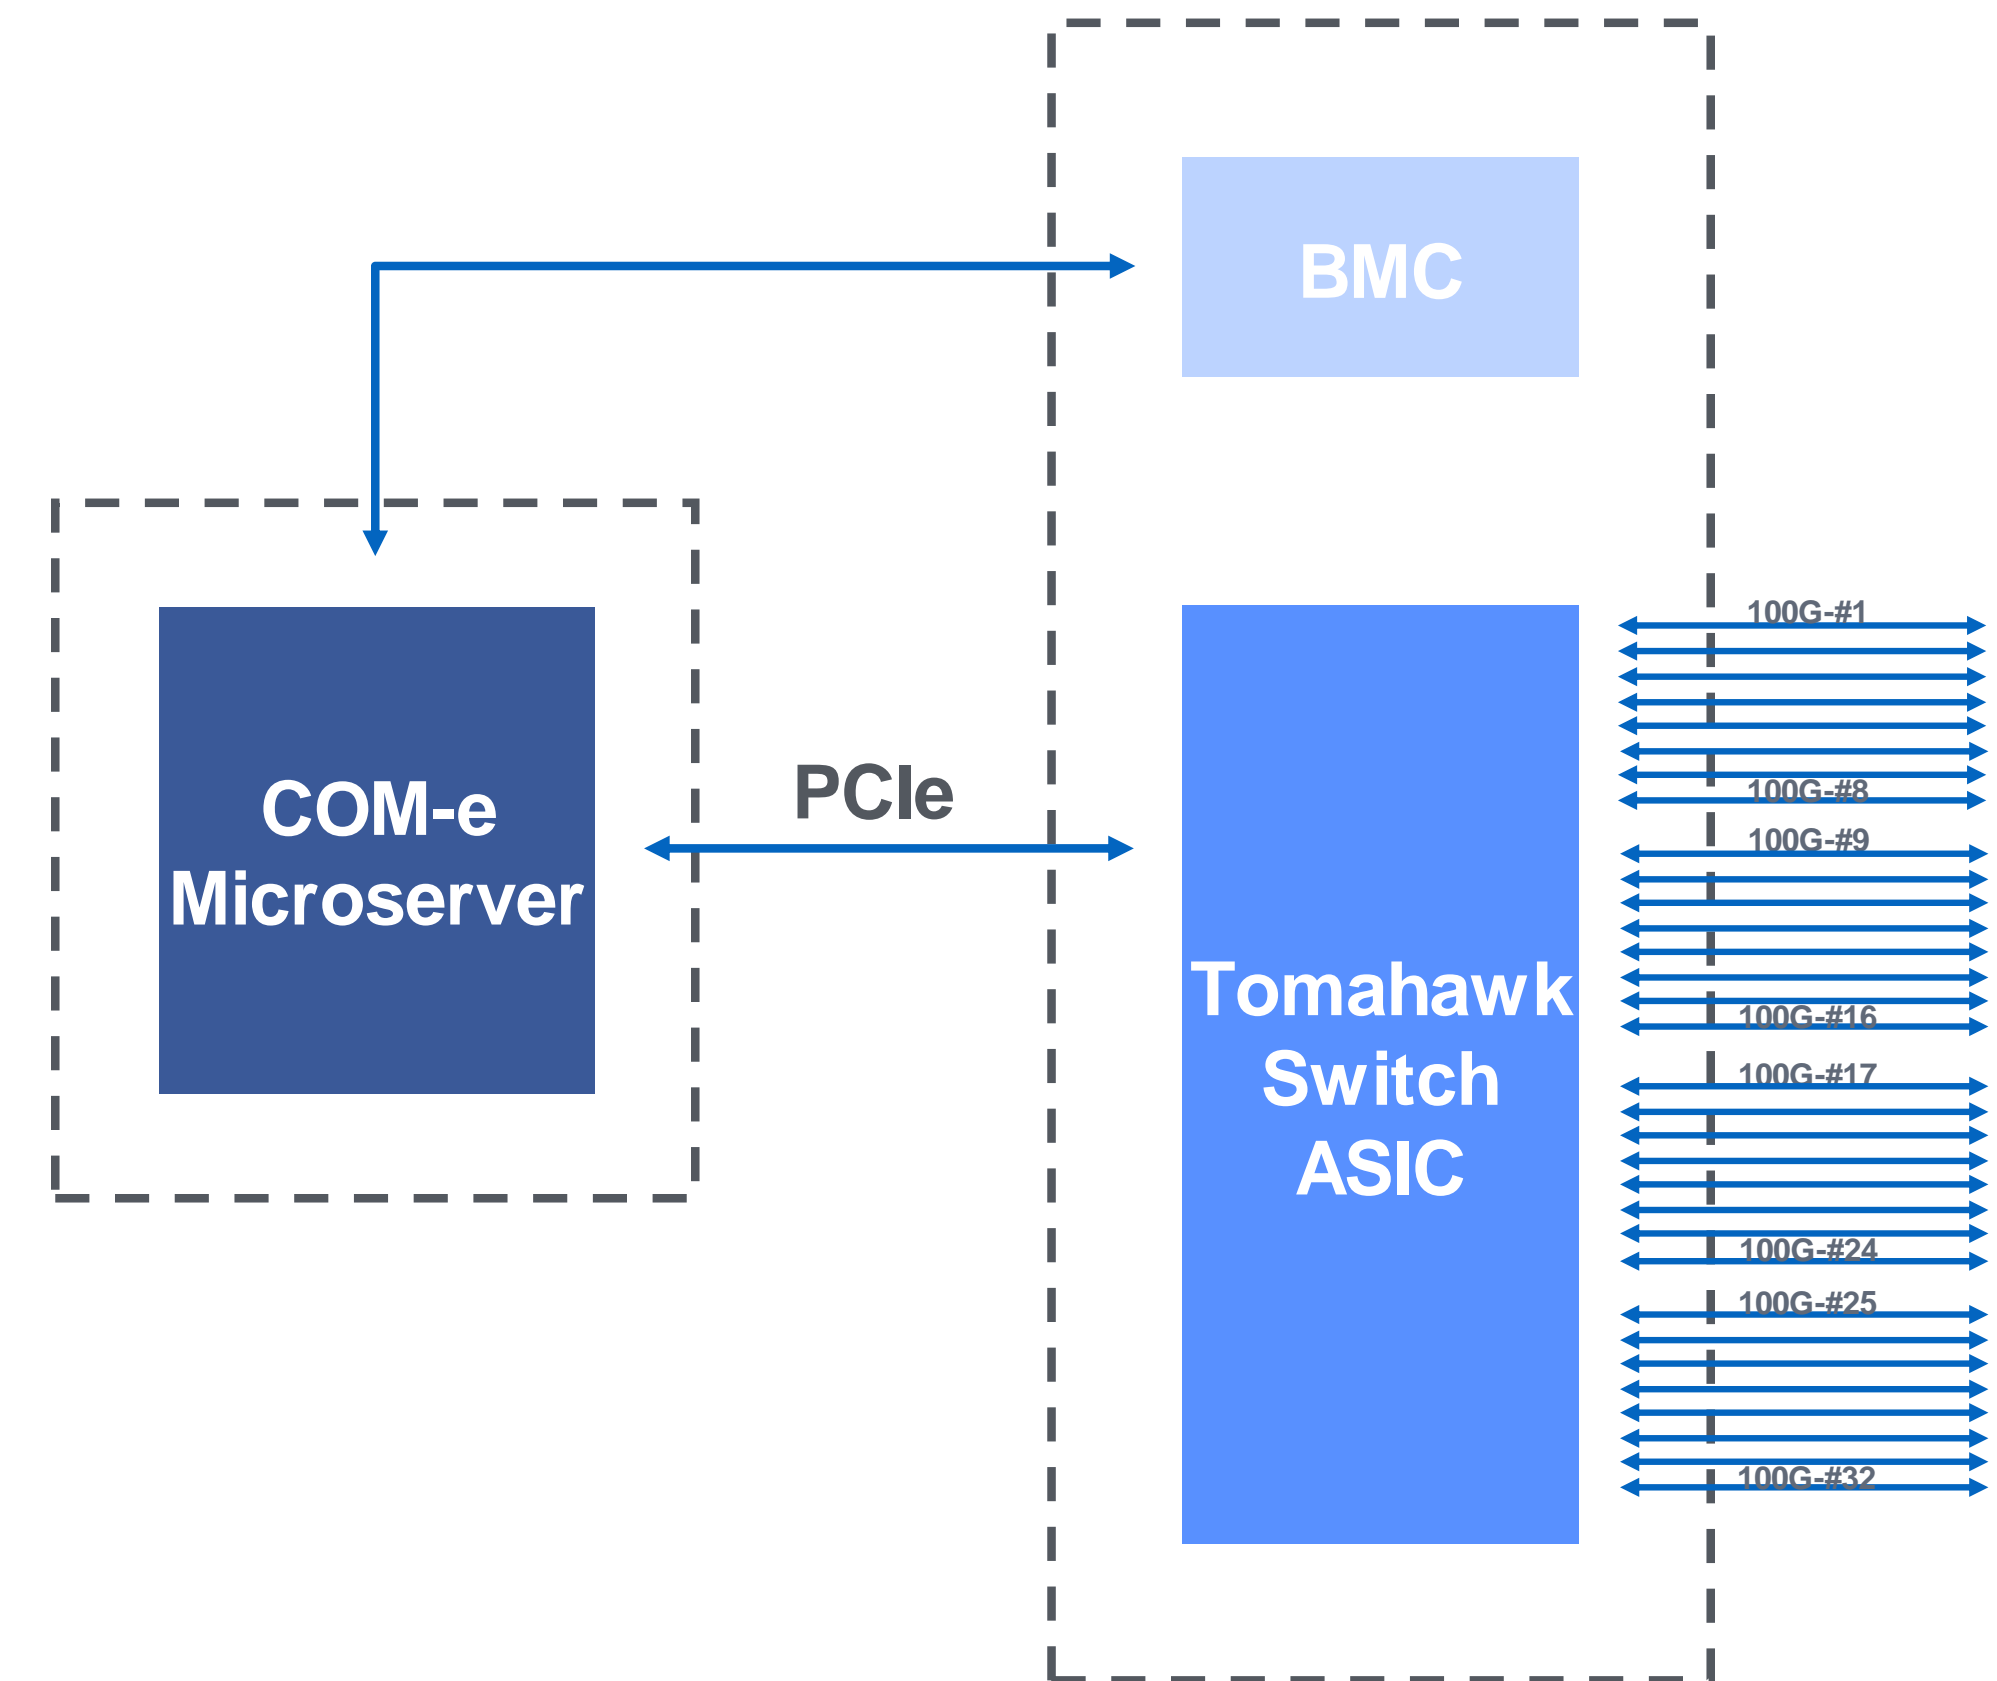
Broadwell-DE micro-server

- Higher performance
- Better RAS
- Intel Trusted Execution Technology (TXT)

