

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

OCP U.S. SUMMIT 2016

Lightning: A flexible NVMe JBOD

Chris Petersen

HARDWARE SYSTEMS ENGINEER, FACEBOOK


Mike Yan

HARDWARE ENGINEER, FACEBOOK

Clark Shao

HARDWARE ENGINEER, FACEBOOK

Why NVMe?



It scales!




It's open
source!



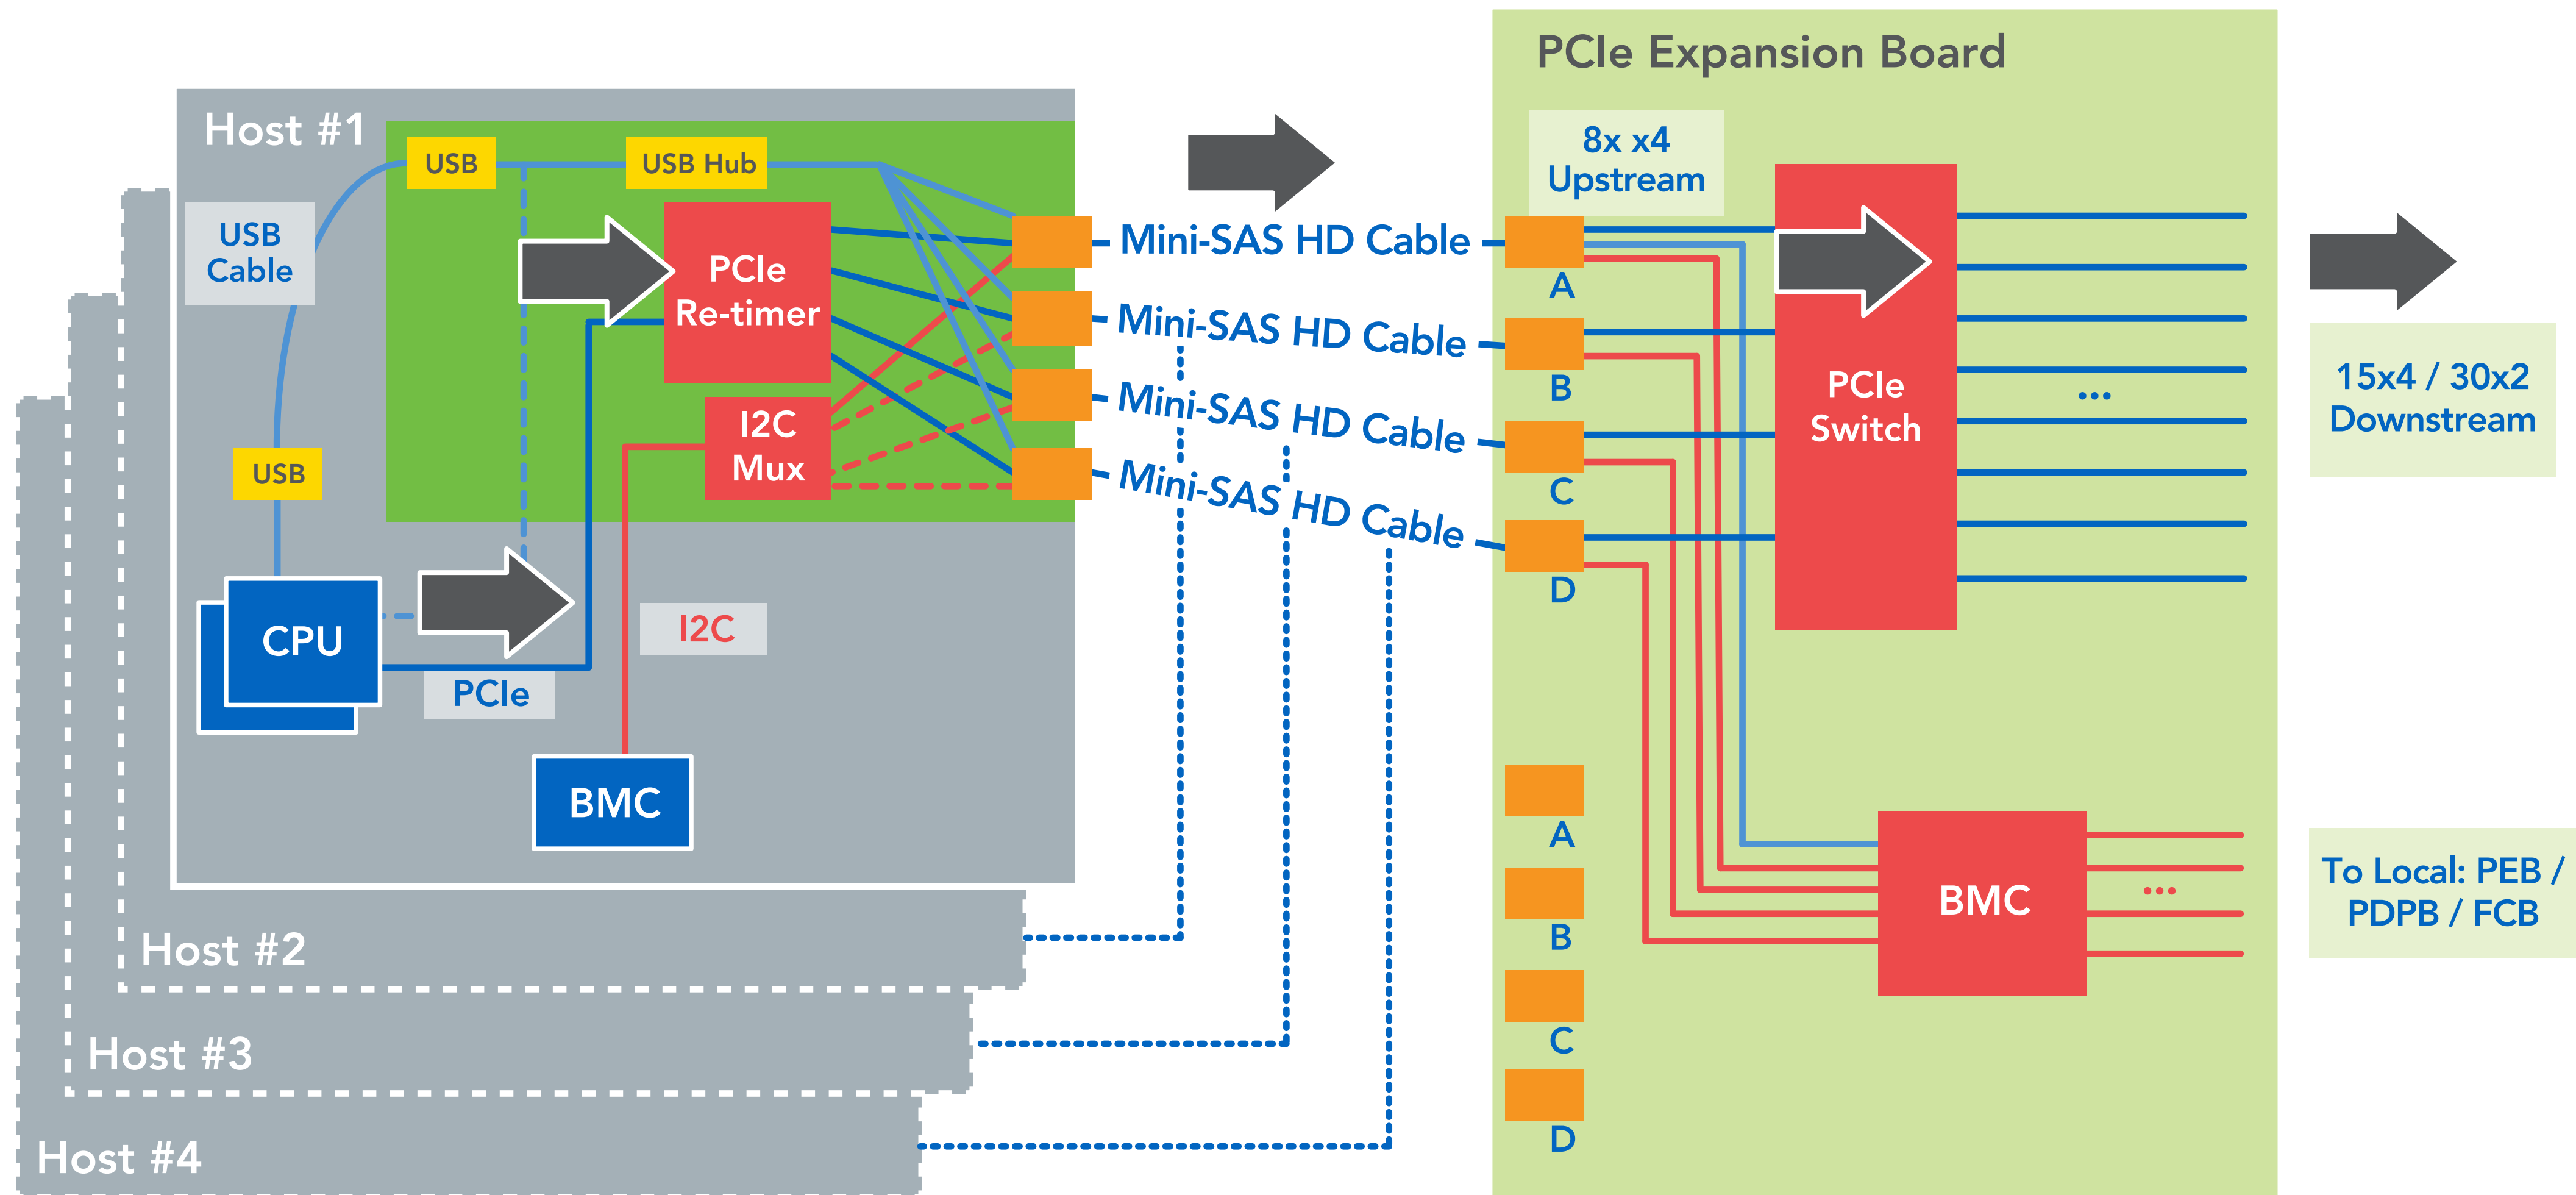
Multiple form
factors!

