

OPEN

Compute Summit
Engineering Workshop
October 30-31, 2014
Paris





Server Committee Workshop

Oct 30, 2014

11:00 – Open CloudServer v2 Overview

14:00 – OCS v2 Chassis & Blade

16:00 – Multi-node management

Oct 31, 2014

11:00 – Facebook v3 Motherboard

13:00 – OCP Mezzanine v2.0





Open CloudServer v2 specification

Submission Overview

Mark Shaw
Director of Hardware Engineering



Open CloudServer OCS features

Chassis 12U, EIA 19" Standard Rack Compatibility

- Highly efficient design with shared power, cooling, and management
- Cable-free architecture enables simplified installation and repair
- High density: 24 blades / chassis, 96 blades / rack

Flexible Blade Support

- Compute blades – Dual socket, 4 HDD, 4 SSD
- JBOD Blade – scales from 10 to 80 HDDs, 6G or 12G SAS
 - Compatible with v1 JBOD Blade

Scale-Optimized Chassis Management


- Secure REST API for out-of-band controls
- Hard-wired interfaces to OOB blade management



Open CloudServer v2 upgrade


Blade upgrade

- Intel E5-2600 v3
- 36% higher performance
- 2.67X more memory
- 4X more flash memory
- 40G networking
- 12G SAS



High Performance Chassis Upgrade

- New 1600W PSU, 20 millisecond holdup
- Blade power >300W
- 40G Ethernet networking
- PCI-Express x16 expansion




Chassis components

8 kW DC Capacity

- >300W DC blades
- Six 1600W PSU with 20 msec holdup
- Higher CFM fans

Tray upgrades

- 1 x 40Gb + 1 x 10Gb
- Mezzanine: x16 Gen3 PCI-Express



Open CloudServer OCS v2.0 blade

Motherboard

- Intel © E5-2600 v3, 2S, 120W
- 16 x DDR4 DIMMs, 512GB memory


Storage

- 4 x 3.5" HDD
- 4 x 2.5" SSD
- 8 x PCIe Gen3 x4 M.2 Flash Modules

I/O

- Blade single 40G or 10G NIC mezzanine
- Tray PCI-Express x16 Gen3 Mezzanine
- Dual 4X SAS 12G


V2 – 36% more
performance over v1¹



M.2 PCI-Express Gen3
NVME Flash Modules



NIC



¹SpecInt Rate E5-2670 v3 120W vs E5-2470 v2 95W



Chassis trays

Blade support

- 12V DC power, management
- Passive PCBA for high reliability


High Speed I/O

- 40G + 10G Ethernet, 12G SAS
- Tray mezzanine: x16 Gen3 PCI-Express

Simplified deployment and operations

- I/O cabling is pre-wired and tested
- Eliminates cabling errors during service
- Reduces need for cabling reseats

