Non-Industrial Use of P802.3cg

Jon Lewis: Dell EMC

Berlin, Germany

12 July, 2017

Current Server Architecture

- Device to device communication in a server uses a variety of bus technologies typically connected using multiple topologies.
  - SMBus/MDIO (or I2C)
  - SPI
  - UART
- These legacy busses while having performed well for many years are a design challenge for baseboard management controllers (BMCs) especially from a SW perspective.
  - Many of the devices while claiming “compliance” require special drivers to operate properly

Block Diagram – Current Architecture

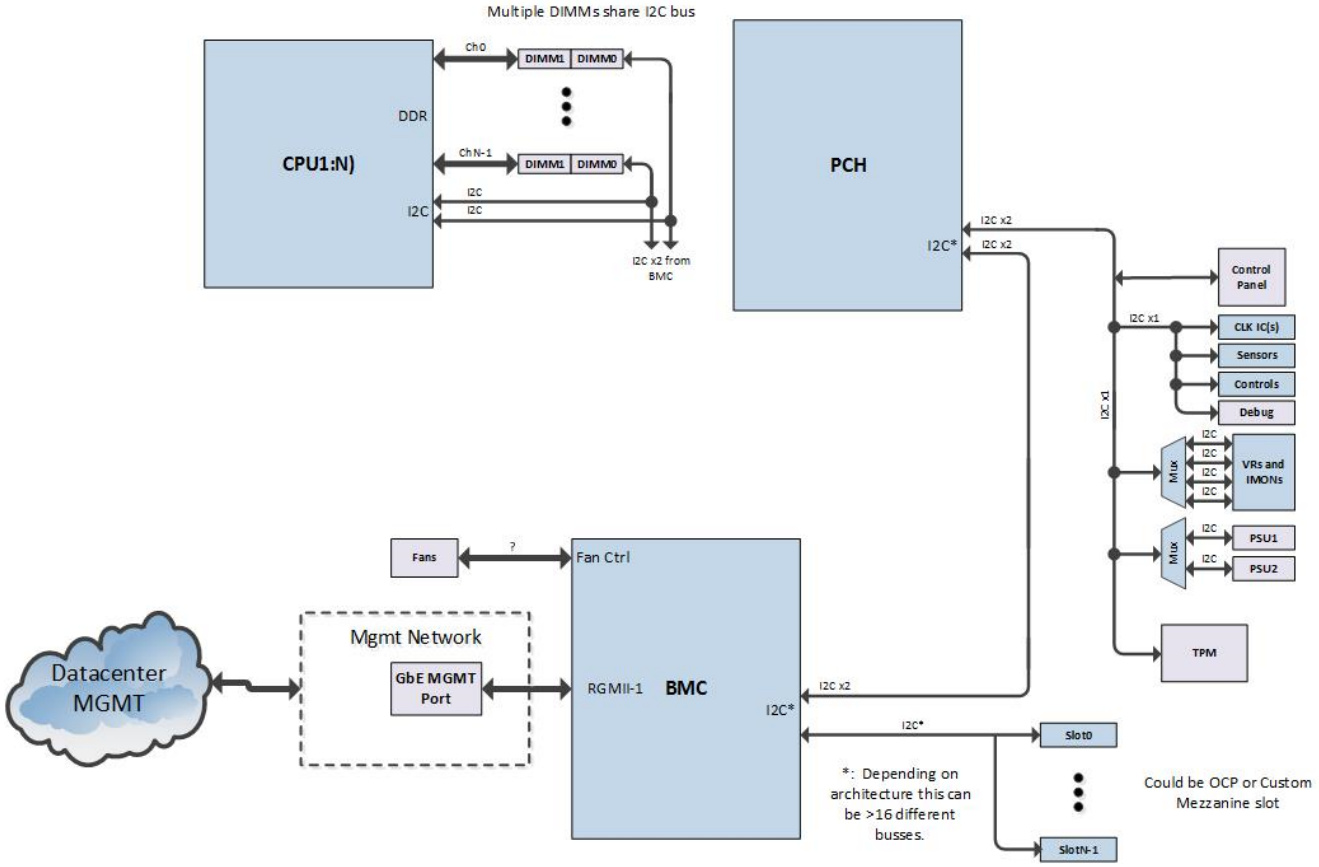

Proposed new architecture (Near Term)