Design objectives

- Modular
- Flexible configurations
- Compatible with upcoming standards

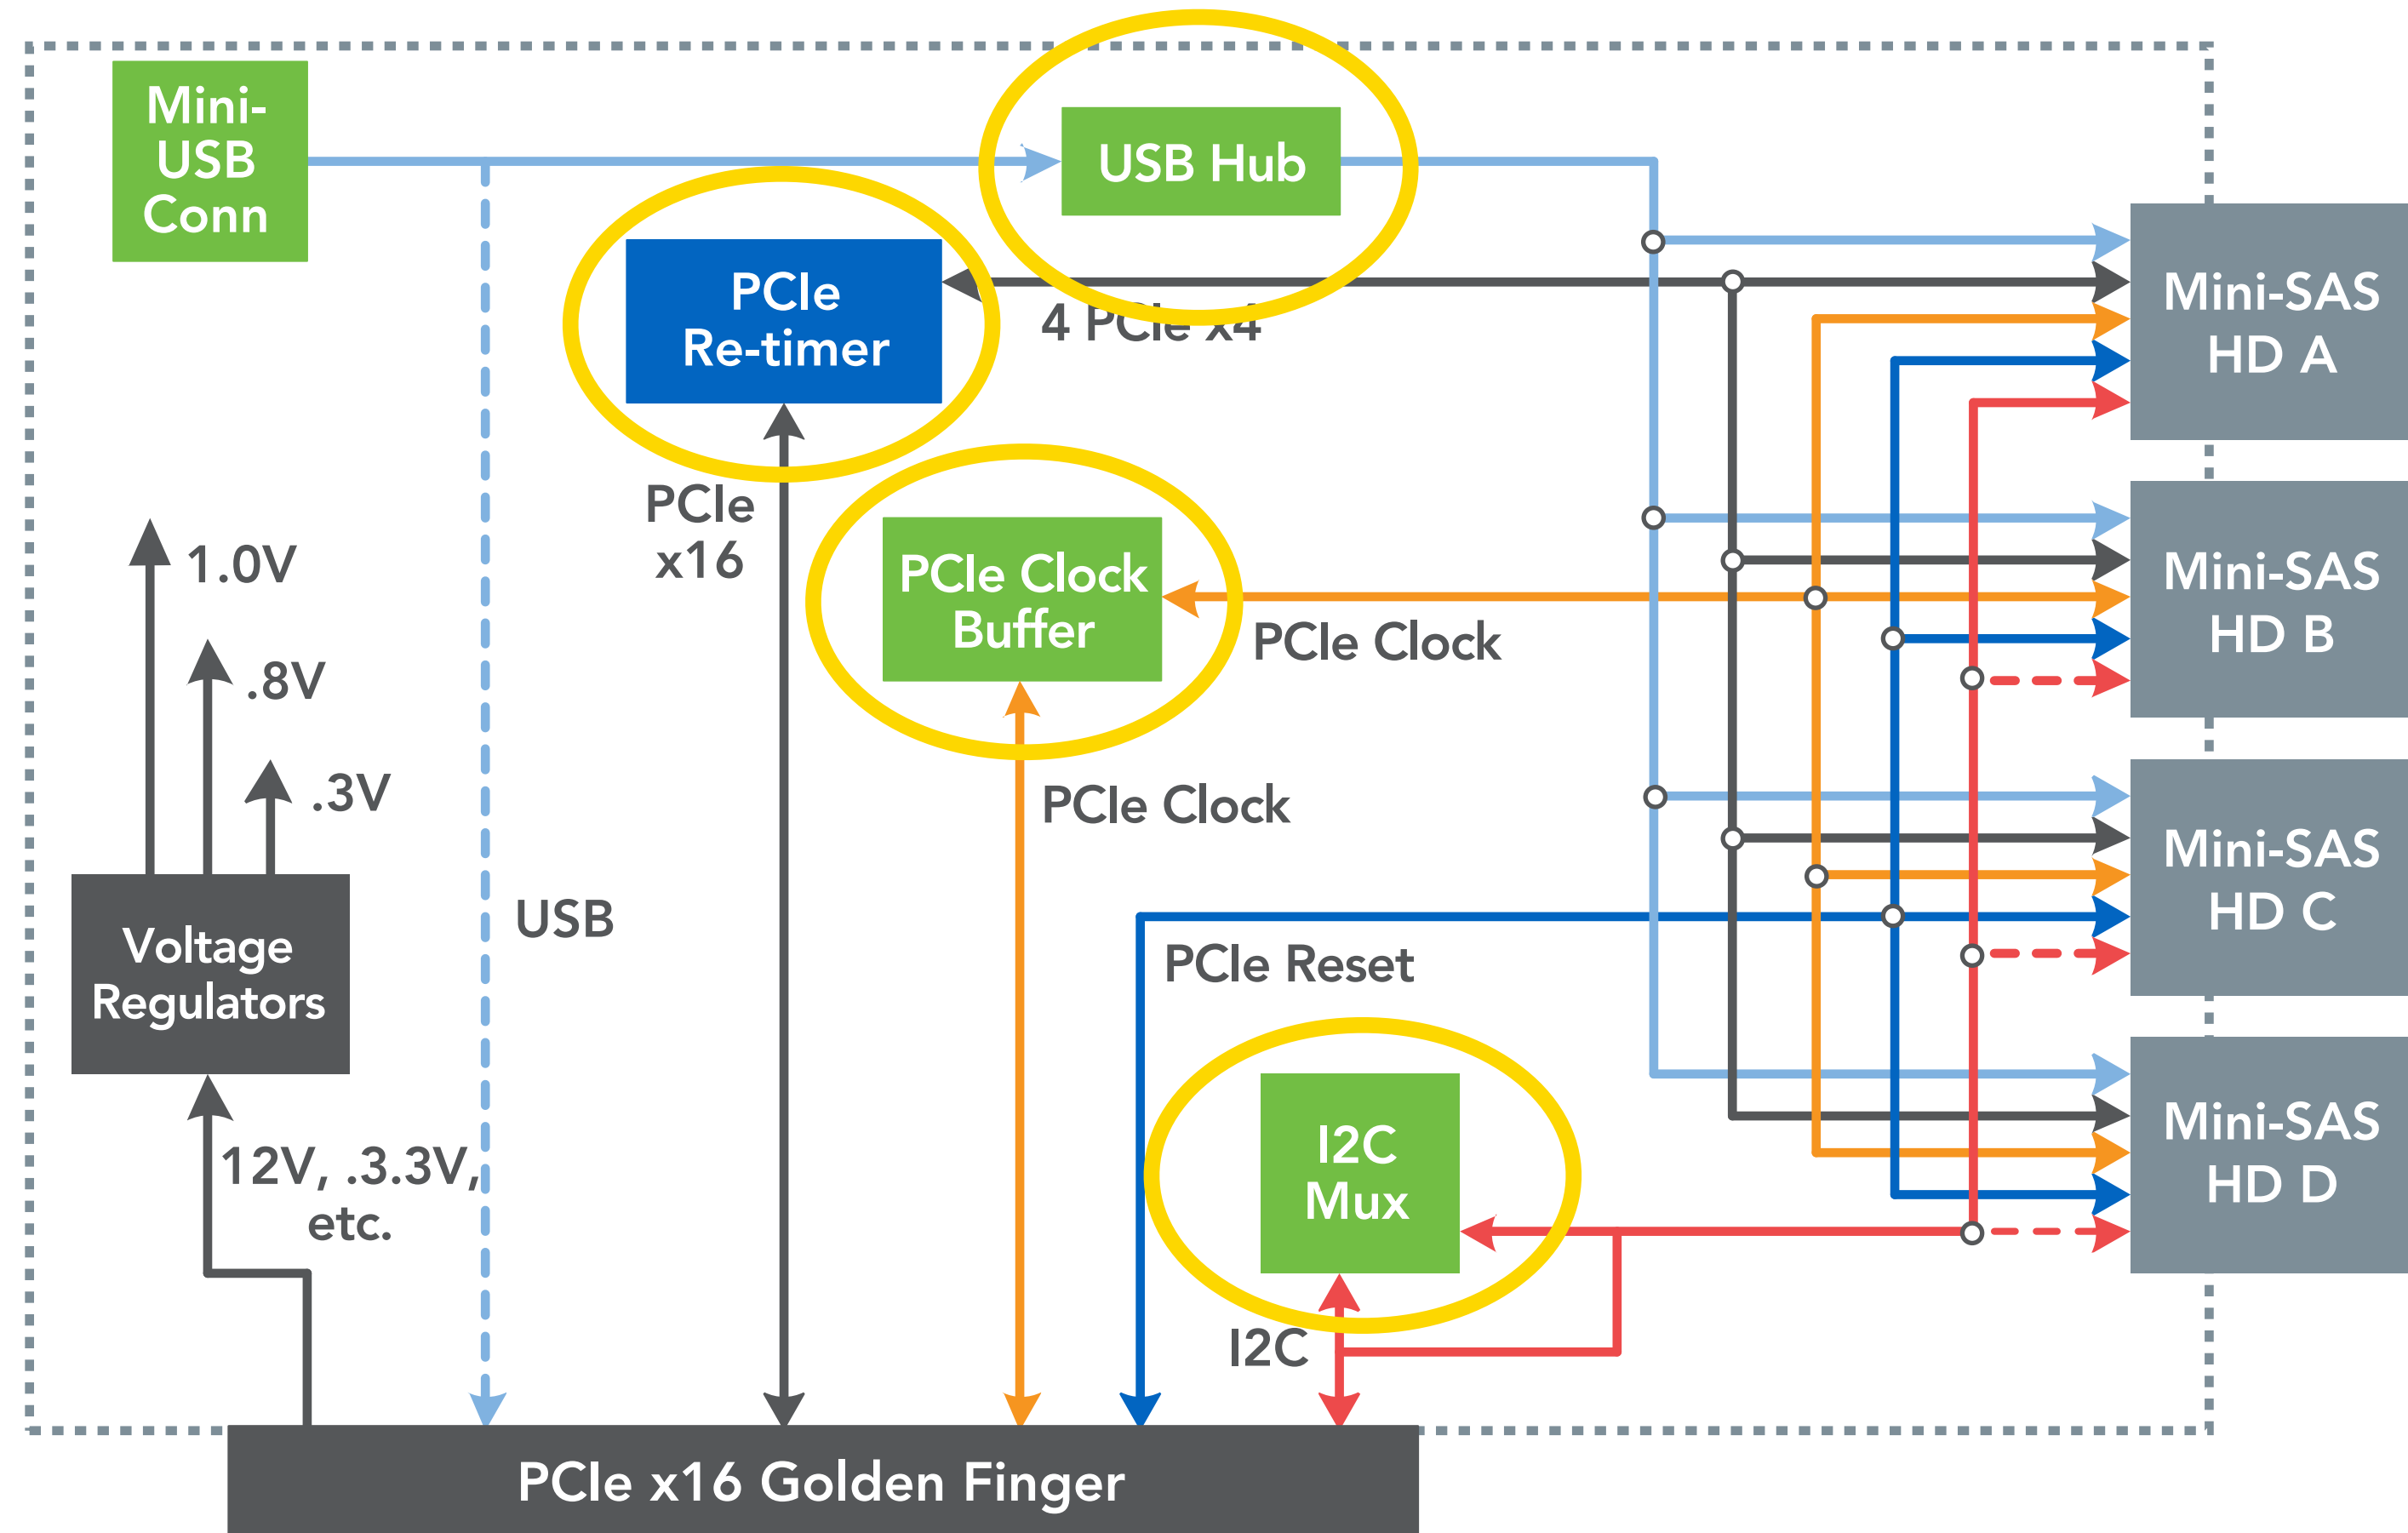


Lightning overview


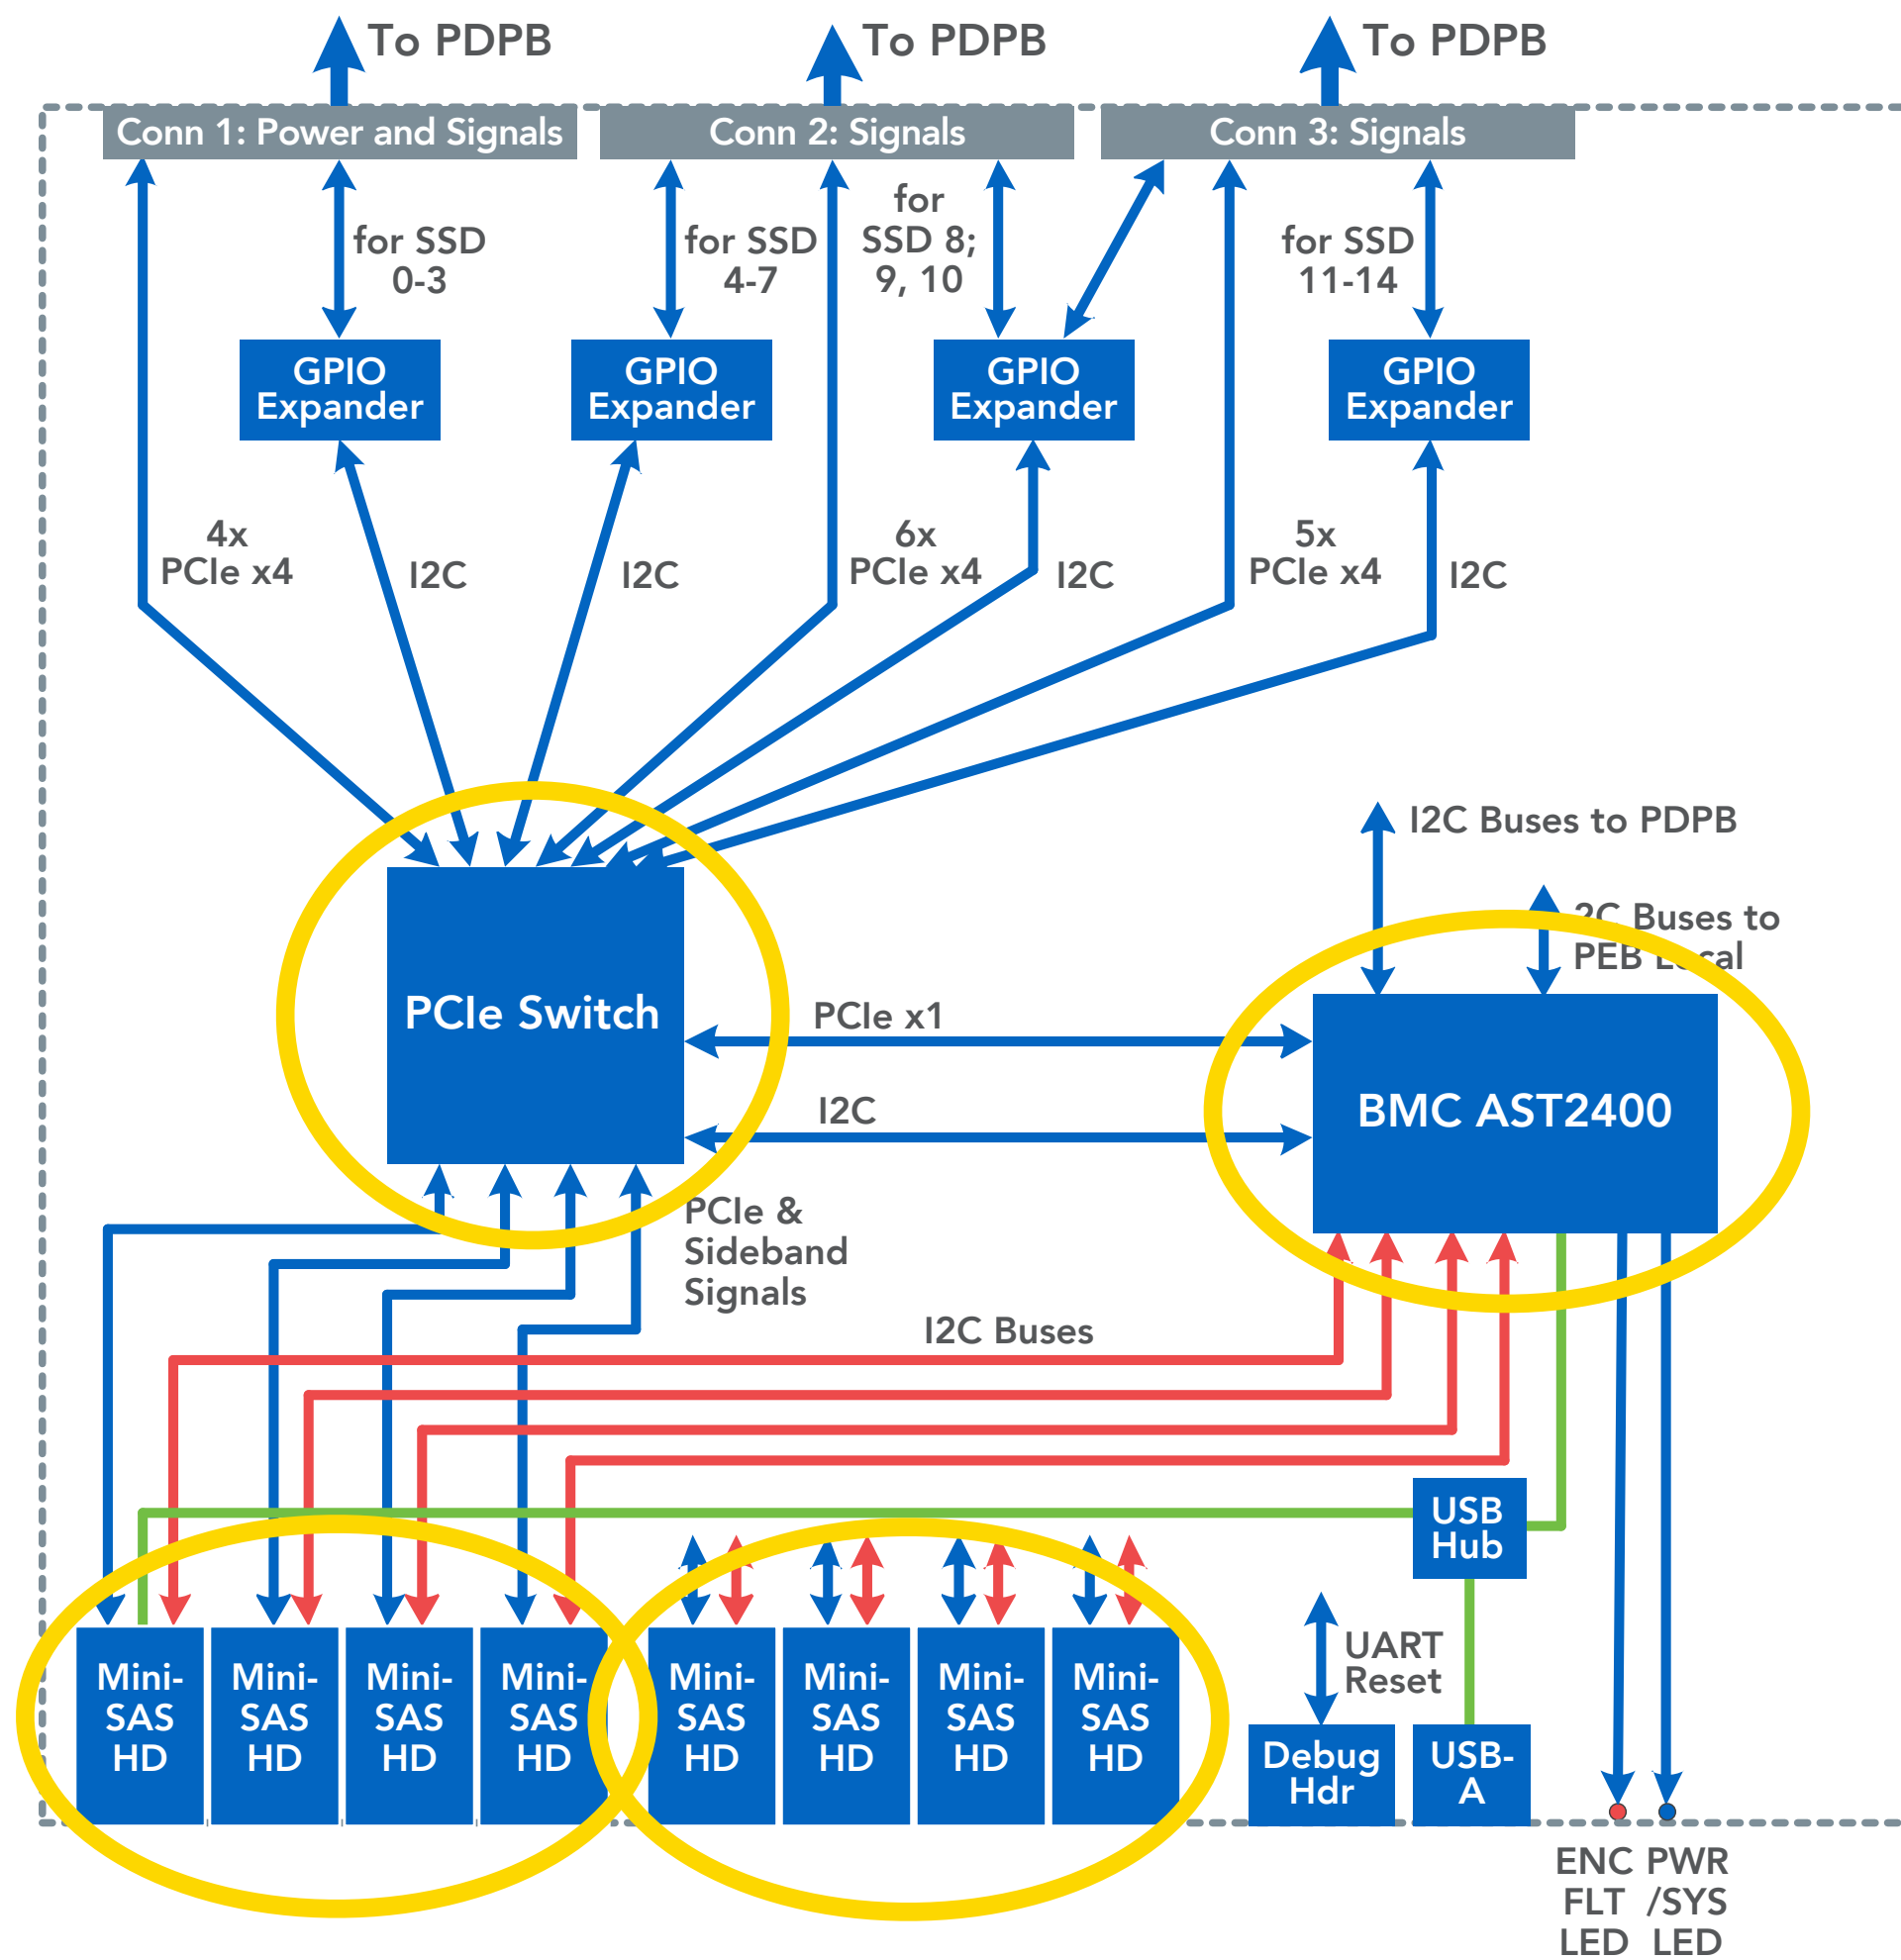


Re-timer card

- PCIe re-timer
- USB hub
- Clock buffer
