

OPEN

Compute Summit

January 28–29, 2014 San Jose





Microsoft's cloud server specification

Hardware Overview

Mark Shaw,
Director
Microsoft Cloud Server Hardware Development



Microsoft cloud server spec features

EIA 19" Standard Rack Compatibility

Chassis 12U


- 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, 2 SSD
- JBOD Blade – scales from 10 to 80 HDDs

Scale-Optimized Chassis Management

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




Chassis components



Chassis assembly


1. 19" EIA 310-D Rack
2. Sidewalls
3. Fan Door
4. Chassis Manager
5. Six Power Supplies
6. Twelve 1U Trays
7. 24 Blades



Key features

Shared infrastructure for efficiency and TCO optimization

- Power delivery, mechanicals, thermals/cooling, management
- Optimized for mass contract manufacturing and assembly
- Up to **40% cost savings** and **15% power efficiency** benefits
- **Saves 10,000 tons of metal** per one million installed servers



Power distribution backplane

Six 1400W Power Supplies for N+1 Capability

- Commodity Off-the-Shelf for low cost
- Three phase balanced for efficiency at the data center

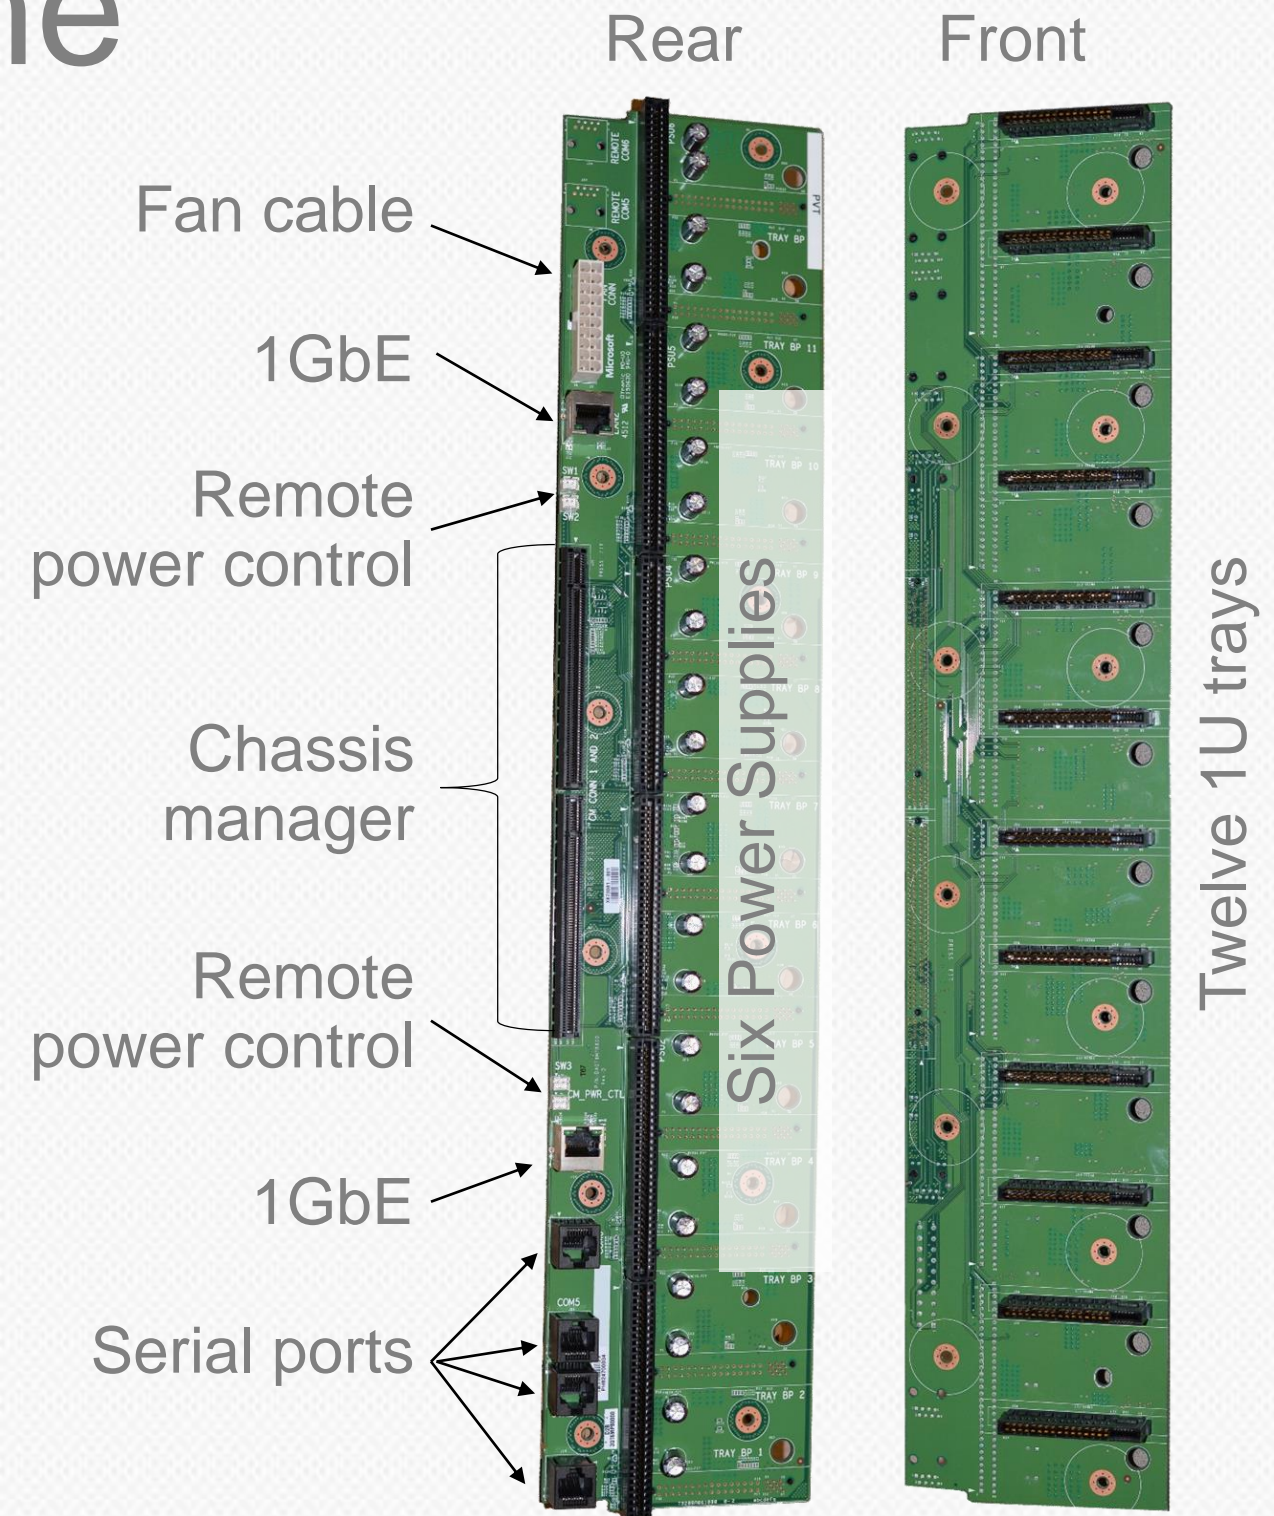
Chassis Management

- Hard-wired serial and on/off to blades
- Control and monitoring of Fans and PSU

Ease of Service

- Hot-swap chassis manager and power supplies for uninterrupted, error-free service

Schematics & gerbers contributed



Chassis and blade power

Optimized for N+1 Capability

- 6800 W at the chassis, 27.2 kW at the rack
- 24 x 250W Blades on average

Flexible for N+N Capability

- 4100 W at the chassis, 16.4 kW at the rack
- 14 x 250W Blades on average

High power blade capable

- Maximum 300W blade 12V DC power
- Average high and low power blades, i.e. storage servers



Fan door

Six shared high-efficiency 140 mm fans

- Sharing reduces peak and average power
- N+1 Reliability

Low-cost door holds all fans


- No hot-swap connectors or mechanics
- Simple lift-off hinge pins
- Cable harness quick disconnect

Serviced only when convenient

- Fans are intended to be fail-in-place
- Service is scheduled, annually if ever

Pro/E CAD contributed

Quick Disconnect





Thermally architected for reliability

HDDs are #1 failure item

- AFR increases with temperature¹

Simplified fan control cools HDDs

- HDDs in front of hot motherboard
- Closed loop fan moderates temperatures




¹DSN 2011: Impact of Temperature on Hard Disk Drive Reliability in Large Datacenters