Schematics and gerbers contributed



Tray mezzanine: designed for advanced networking such as 25 Gbps



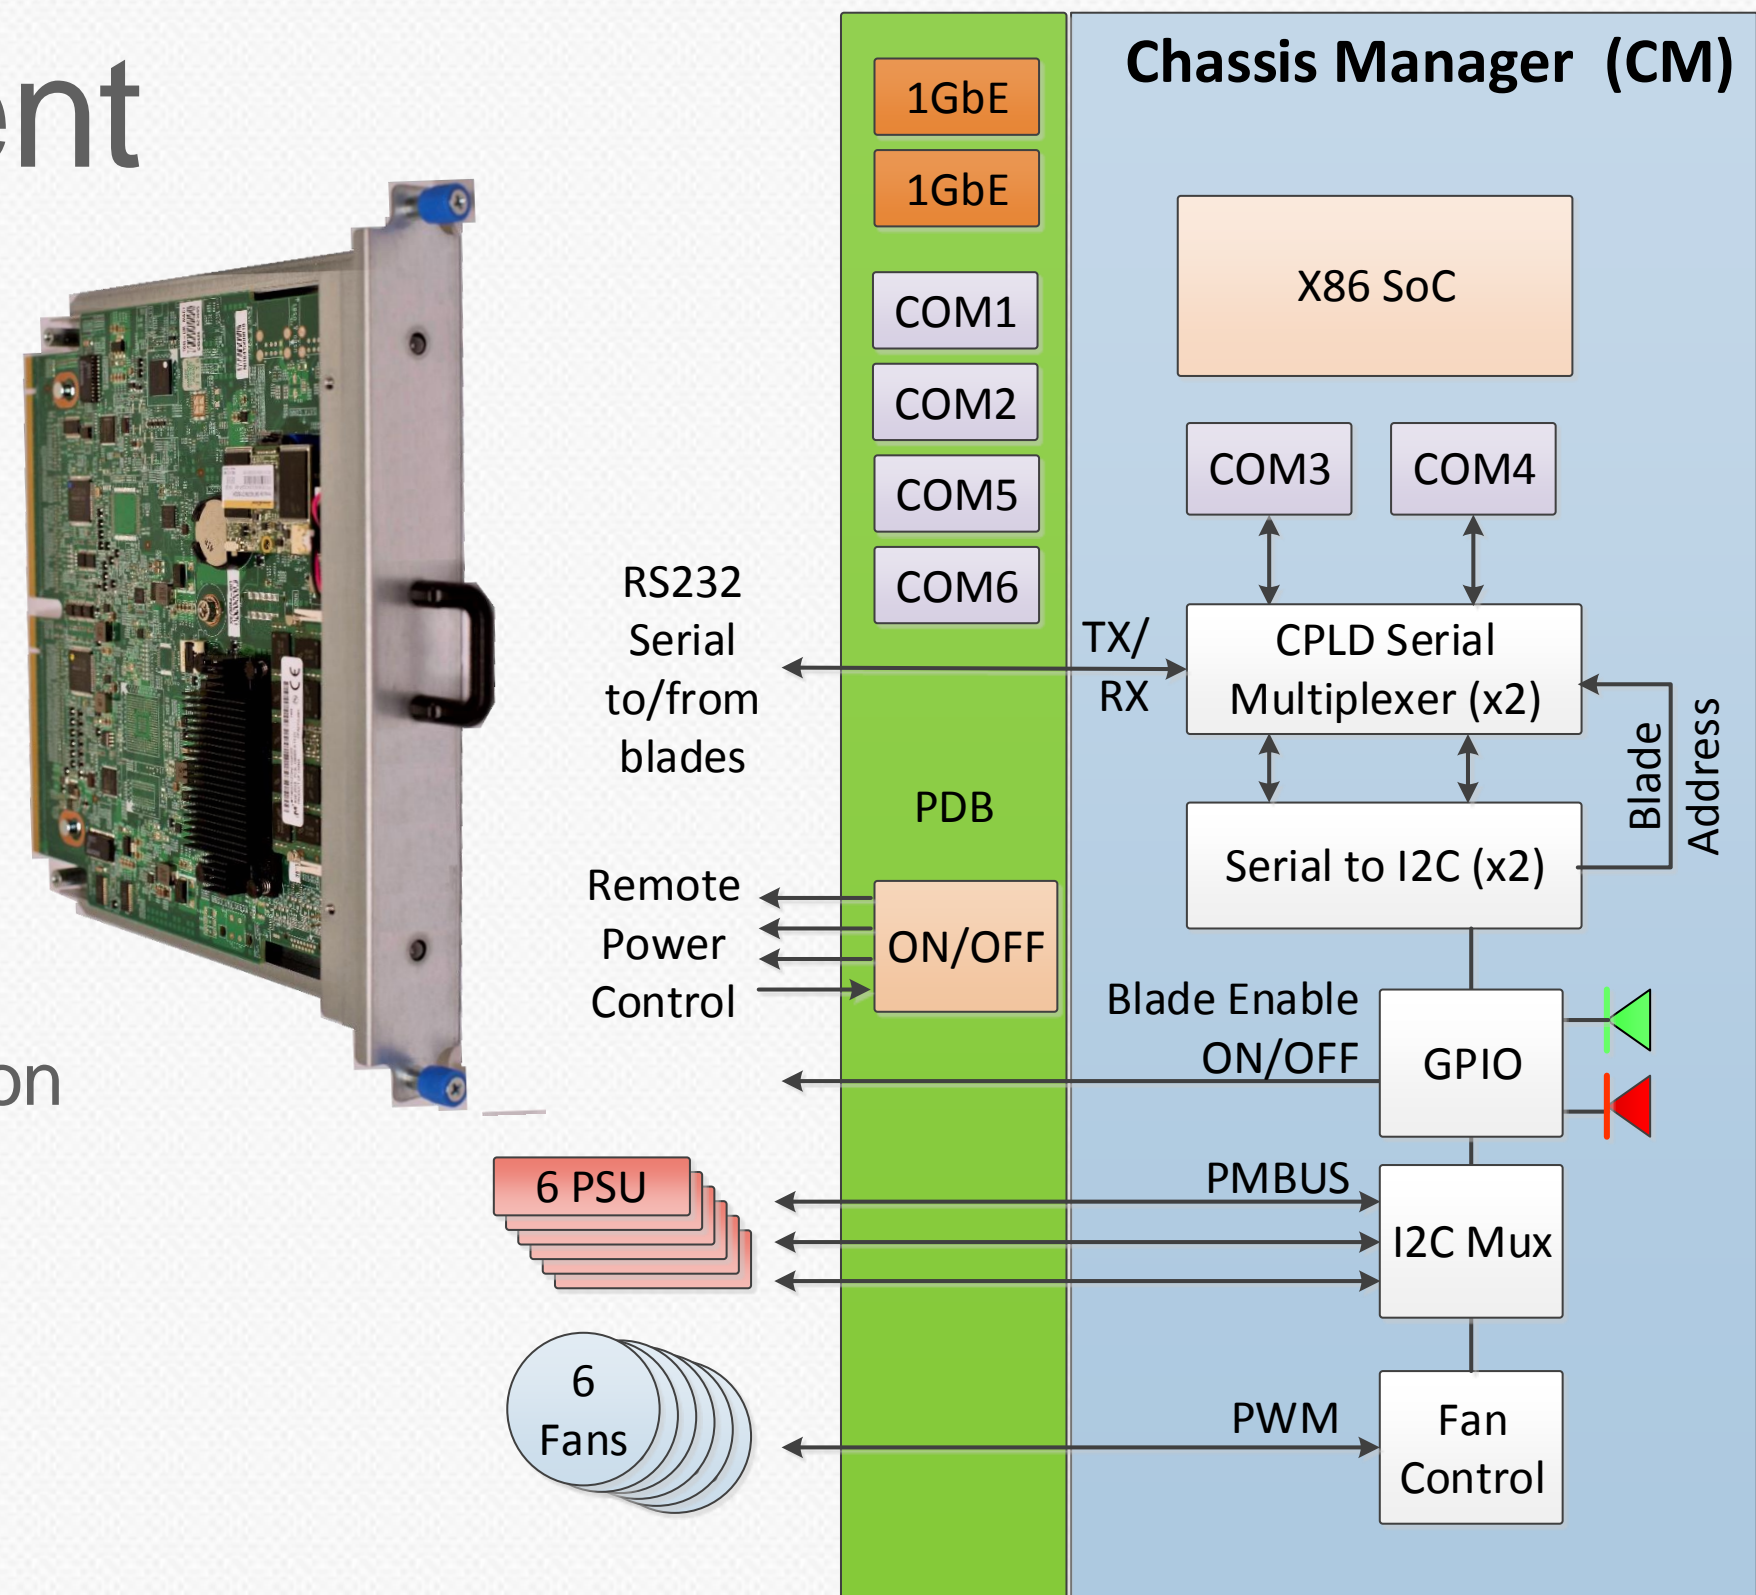
Chassis management

Secure OOB management

- Low-cost embedded x86 SoC
- REST API for machine management
- CLI interface for human operations