- I2C Mux




PCIe Expansion Board (PEB)



Head node connections

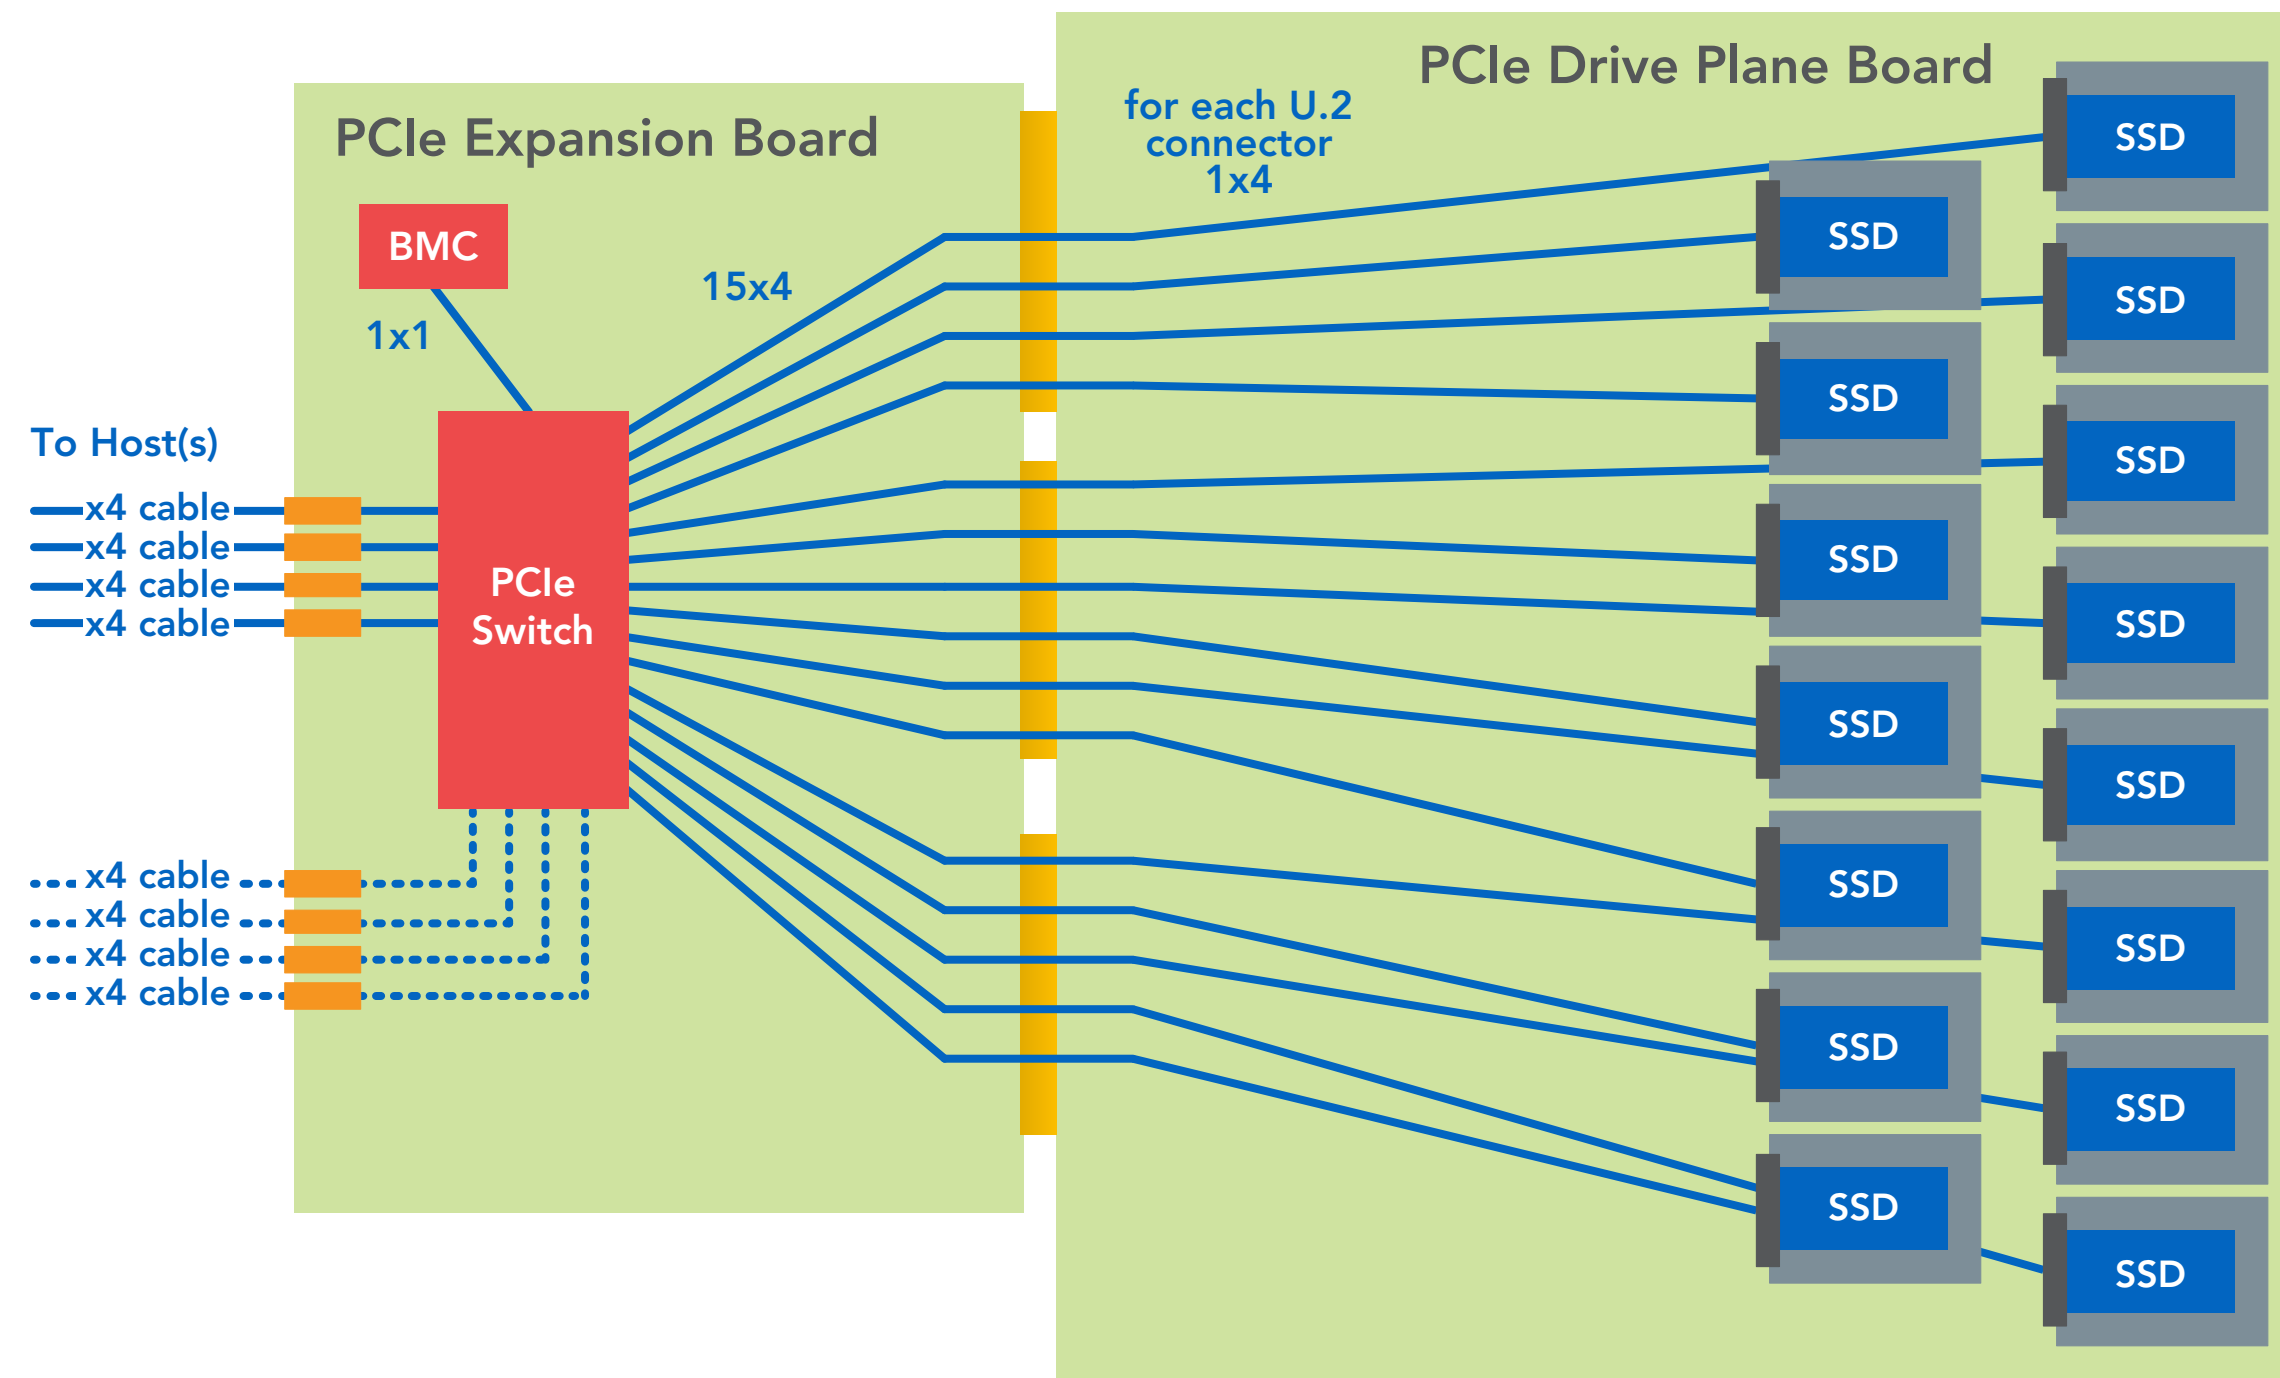
- Dual x4 Mini-SAS HD cables
- Pin-out is compatible with PCI-SIG external PCIe cable specification
- Custom cable to include full side-band