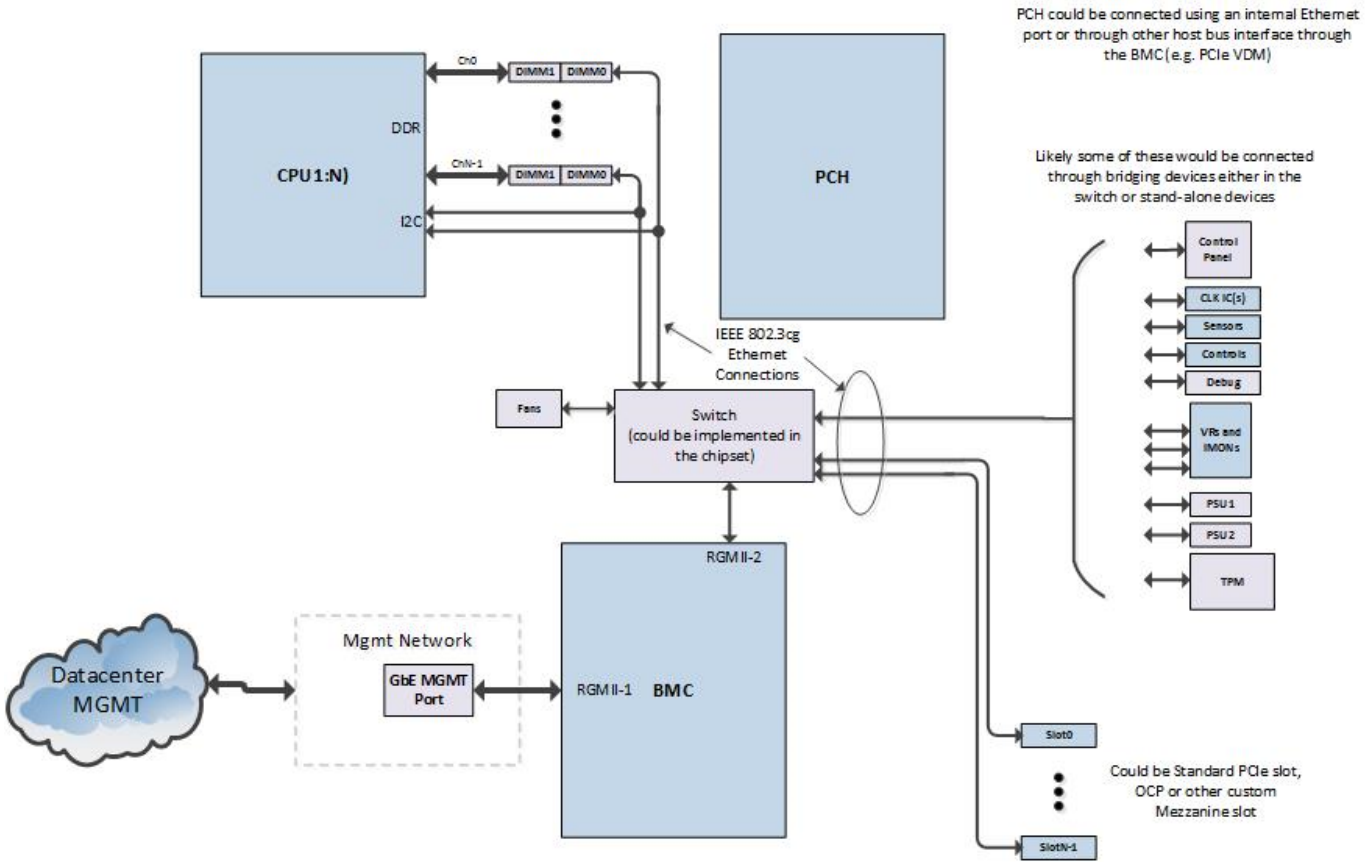

Future new architecture (Long Term)