Wedge 100S

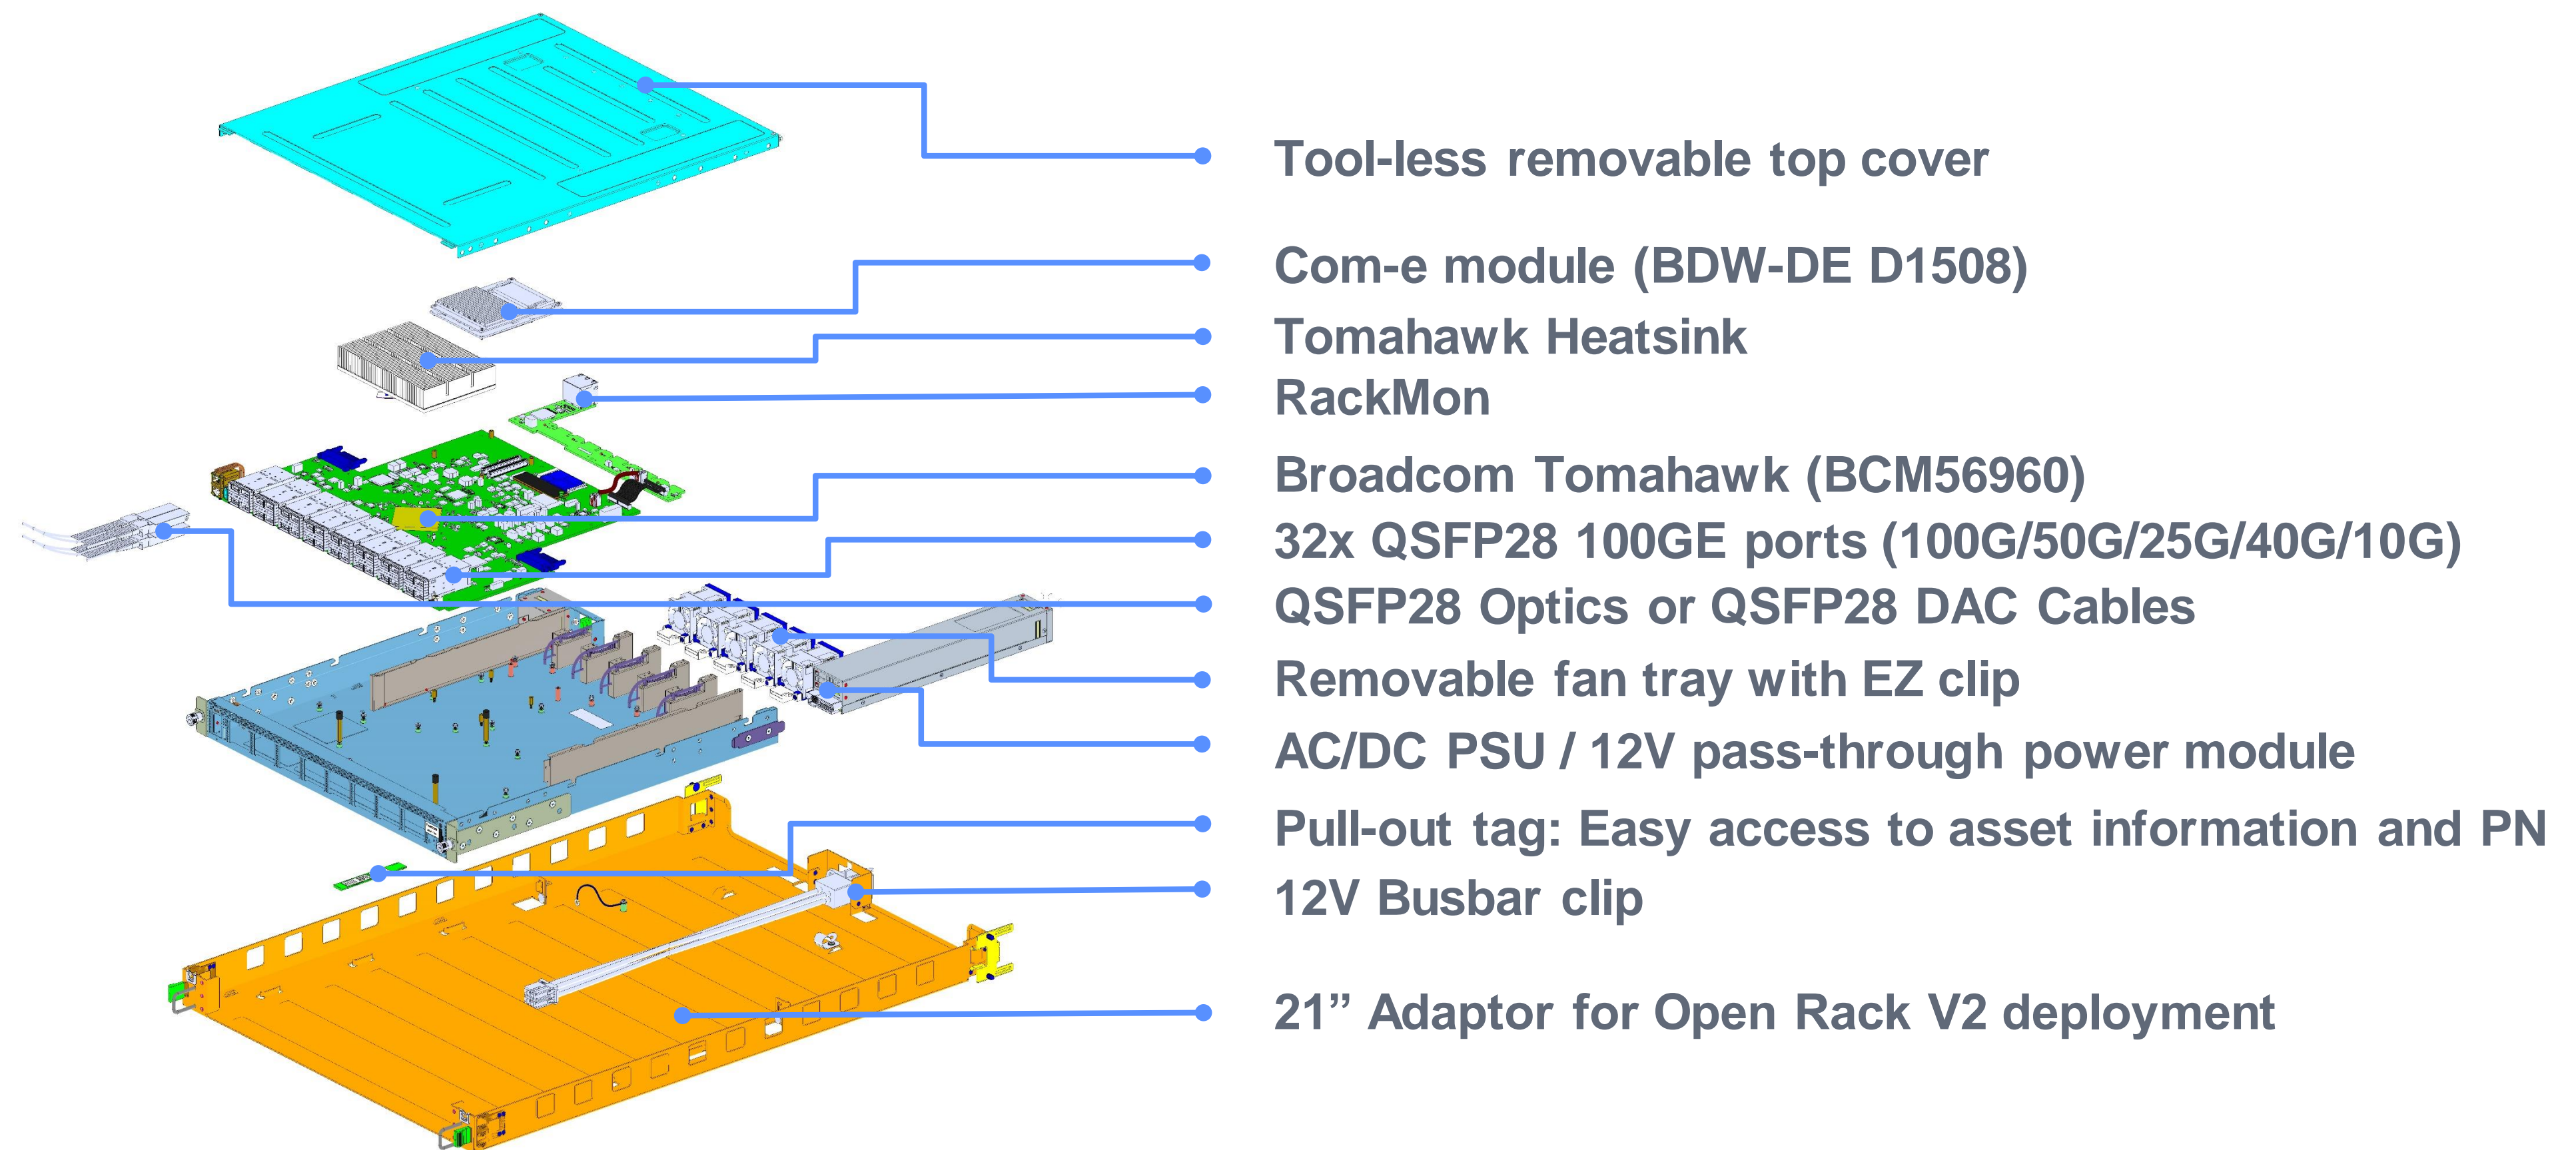
Open Hardware

- Hardware architecture: fully open
- Software: FBOSS and OpenBMC
- Manageability: operated just like a server
- Density: 32x QSFP28 100GE support
- Hardware building blocks
 - Broadcom Tomahawk ASIC, 32x 100GE, 3.2Tbps
 - COM-Express Microserver, Intel Broadwell-DE D1508
 - BMC, Aspeed AST1250




Wedge 100S

Component View




Wedge 100S

Top View



Wedge 100S


Front View



Wedge 100S


Rear View

- Fan Tray changed to latch design for tool less insertion or removal
- Fan LED now can be displayed on rear panel to indicate fan tray failure.