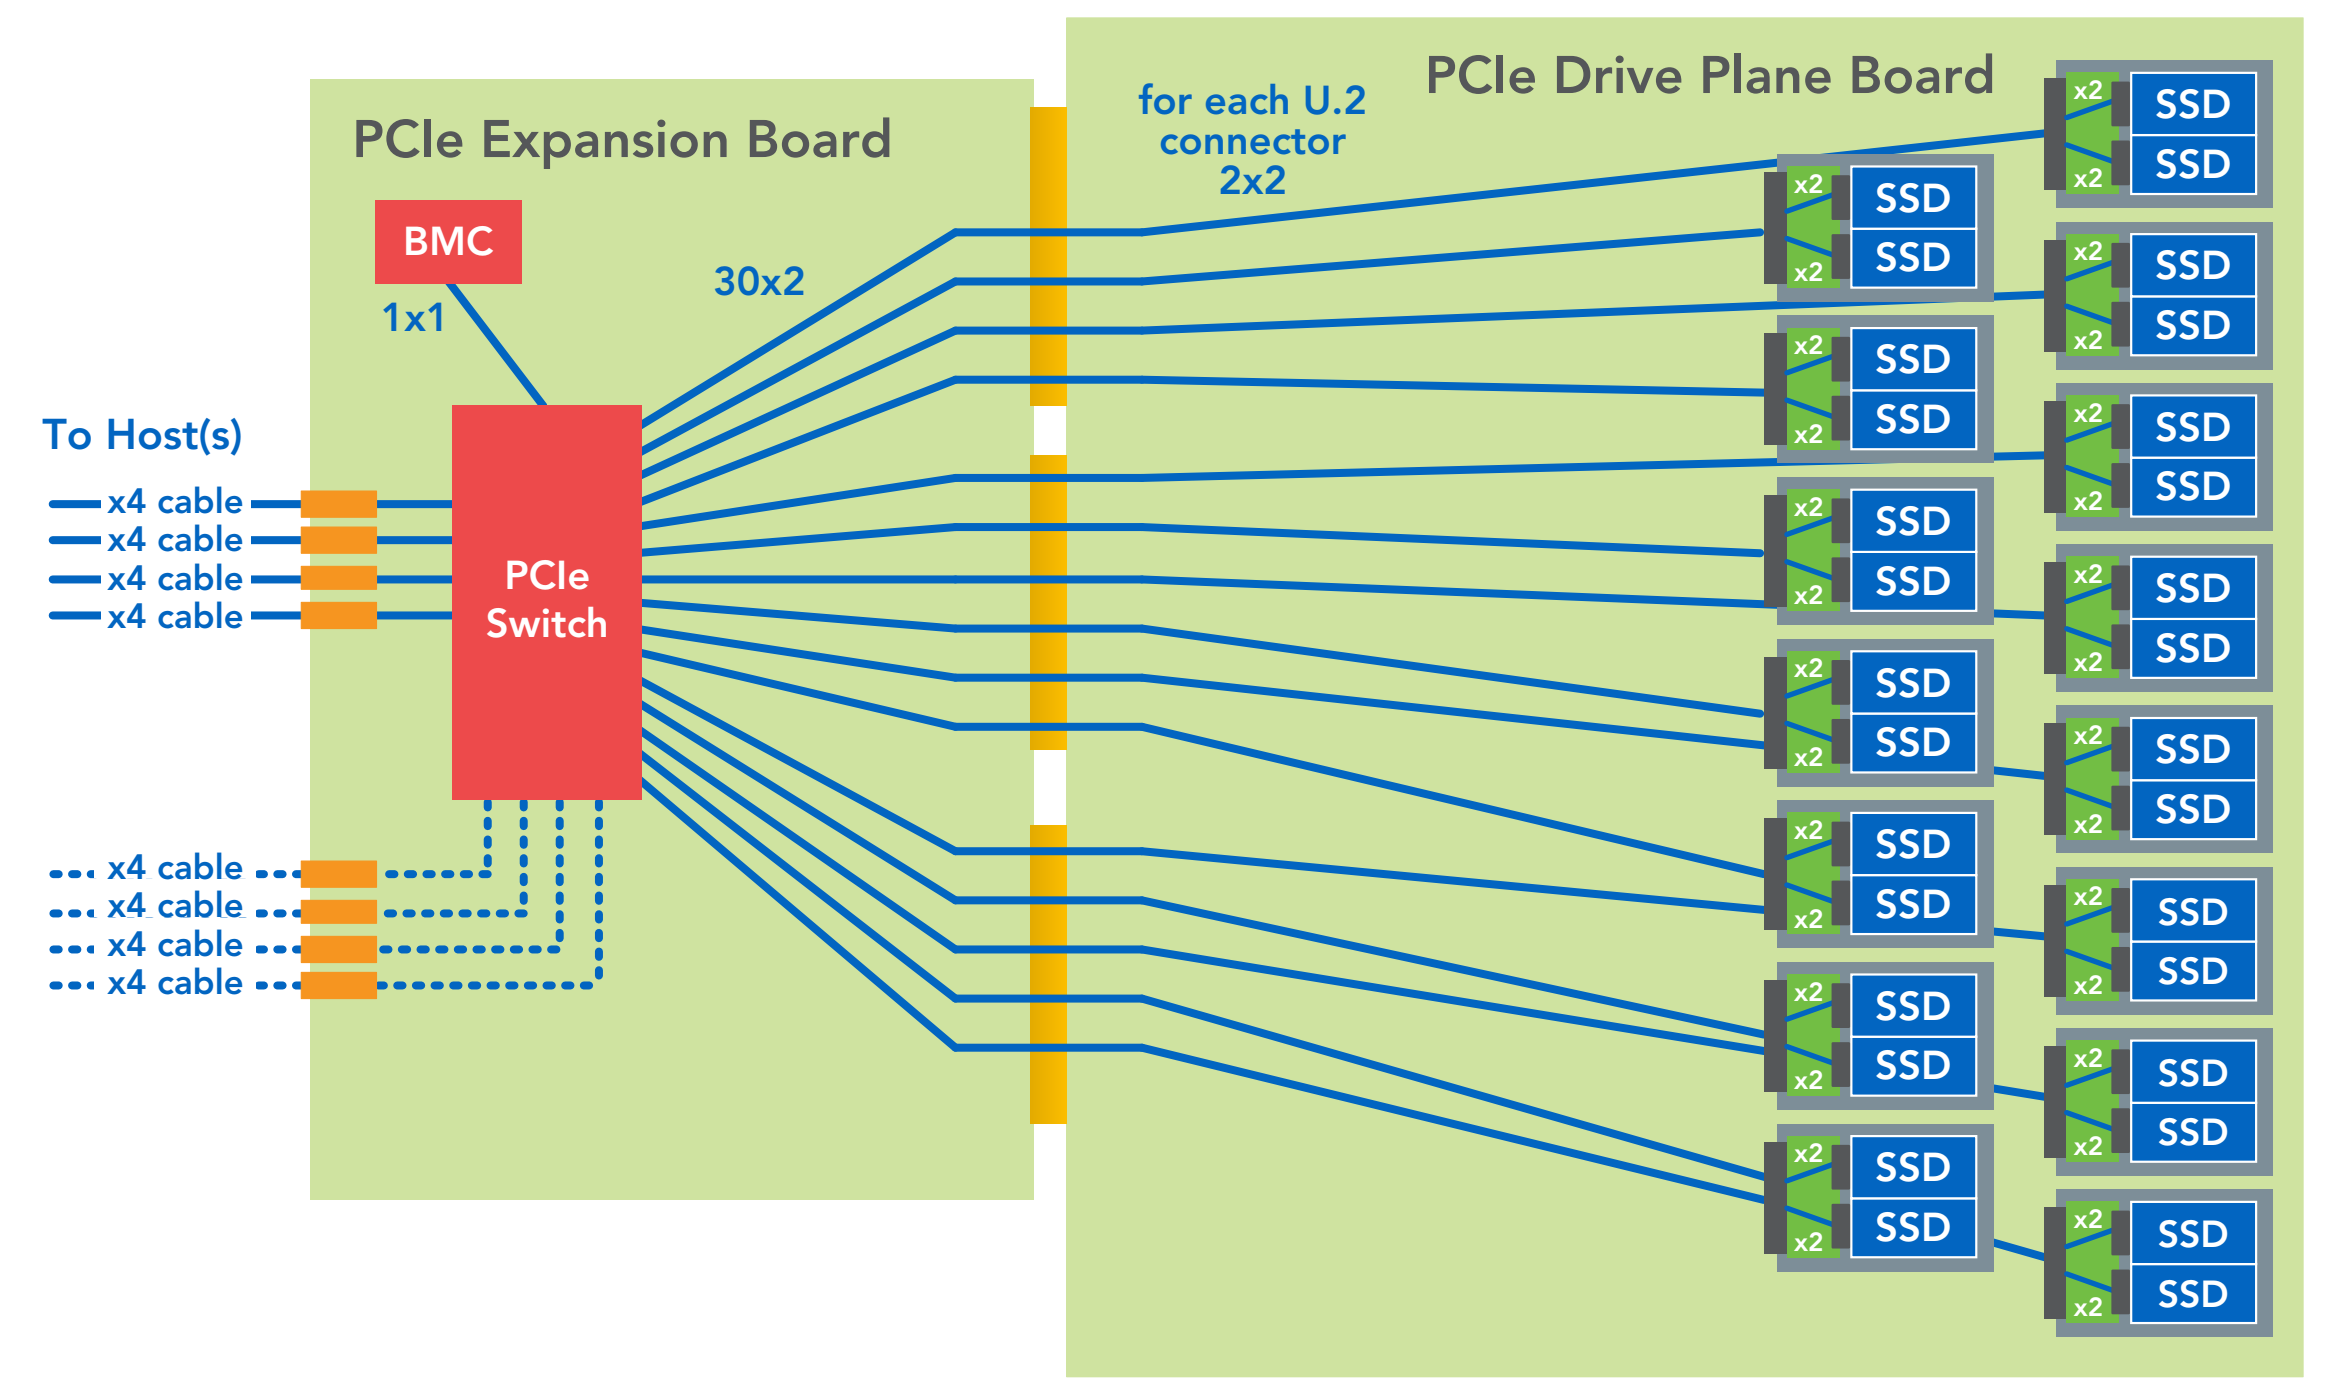
* - P1-Z WIRING IS THE SAME AS P1-Y

* - P2-Z WIRING IS THE SAME AS P2-Y