Hard-wired management

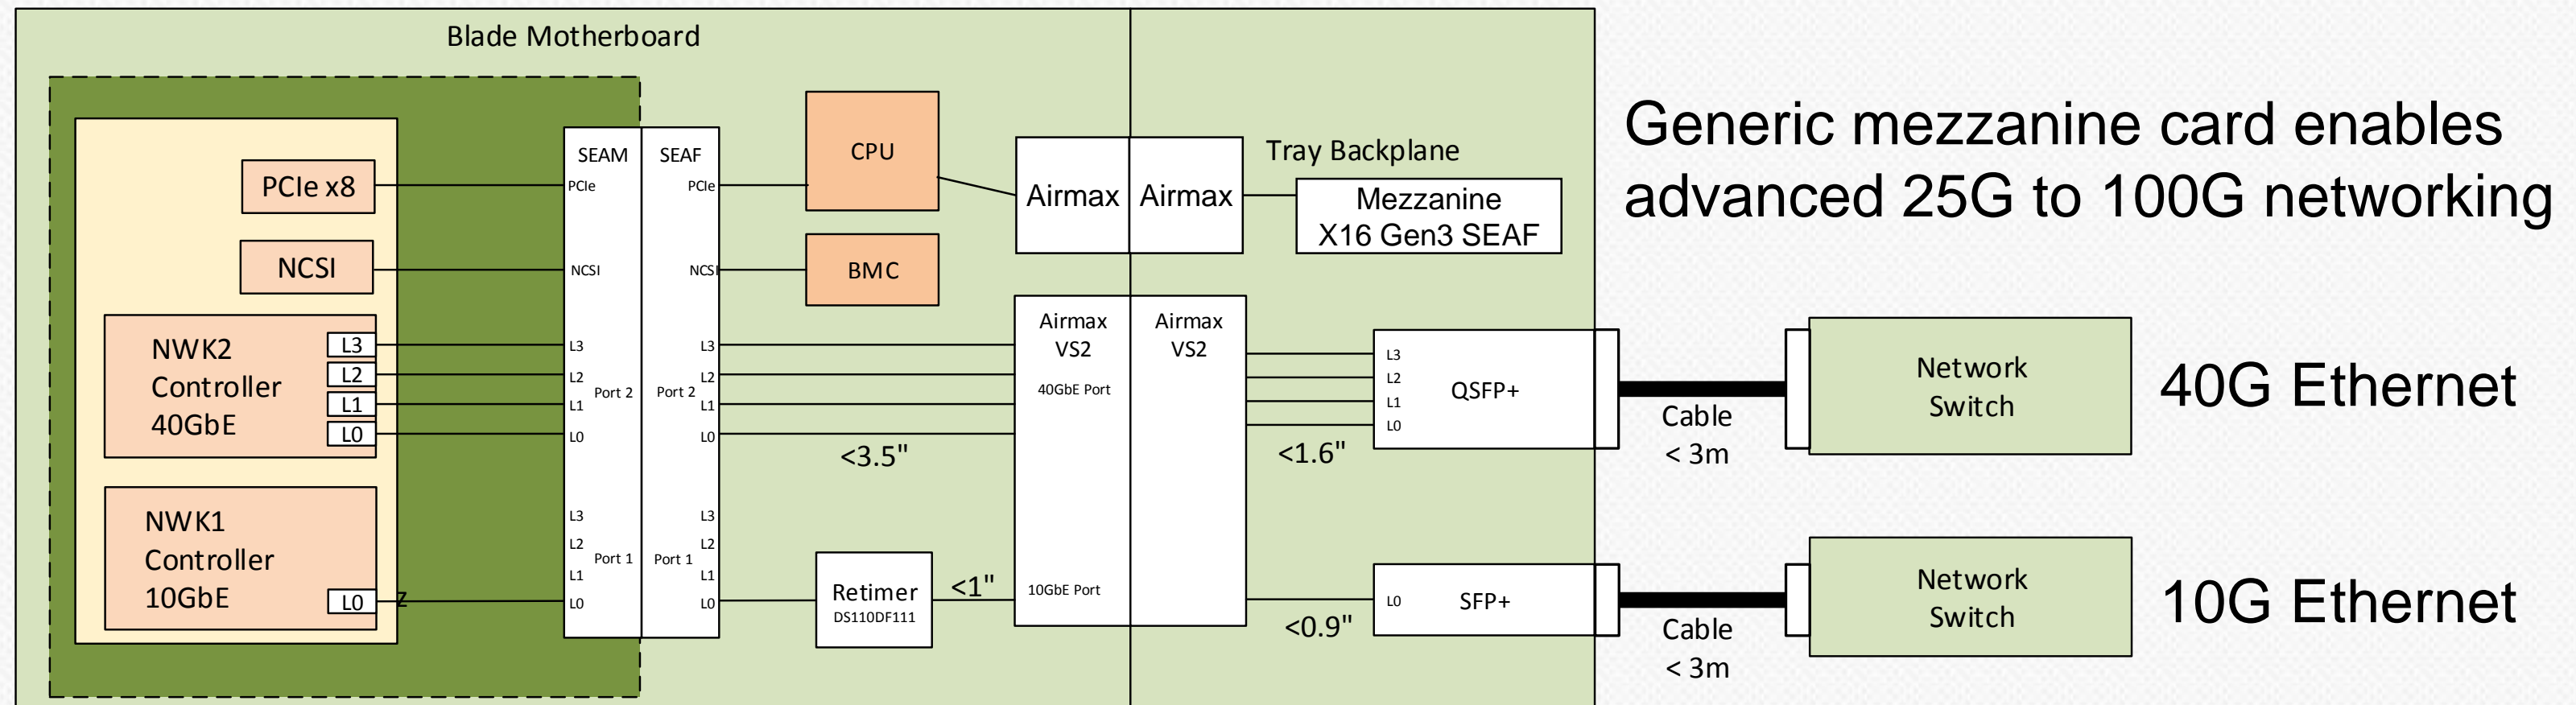
- On/Off to blade power cut-off circuit
- IPMI-over-serial out of band communication
- Fan and PSU control and monitoring
- Remote switch and CM power control
- Software is being open sourced
- Same hardware as OCS v1