Wedge 100S


21-in Adaptor Tray



Wedge 100S

OCP Power Pass-through Card

- Power pass-through card with the same form factor as the AC/DC PSU



12V bus bar power
input connector



Connector to
Wedge-100 main board



Wedge 100S

QSFP28 Port Modes

1x100G: use CWDM4 QSFP28
optic to connect to 100G
aggregation switch

2x50G: use QSFP28-
2xQSFP28 DAC cable to
connect to 50G NIC

4x25G: use QSFP28-4xSFP28
fan-out cable to connect to
25G NIC

1x40G: support QSFP+ 40G
SR4 optic and QSFP+ 40G
LR4 optic

4x10G: use QSFP-4xSFP fan-
out cable to connect to 10G
SFP+ NIC



Wedge 100S

Optical Transceivers


- **QSFP28 100G optics**
 - CWDM4
 - Can support other MSA, such as SR4, LR4, CLR4, etc.
- **QSFP+ 40G optics**
 - QSFP+ 40G SR4 optic (multi-mode fiber OM3/OM4)
 - QSFP+ 40G LR4 optic (single mode fiber)



Wedge 100S

DAC Cables

- **25Gbps/lane class cables**
 - QSFP28 100G to QSFP28 100G cable, 1M, 2M, and 3M
 - QSFP28 100G to two QSFP28 50G cable, 1M, 2M, and 3M
 - QSFP28 100G to four SFP28 25G cable, 1M, 2M, and 3M
- **10Gbps/lane class cables**
 - QSFP+ to 4xSFP+ 40G fan-out cable