Key features

Blind-mated signal connectivity

- Decoupled architecture for server node and chassis enabling simplified installation and repair
- Cable free design for significantly fewer operator errors during servicing
- Up to **50% improvement in deployment and servicing time**



Chassis trays



Blade support

- 12V DC power, management
- High speed I/O (10G Ethernet, 6G SAS)
- Passive PCBA for high reliability

Simplified deployment and operations

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

Schematics & gerbers contributed





¹ Understanding Network Failures in Datacenters: SIGCOMM 2011

Chassis tray mechanical


Designed for flexibility

Half width blades


Full width blades



Mechanical alignment features



1. Rear alignment post
2. Guide pins
3. Front alignment slot
4. Latch notch
5. Power and signal connectors



Key features


Secure and scalable systems management

REST API for chassis infrastructure

- TPM boot, SSL transport and Role-based authentication built-in
- Simplified server management via extremely small subset of IPMI commands
- Support for server diagnostics and self-health checks

For more information attend

- Software Management Overview
- Chassis Manager Hardware Overview



Chassis management

Secure OOB management

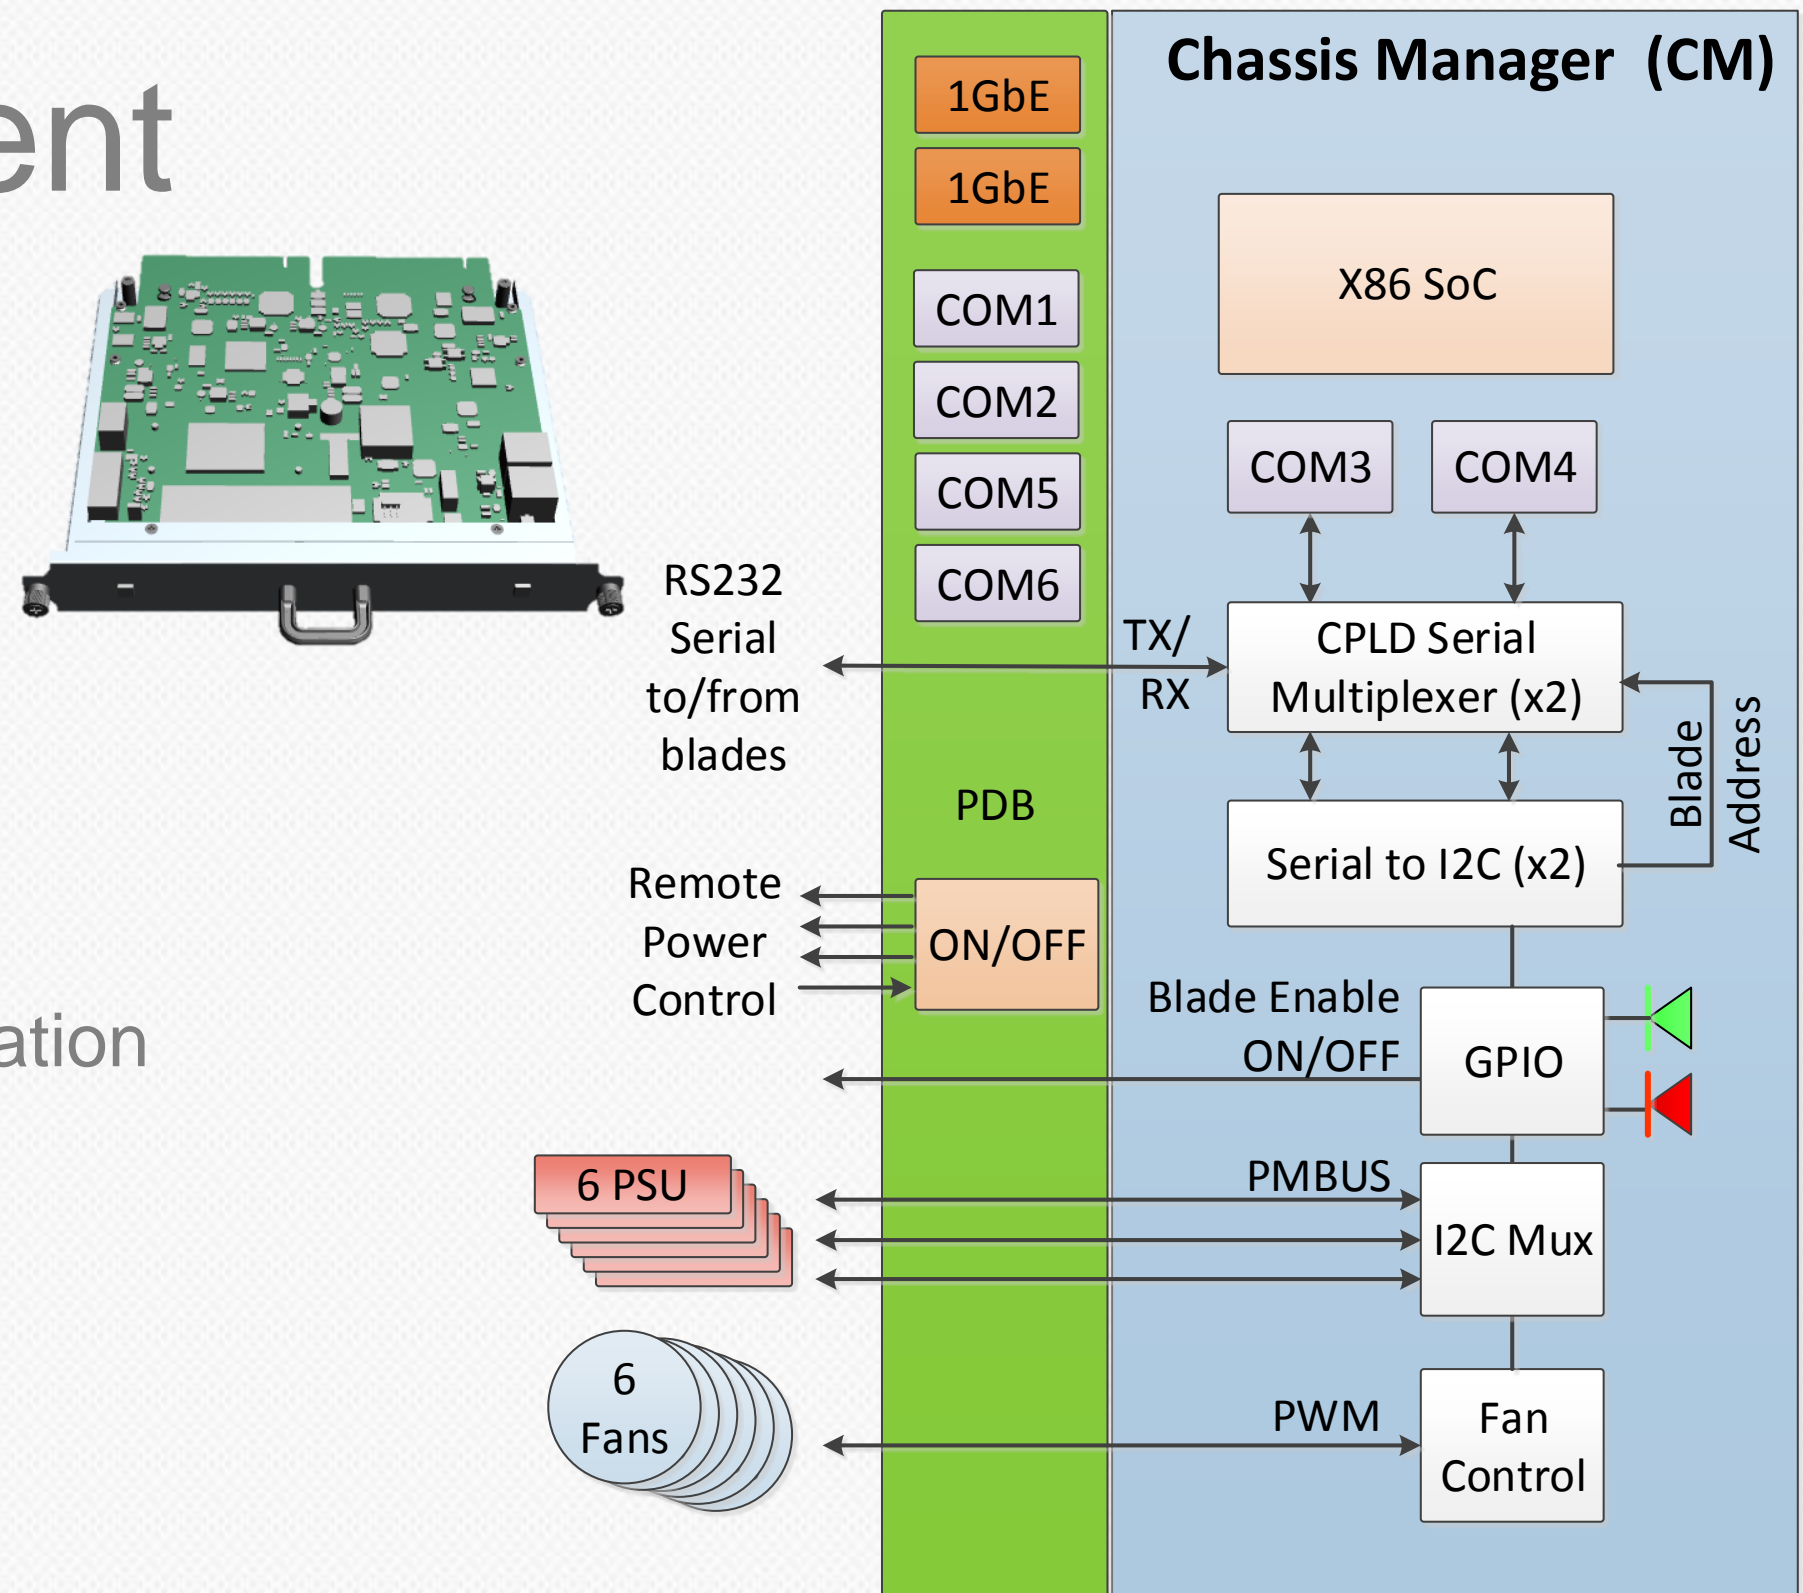
- Low-cost embedded x86 SoC
- REST API for machine management
- CLI interface for human operations