Compute blade Networking

Flexible options to transition from 10G to 40G

Optional NCSI signals for side-band management



Expansion v1 JBOD reference design



20-lane SAS expander

- 10 internal lanes connect to LFF SATA HDDs
- 8 external lanes connect to tray backplane

Expander connects to chassis manager via RS-232 port

- Managed with the same command set as the compute blade



Safety and compliance

Ready for data centers world-wide



- Microsoft requires full compliance
- Containment at blade and tray
- Chassis is contained for use in EIA racks

Safety is Microsoft top priority

- UL, IEC, CSA standards among others

EMI Compliance is important

- CISPR, ANSI, IEC standards to start with



SW submissions

Operations Toolkit

Chassis Manager Service v2

QA and Interoperability test suite



Operations Toolkit

Component Updates

- Update programmable components such as BIOS and BMC
- Batch update of all programmable components

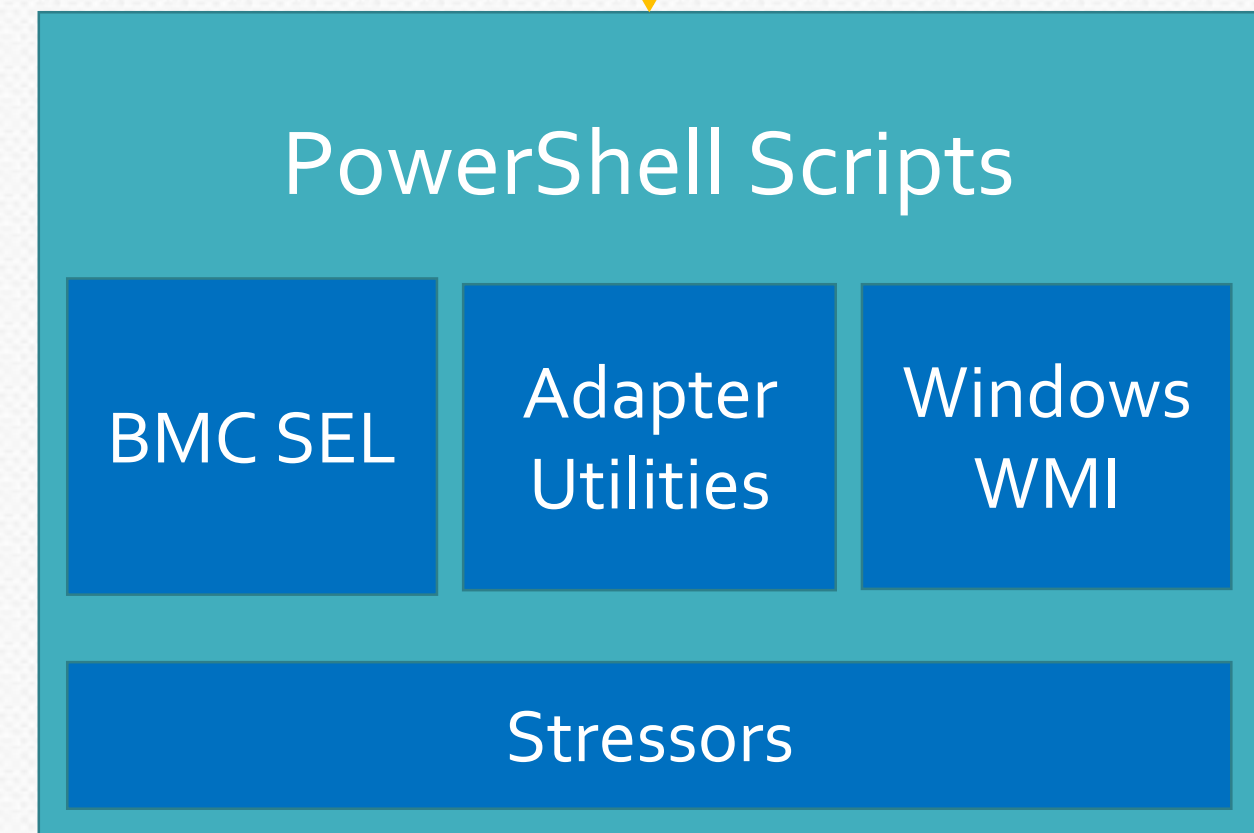
Diagnostics

- Identify defective components (HDD, DIMM, CPU)
- View, Log, and Compare Configurations
- Read, clear and log errors
- Execute IPMI and REST commands

Functional Tests

- System stress tests to identify intermittent problems
- Component specific stress tests (HDD, DIMM, and CPU)

OCS Commands



Chassis Manager Service v2

Better User Interface for CLI

- CLI auto-completion/scroll up-down/Function keys enabling
- VT100 support for serial sessions/DIGI sessions

Adding DIGI support

- CLI over serial
- Serviceability over Serial (crash cart)

QA enhancement

- Serial session hardening and state management
- Bug fixes

Adding QA test suite for submission