Wedge 100S

Thermal Design

1

Supports 55°C optic in 35°C environments

2

Five fan-trays on the rear panel

3

PSUs have a separate air channel to avoid air recycling

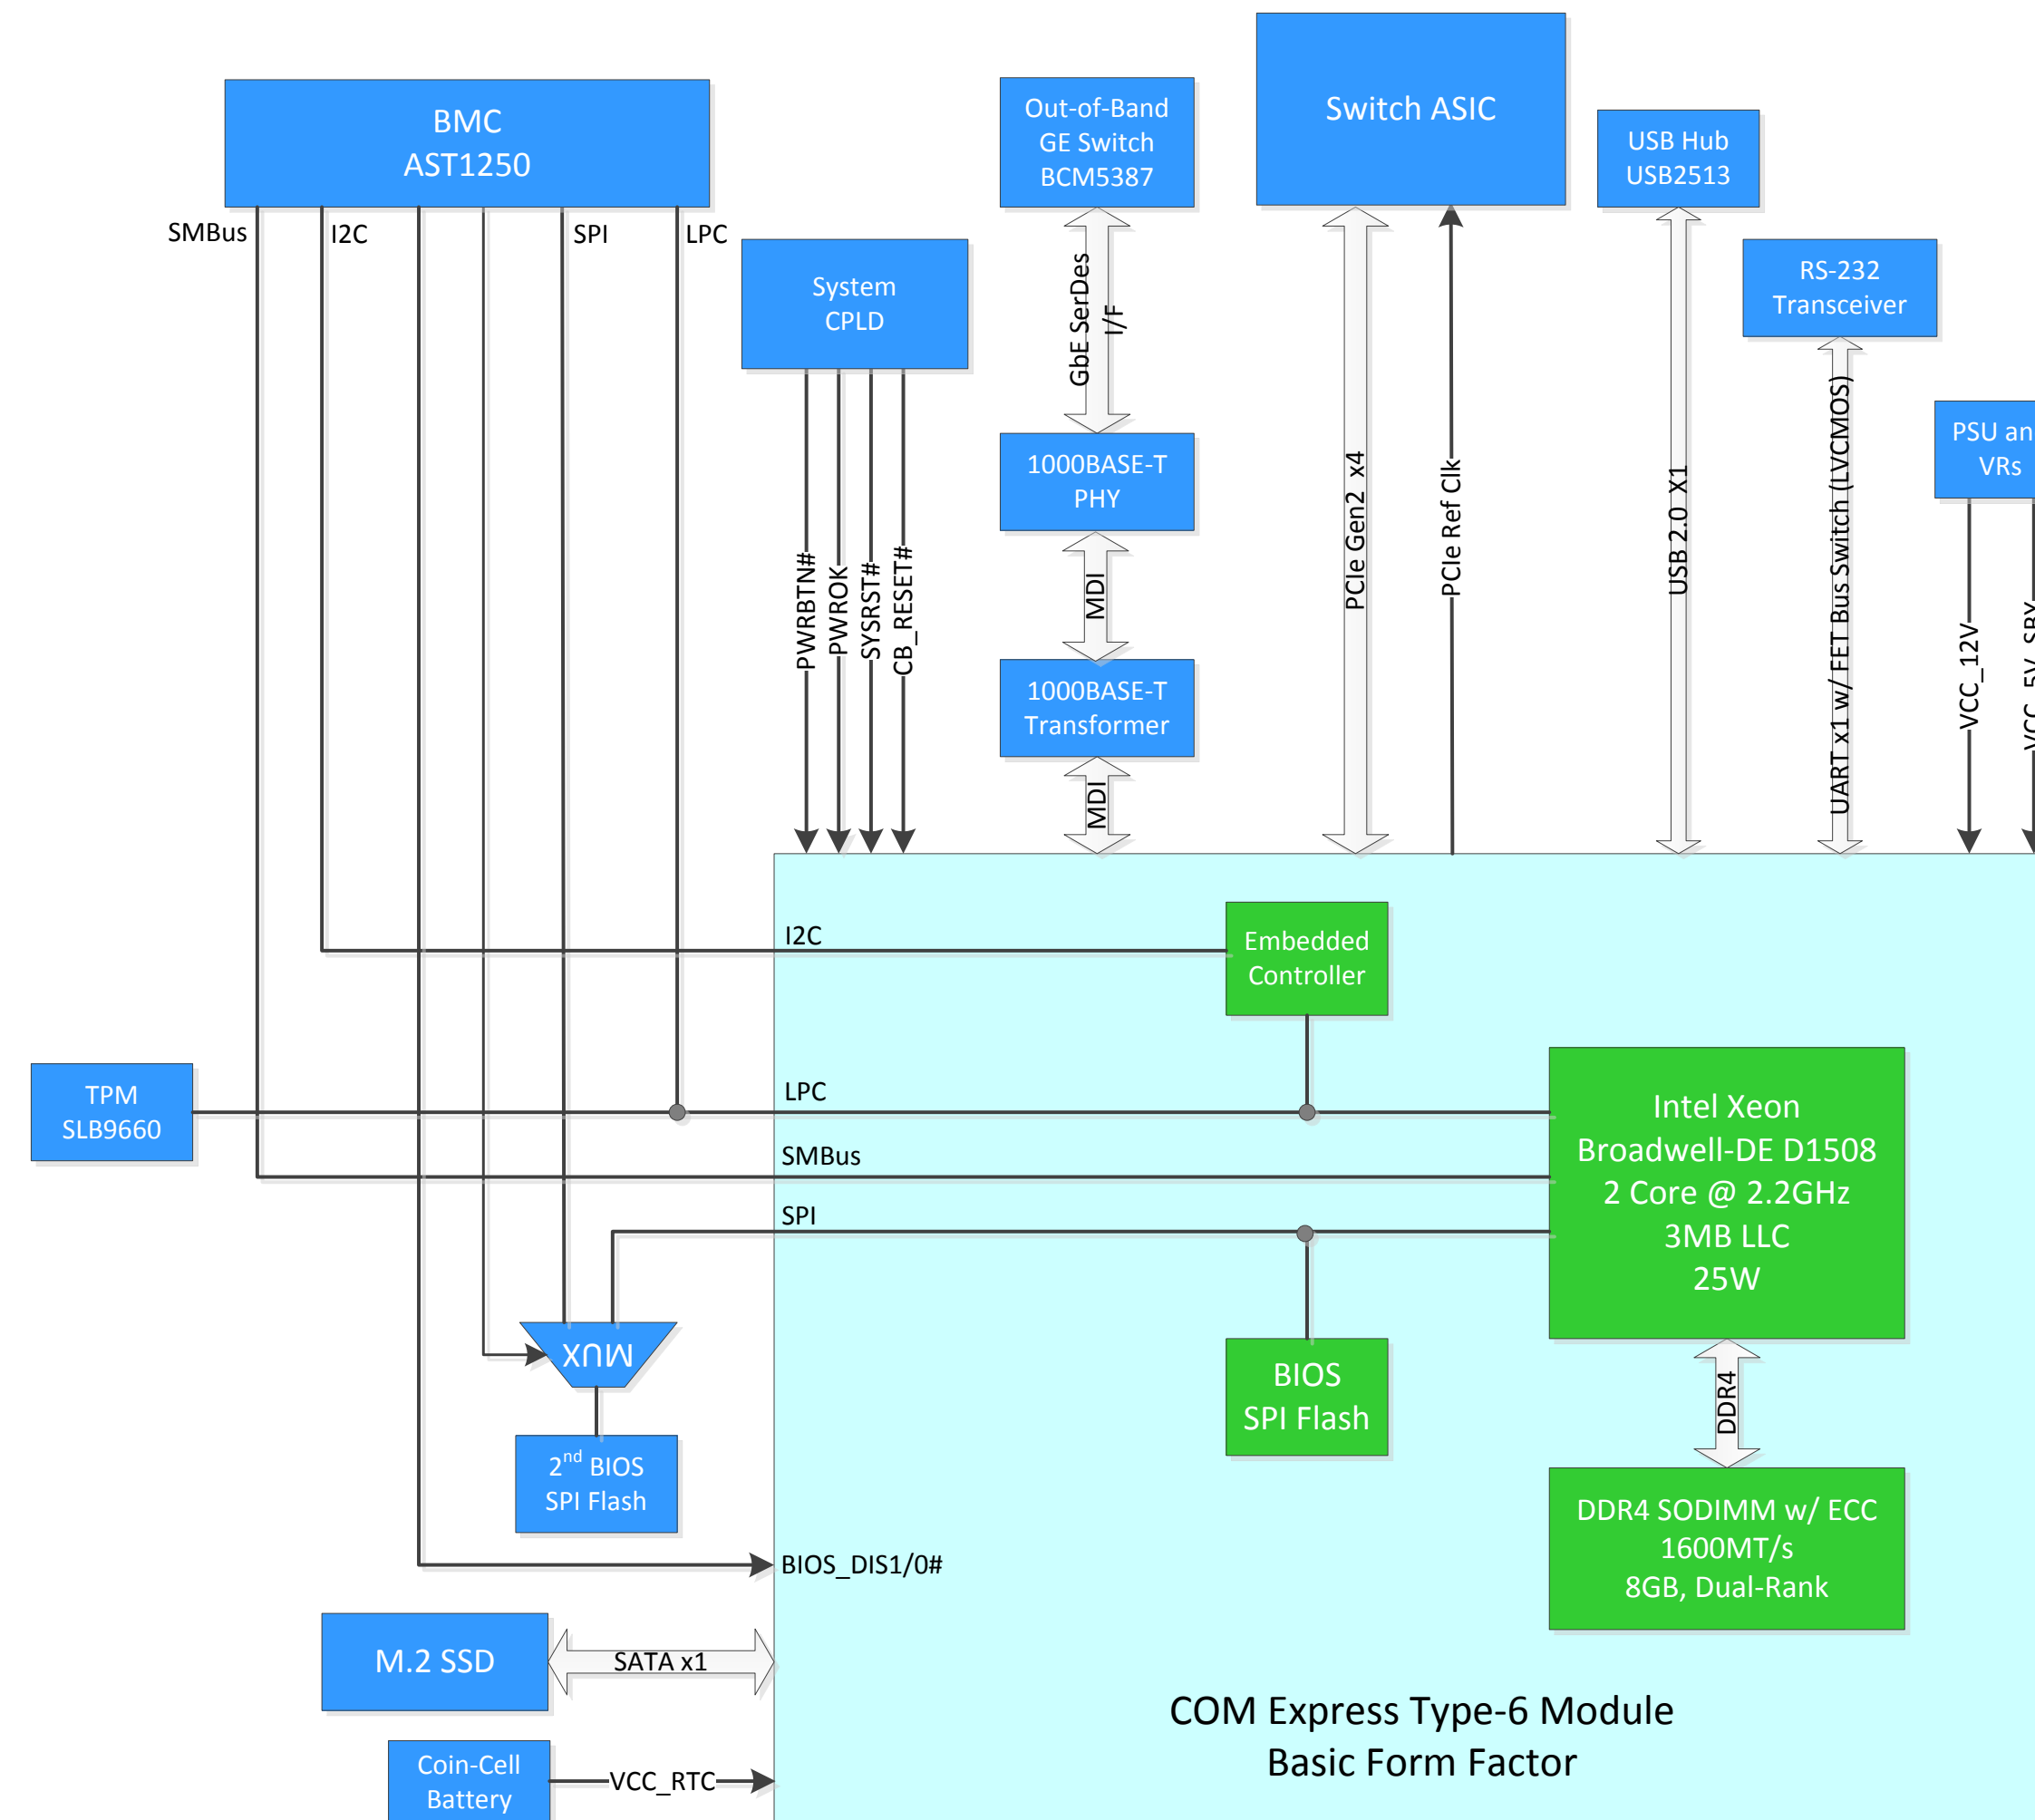
4

Multiple on-board temperature sensors to monitor thermal status of the system



Wedge 100S

COM-Express Module Interfaces



Wedge 100S


CPU SKU Selection Criteria and Process

- ToR power efficiency – important for data center applications
- Minimize increase in power consumption while migrating to Broadwell-DE
- No performance degradation
- Better performance if needed

Wedge 100S

OCP Contribution

- Wedge100S specification
- Wedge100S design package



OPEN
Compute Project

<http://www.opencompute.org/wiki/Networking/SpecsAndDesigns>

Questions?



OPEN

Compute Project



Backup


Wedge 100S

Fan Tray

- Screwless latch design for easy maintenance
- Powerful Counter-Rotating fans
- LED on the rear panel