PCIe switch configurations



15 x4 SSDs



30 x2 SSDs

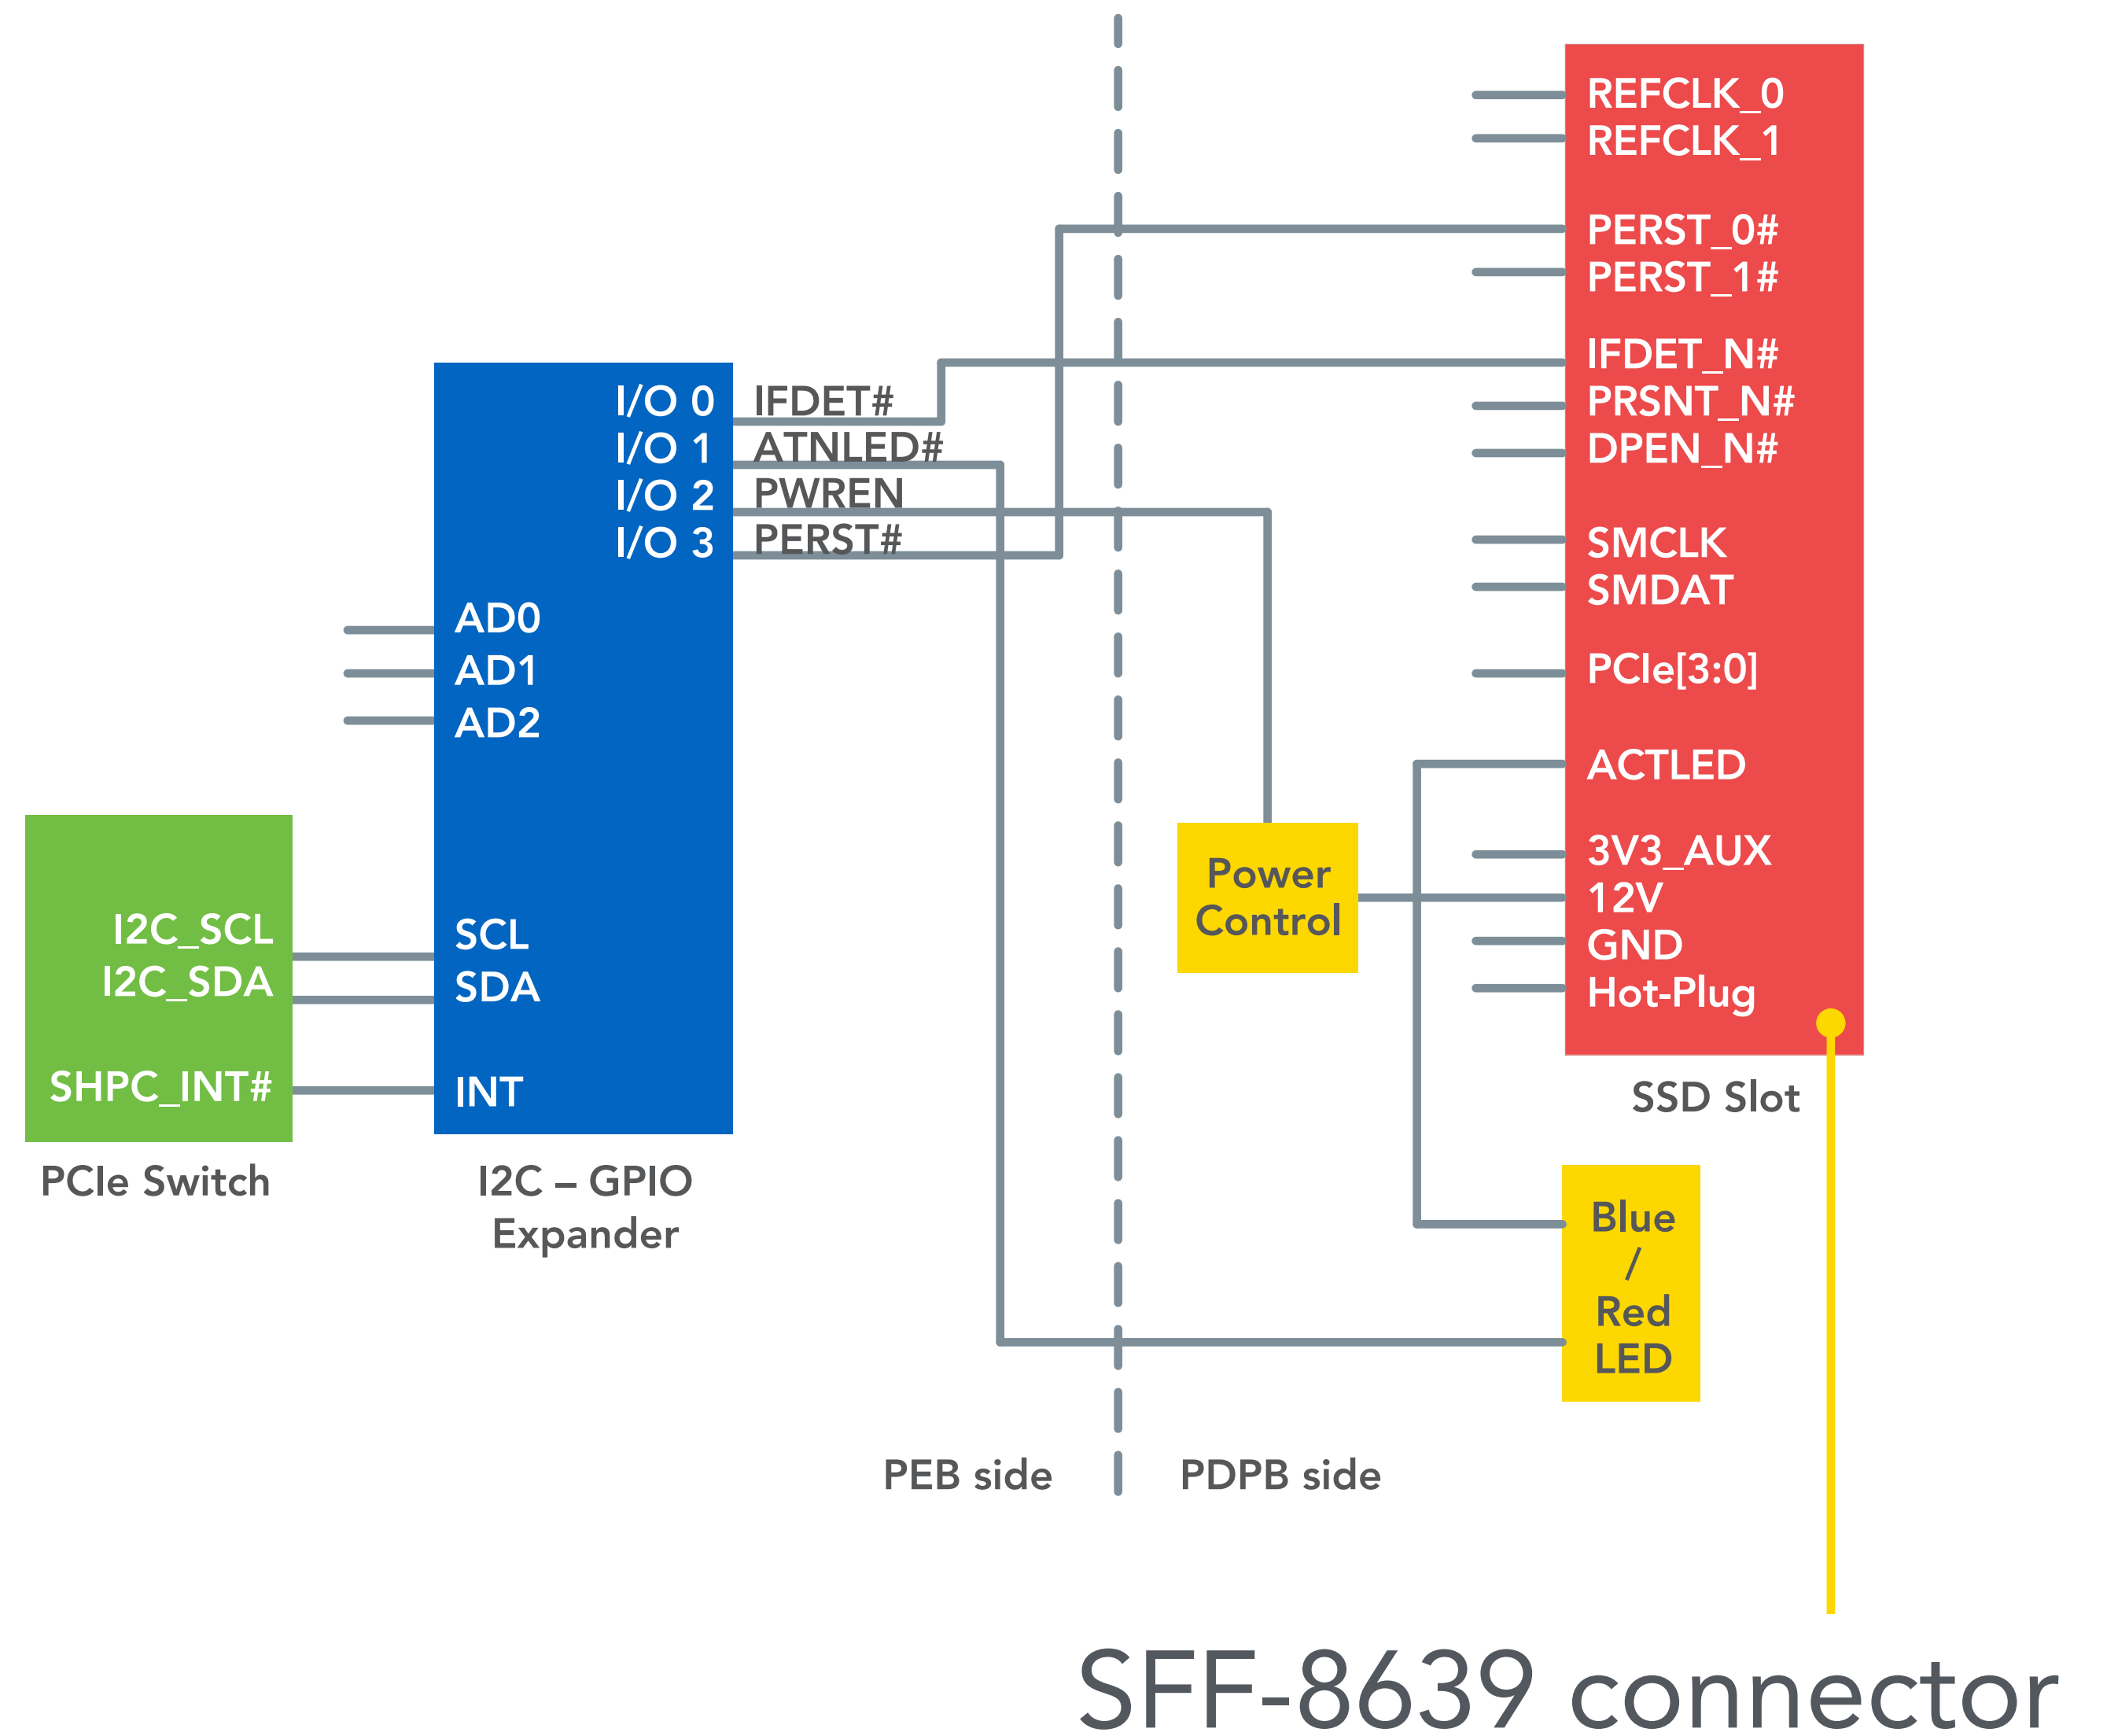
PCIe sideband