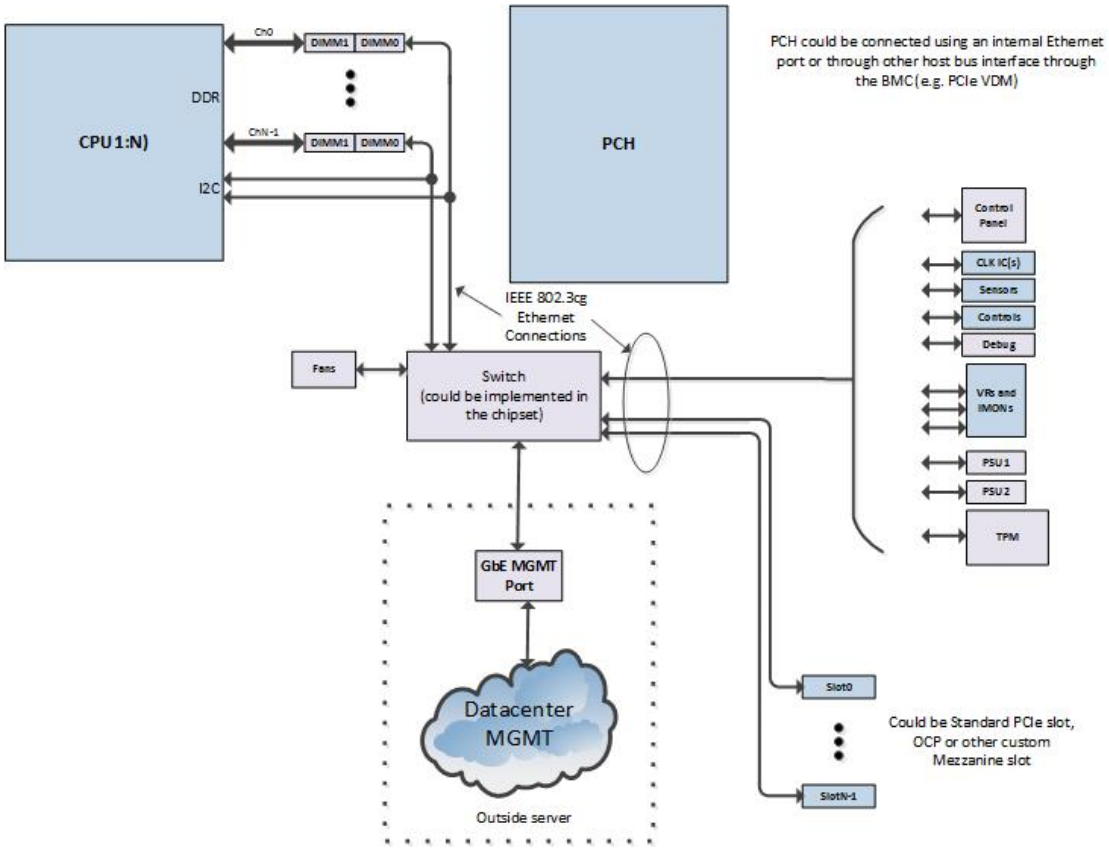

Why 10 Mbps Single Twisted Pair Ethernet?

- Same number of pins as SMBus which would maintain current PCIe standard pin count while adding network functionality.
- Using Ethernet allows for discovery of devices using a common BMC driver.
  - Going from >10 custom drivers to a standard Ethernet driver reduces complexity on the BMC coding and will greatly reduce validation time now required for all custom implementations.
- As the compute node and networking “converge” there is a fine line between what is in a traditional server and what is in the networking “Ether”
  - External PCIe Expanders
  - Chassis servers
  - Modular server implementations. Ethernet provides a standard ubiquitous management communication path
- Easier for automated alerts than multi-master SMBus.
  - Using SMBus multiplexers makes multi-master difficult → impossible. Thus scanning the bus continuously is required.
  - Using Ethernet the endpoint could transmit the alert at any time without a master/slave relationship.

What is being requested?

- PHY capable of supporting at least 24” of PWB trace on standard FR4 material using standard routing guidelines.
- Potential for silicon vendors to provide bridging devices to expedite moving to this topology.
- Consider multi-drop to reduce the total number of required segments.
  - Some manageable devices do not need a substantial amount of data transfer. These devices would still benefit from moving away from the master/slave topology and allow automated alerts.
    - DIMMs are one example where this could be used. Alternatively a bridge device could be designed that would act as a DIMM controller for management features (temp, …)


Thanks ☺
DELL EMC