Wedge 100S Temperature Sensor Placement



QSFP28 100G to QSFP28 100G Cable

Four lanes, QSFP28 at both ends
25Gbps per lane

QSFP28 100G
All 4 lanes used


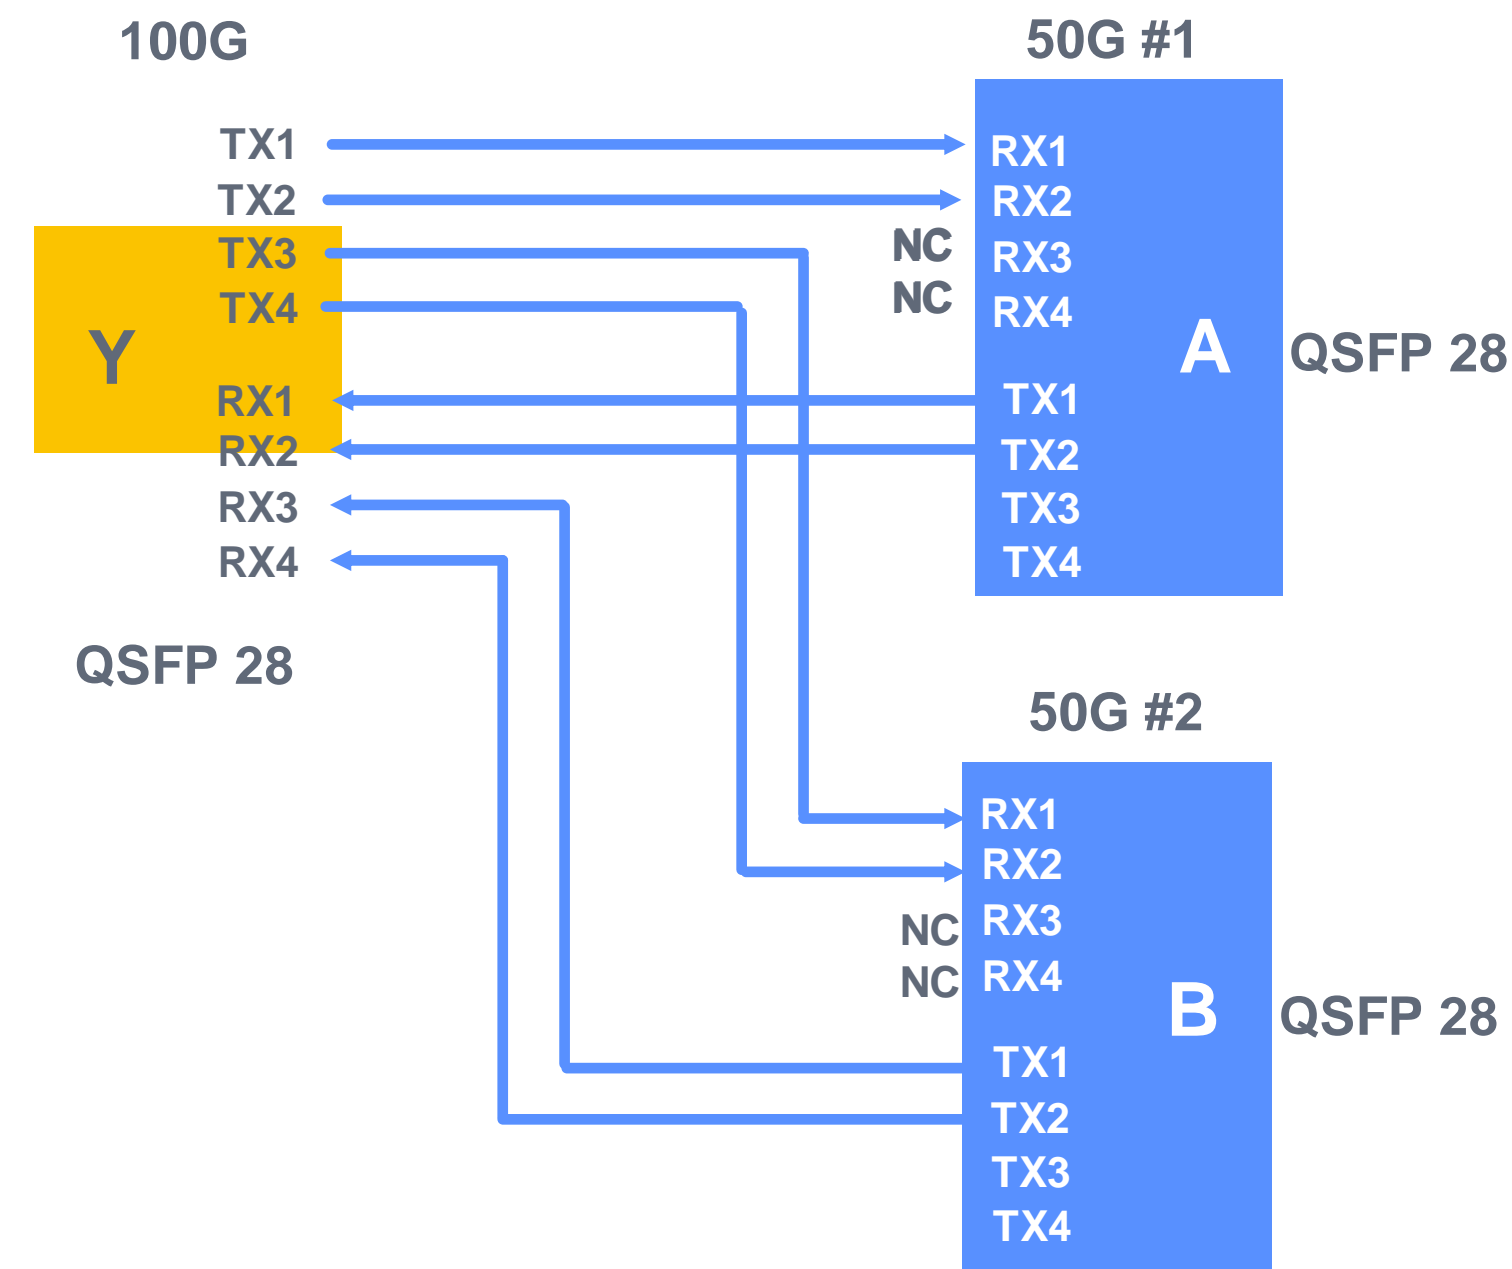



QSFP28 100G
All 4 lanes used

QSFP28 100G to 2x QSFP28 50G Cable

- a.k.a Y-Cable, this was for Wedge 100/S to inter-operate with Yosemite 50G NIC.
- Yosemite 50G CX4 NIC uses QSFP28 form factor, but only 2 lanes used.
- Header Y use all 4 lanes, header A and B use 2 lanes.
- All lanes are 25Gbps.

Facebook 100G to 2x50G Y Cable Pinout Diagram



QSFP28 100G to 4x SFP28 25G Cable

- Four lanes in QSFP28 side, 25Gbps per lane
- One lane at SFP28 side, 25Gbps
- 1M, 2M, 3M

**Four SFP28 25G
header, each SFP28
supports
25Gbps single lane**

**QSFP28 100G
All 4 lanes used**