→ Present

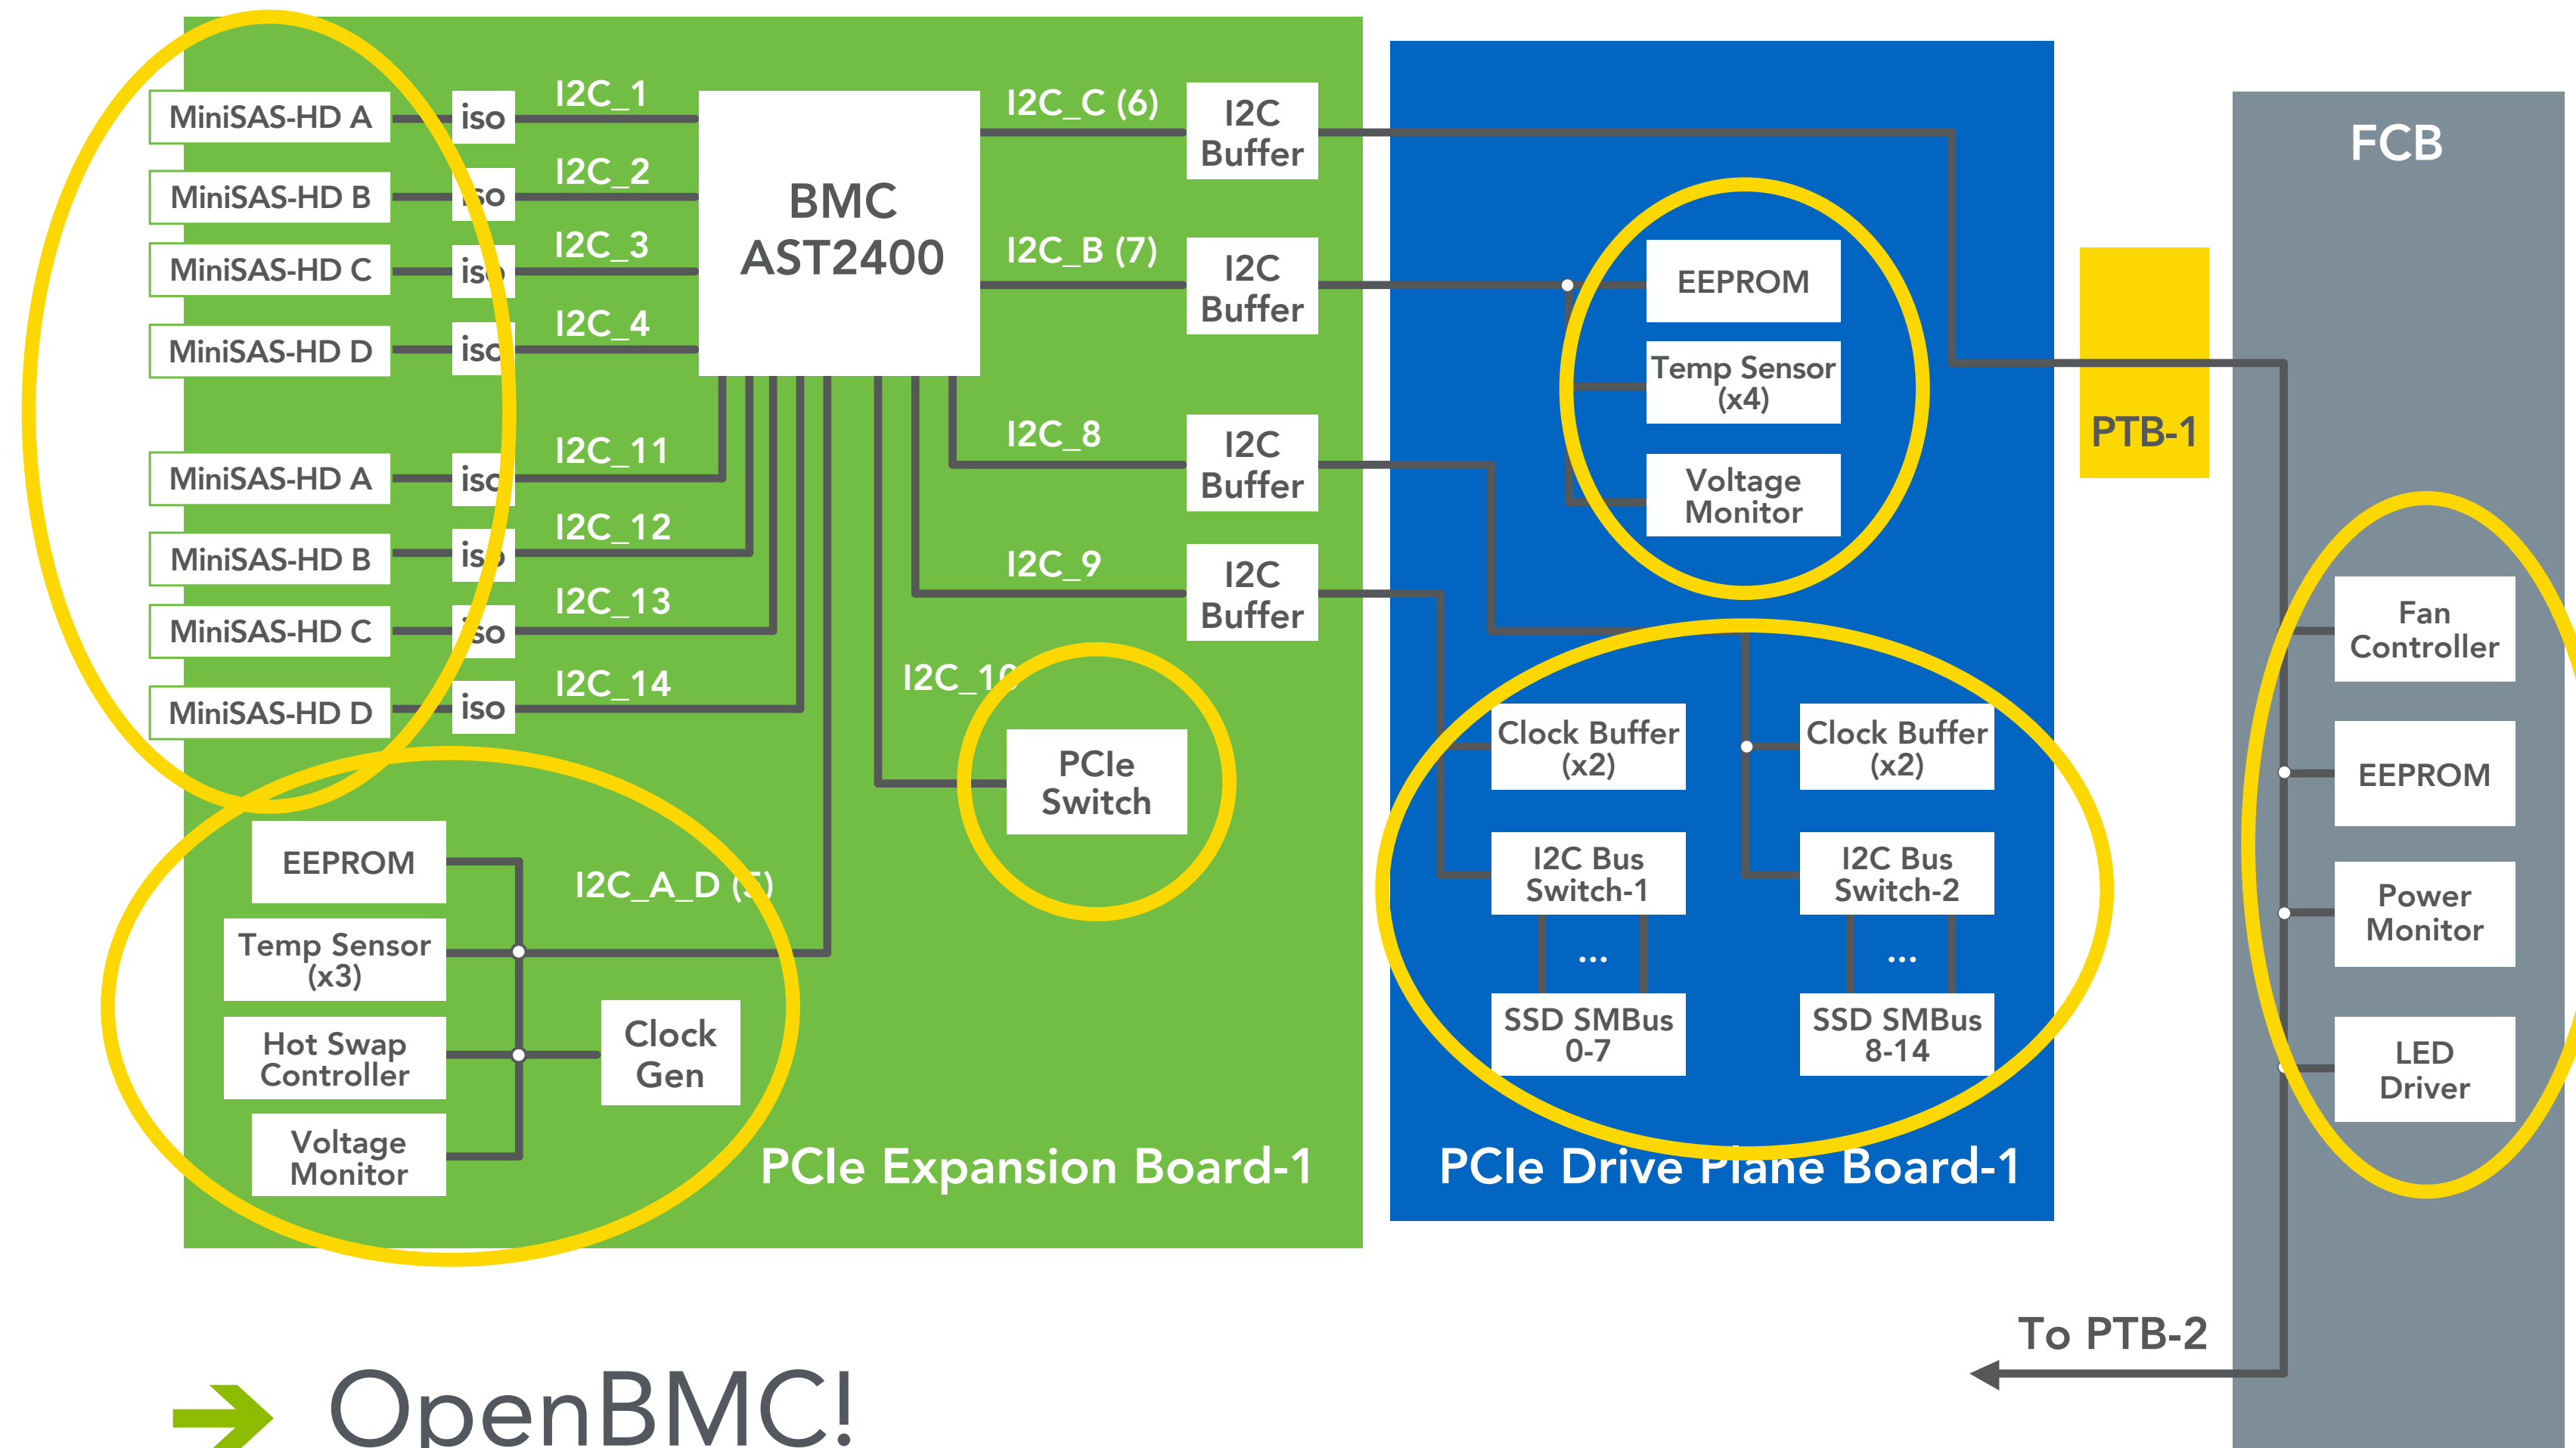
→ Attention LED control

→ Slot power control