Certification Requirements

Specification and test methodology to be submitted

- Based on Microsoft's experience with validating multiple vendor blades

Software

- Certification of the blade-chassis interface via REST API test scripts
- BIOS requirements
- Important tests – power cycling and PXE boot

Hardware

- Mechanical
 - Form / fit / function tests
 - Thermal test results
- Electrical
 - High speed signal integrity tests
 - Targeted power / chassis level tests
- Certifications
 - EMC and Safety



Comprehensive Contribution

Open Source Code


Chassis management
Operations Toolkit
Interoperability Toolkit

Specifications

Chassis, Blade, Mezzanines
Management APIs
Certification Requirements


Mechanical CAD Models

Chassis, Blade, Mezzanines



Board Files & Gerbers

Power Distribution Backplane
Tray Backplane



Learn more

Visit Microsoft booth


- OCS v2 Systems on display
- Operations Toolkit Demo (every 30 minutes)

Attend executive track session:

- Growing OCS Ecosystem and Choice, 11:00AM, Oct 31

Attend technical workshops (Oct 30th)

- OCS v2 Hardware Overview, 11:00AM
- OCS v2 Chassis Management Overview, 11:00AM
- OCS v2 Operations Toolkit, 2:00PM
- OCS v2 Chassis and Blade Overview, 2:00PM
- Server and HW Management shared workshop (multi-node management), 4:00PM



Q&A





© 2014 Microsoft Corporation. All rights reserved. The information herein is for informational purposes only and represents the current view of Microsoft Corporation as of the date of this presentation. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information provided after the date of this presentation. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS PRESENTATION.