→ PCIe reset

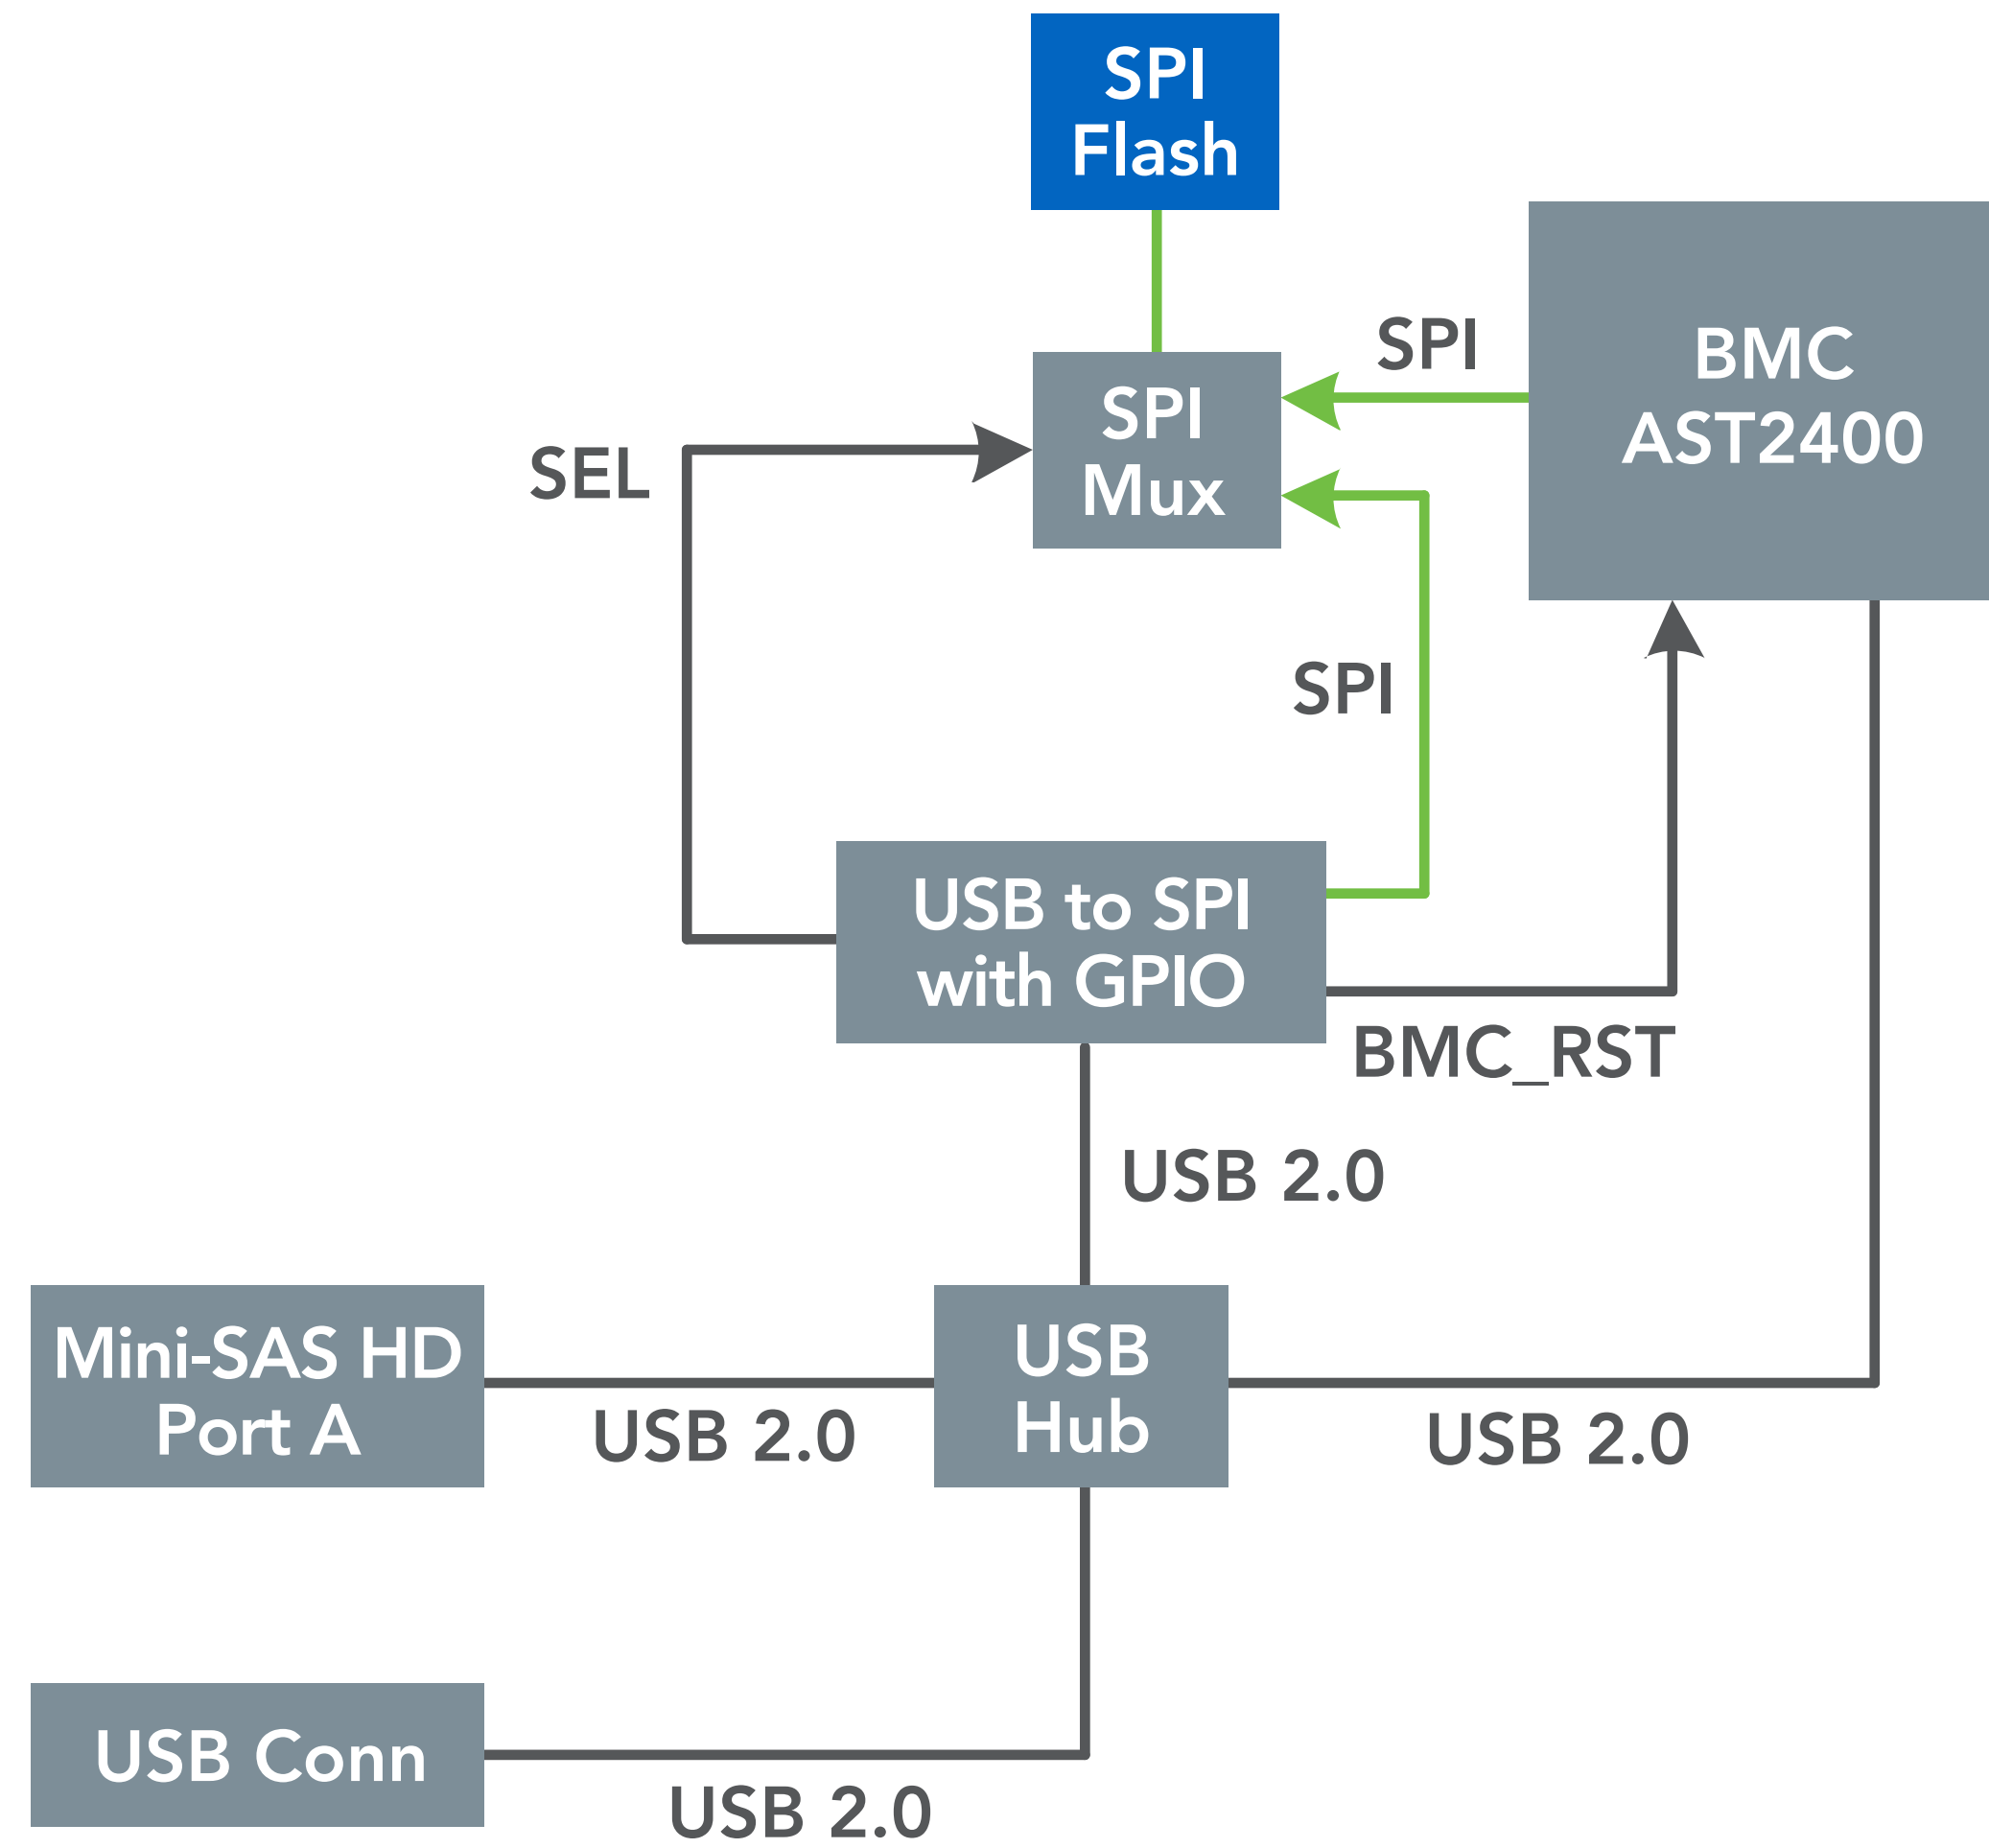


BMC

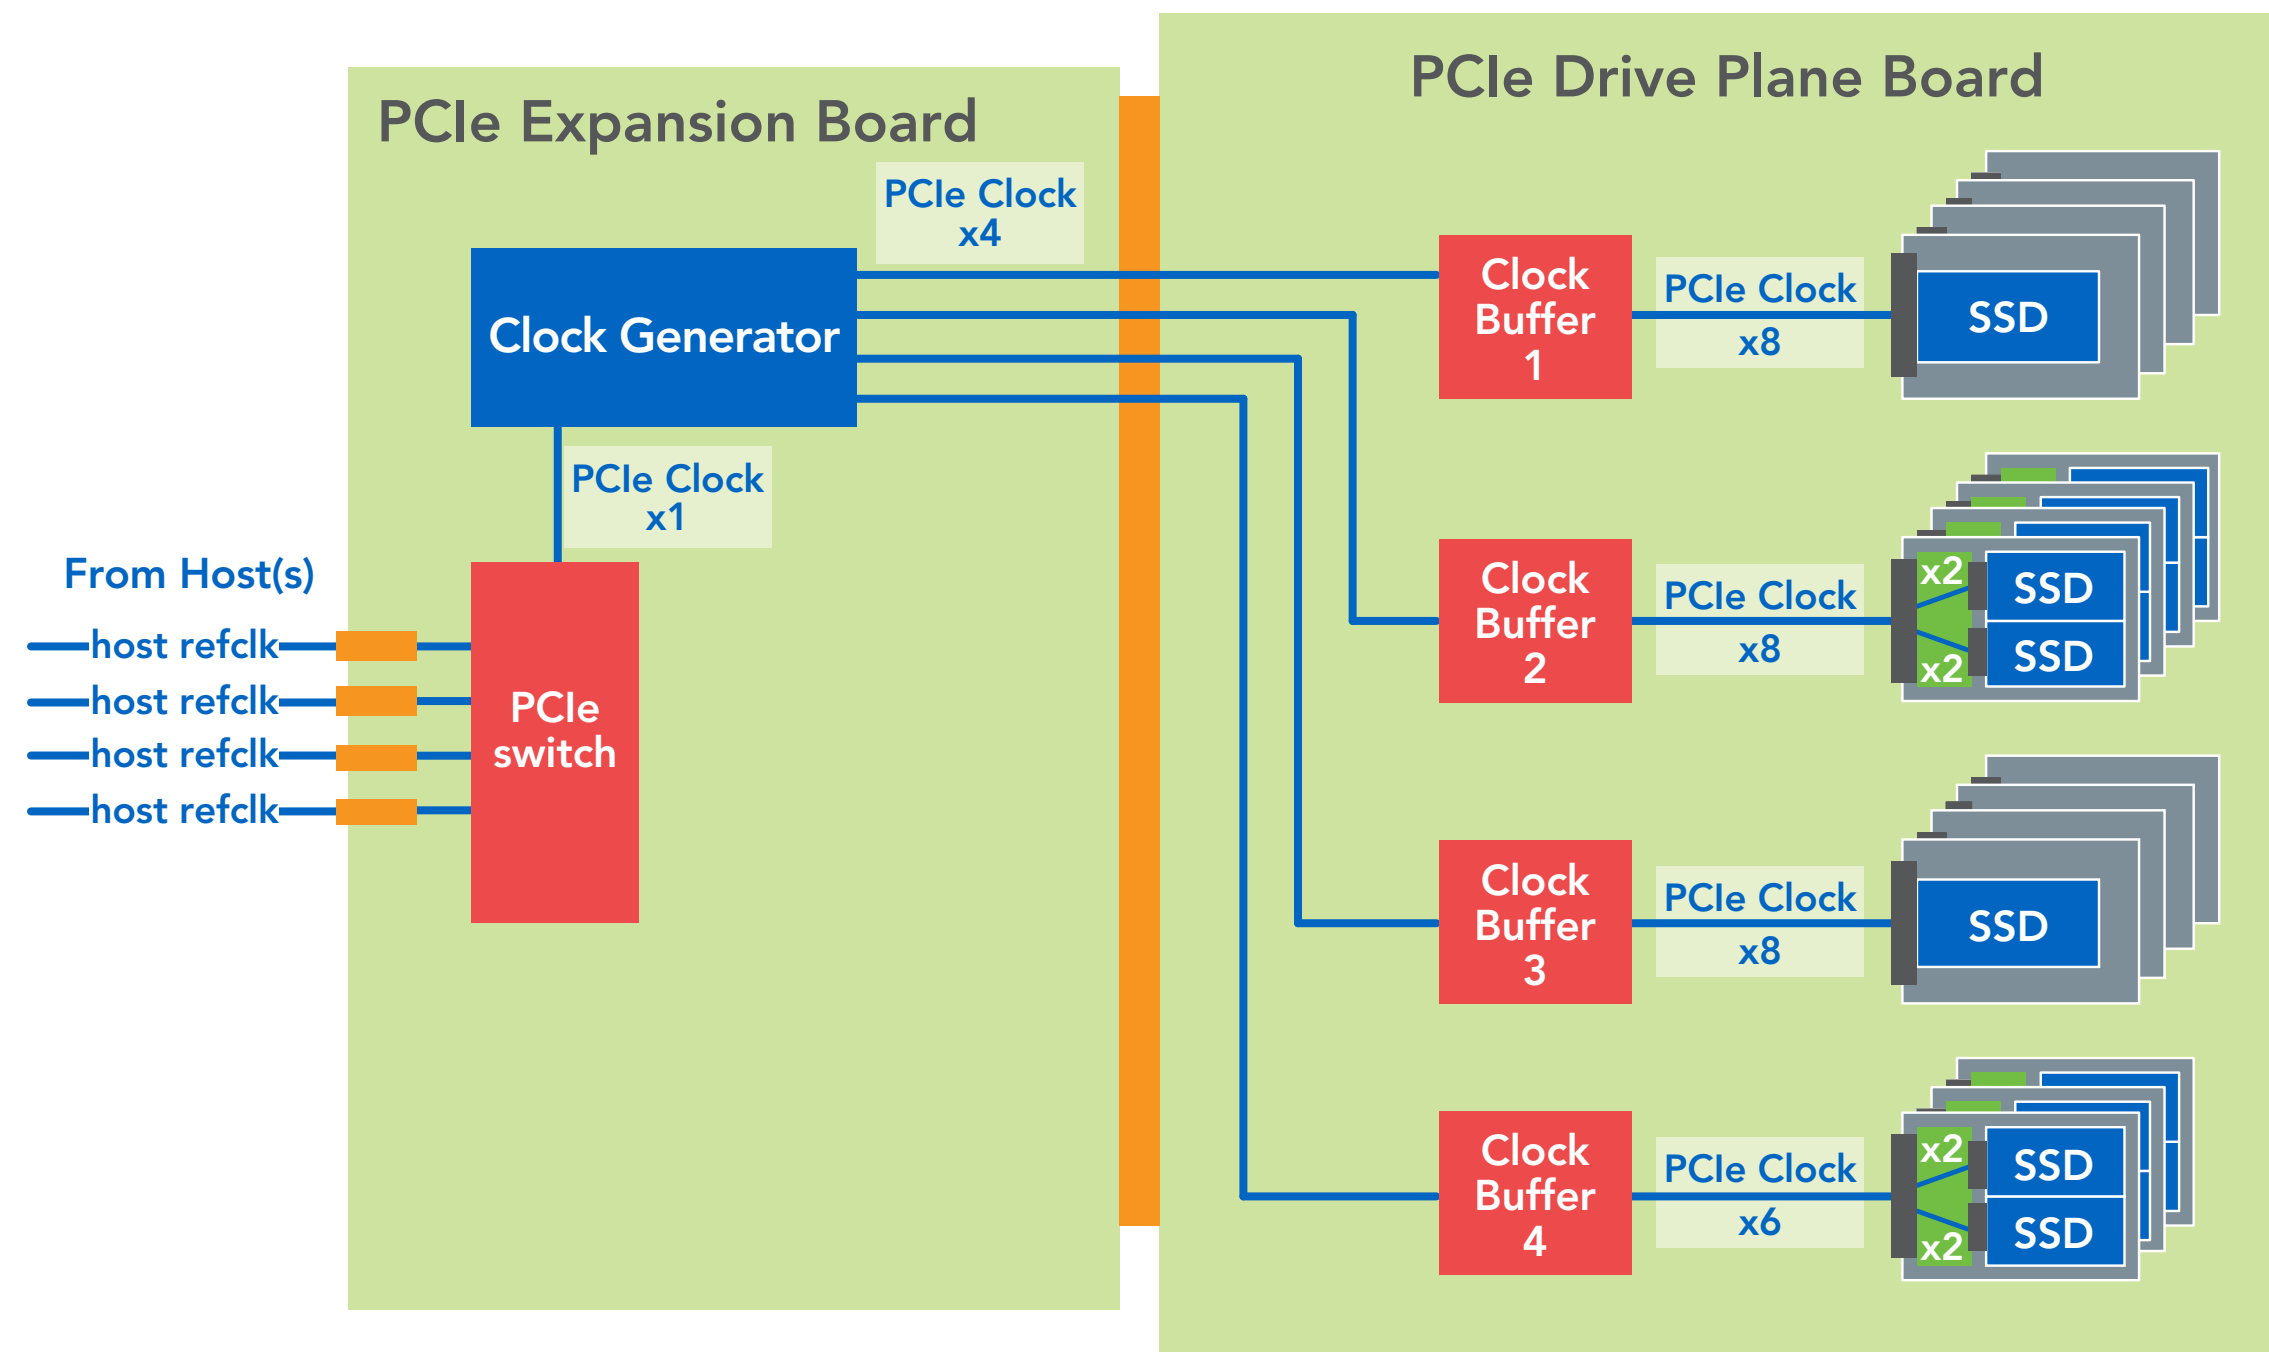


→ OpenBMC!

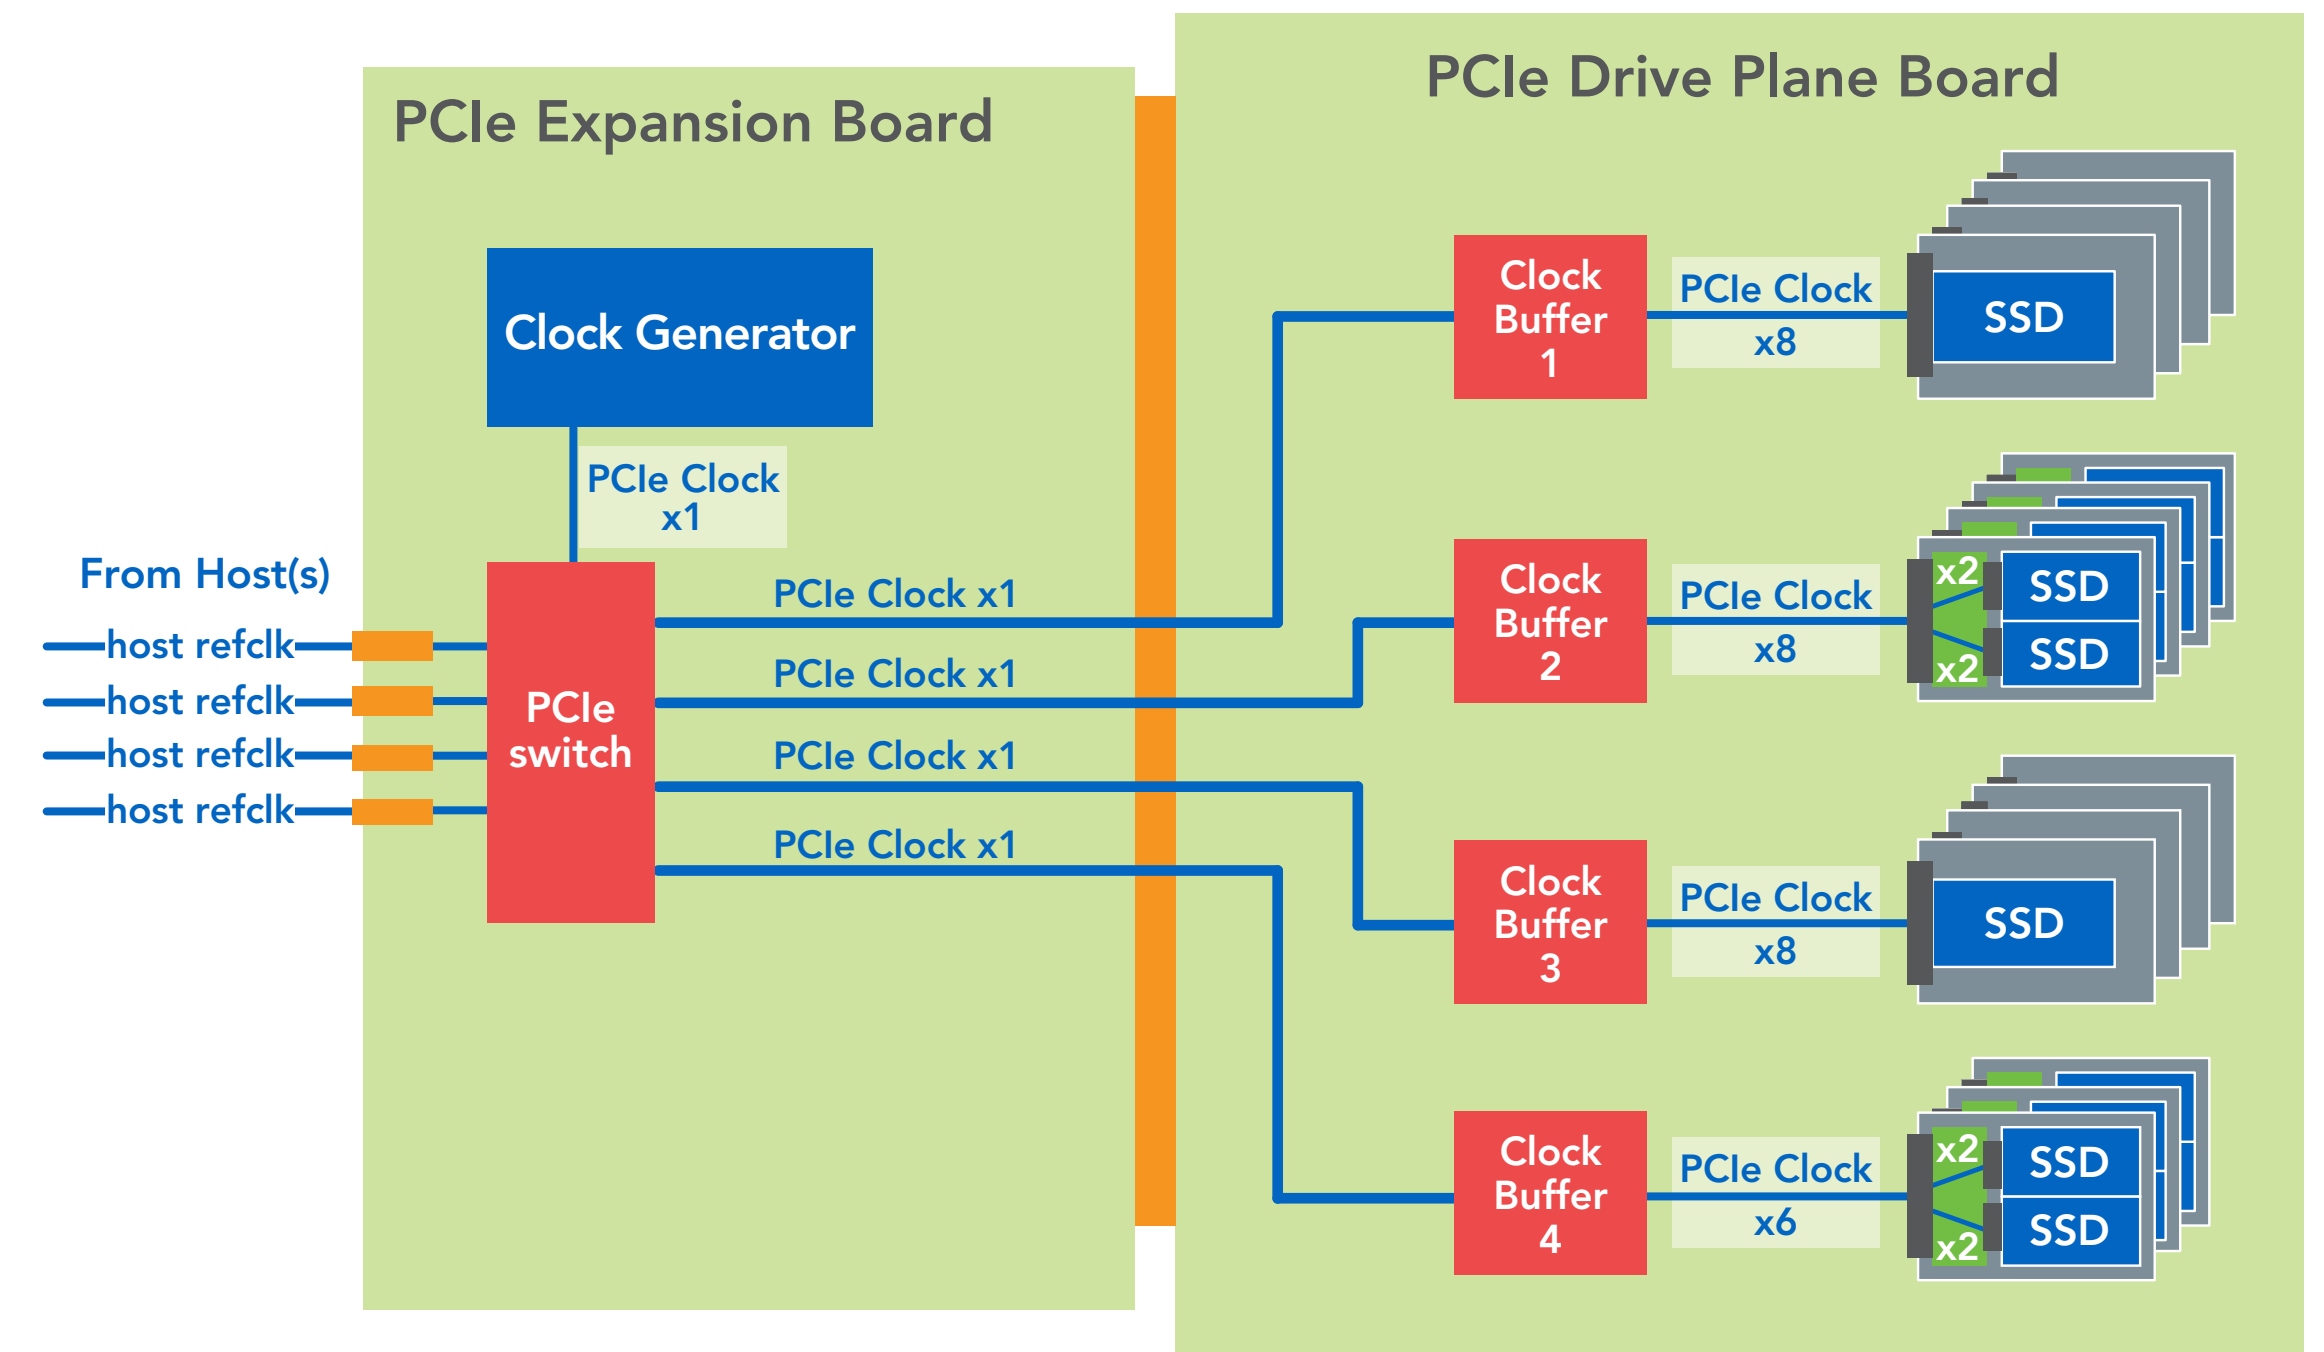
USB



Clock topology



Option 1



Option 2

SSD options



2.5" – 15mm



2.5" – 7mm



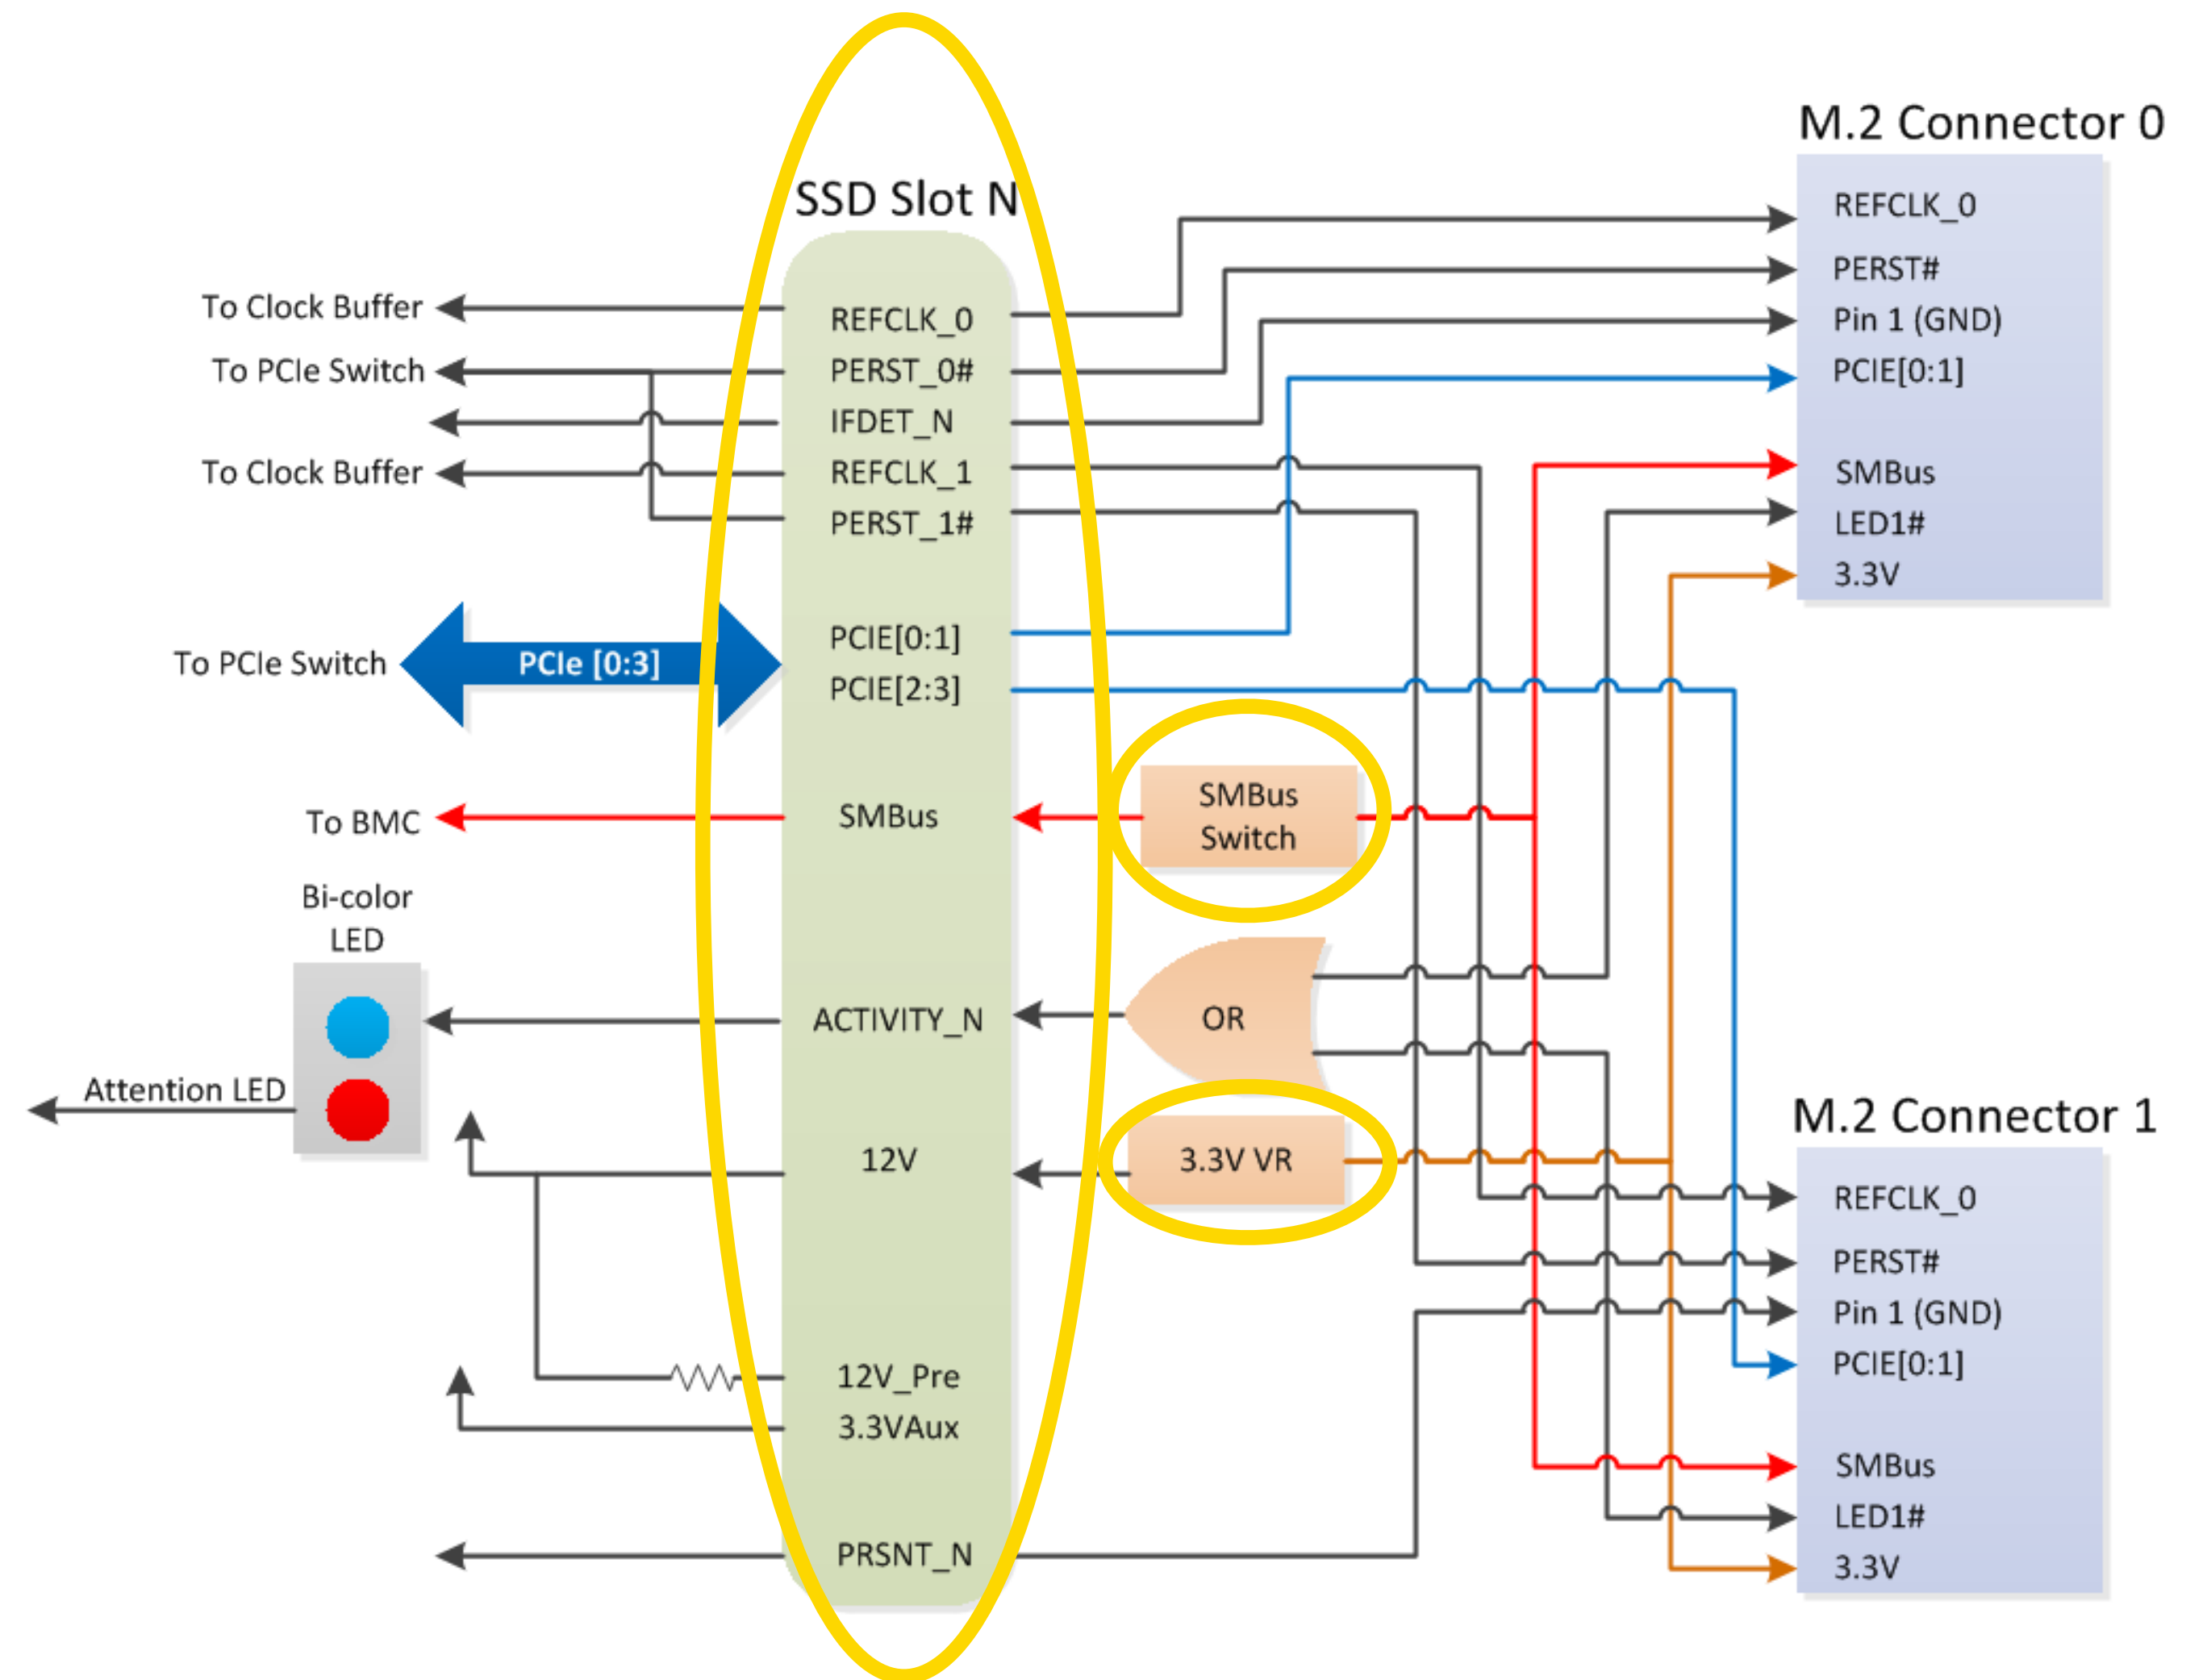
M.2 22110 or 2280

M.2 Carrier

→ SFF-8639 connector


→ 3.3V VR

→ SMBus Mux




Power


- SSD "slot" = 14W total
- PEB power \leq 130W
- System power \leq 770W



Thermal




Temperature(°C)




More information

- See the OCP spec in the server working group!
- Come see the systems in the Facebook, Wiyynn or Intel booths!

**A special THANK YOU to the Facebook,
Intel, and Wiwynn teams!**





OPEN

Compute Project