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

Schematics & gerbers contributed

Software is being open sourced

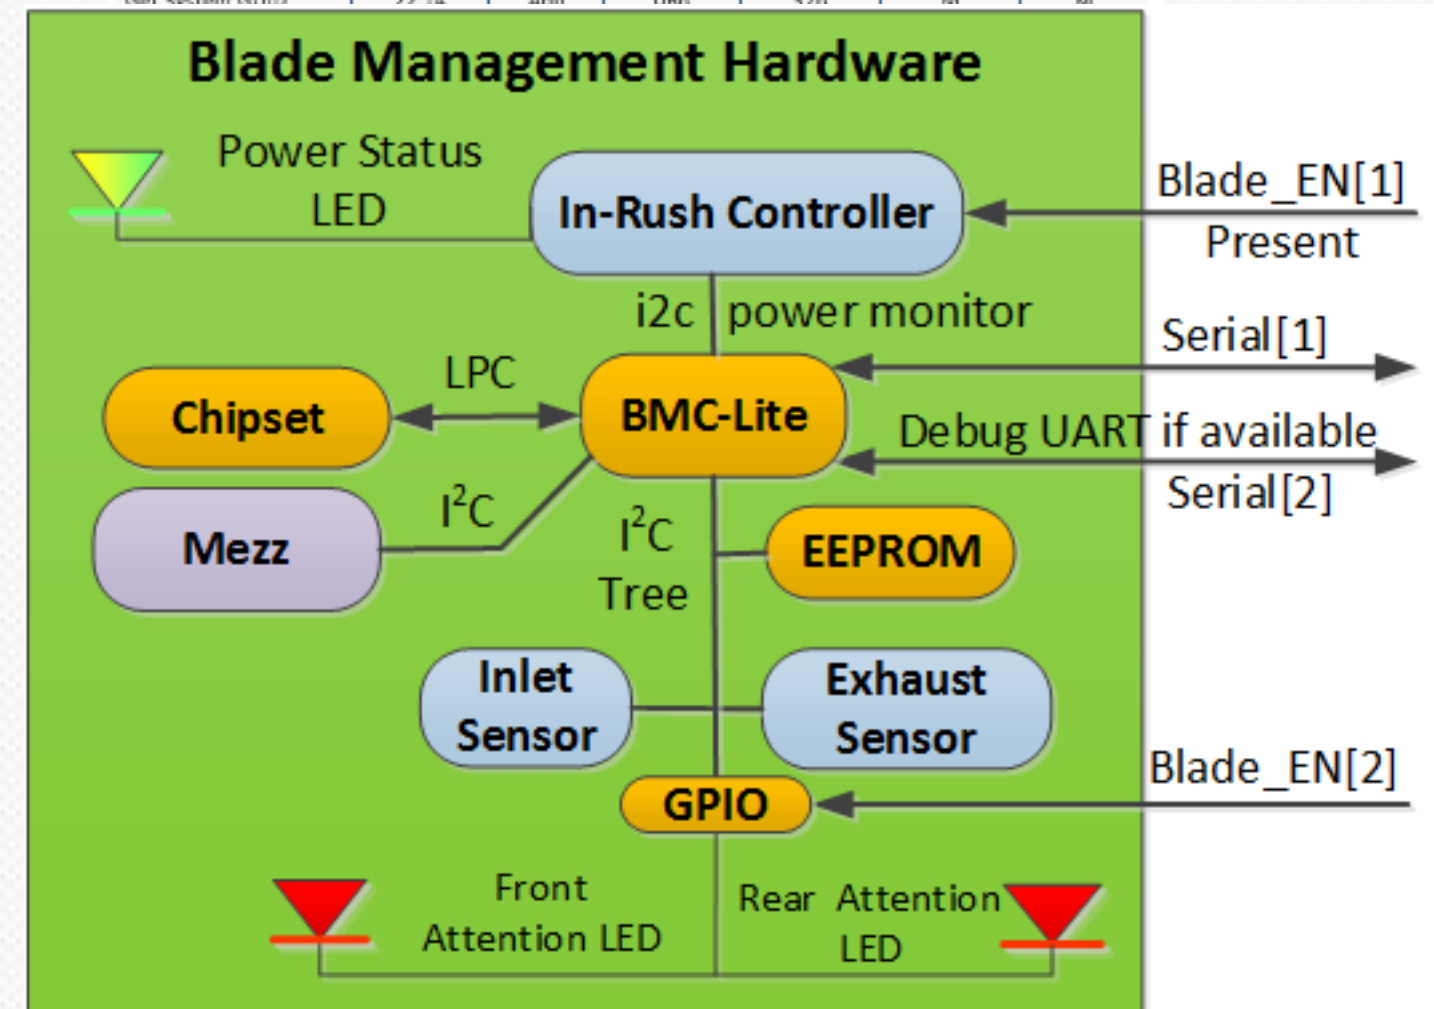


Compute blade BMC-Lite

BMC-Lite

- ✓ IPMI basic mode over Serial
- ✓ I²C Master (SDR)
- ✓ UART I/O
- ✓ System Event Log
- ✓ Power Control
- ✗ ~~KVM, Video drivers~~
- ✗ ~~Ethernet, Network Stack or SOL~~
- ✗ ~~USB~~
- ✗ ~~Full IPMI Command Set~~

Command name	Reference	Type	Fn	Cmd	Compute blade	JBOD blade
Get Device ID	20.1	App	06h	01h	M	M
Set ACPI Power State	20.6	App	06h	06h	M	N/A
Get ACPI Power State	20.7	App	06h	07h	M	N/A
Get System GUID	22.14	App	06h	17h	M	M




Get Chassis Status	28.2	Chassis	00h	01h	M	M
Chassis Control	28.3	Chassis	00h	02h	M	N/A
Chassis Reset	28.4	Chassis	00h	03h	N/A	N/A
Chassis Identify	28.5	Chassis	00h	04h	M	N/A

Microsoft cloud server spec blades

Compute blade


- Dual socket Intel E5-2400 v2 Family
- 4 x LFF HDD, 2 x SFF SSD
- Dual 10G networking
- Dual 4X SAS 6G
- Pro/E CAD contributed

For more information:
Attend Blade Overview -
Compute & Storage



Storage 6G JBOD

- 10 x LFF SATA HDD
- Dual 4X SAS 6G
- 6G SAS expander
- Pro/E CAD contributed



Safety and compliance

Ready for datacenter use worldwide



- 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



Security: defense in depth

Security at all layers

- Hardware, UEFI, APIs, User Management

Trusted Platform Monitor v1.2

- Blades and Chassis manager

UEFI Firmware v2.3.2

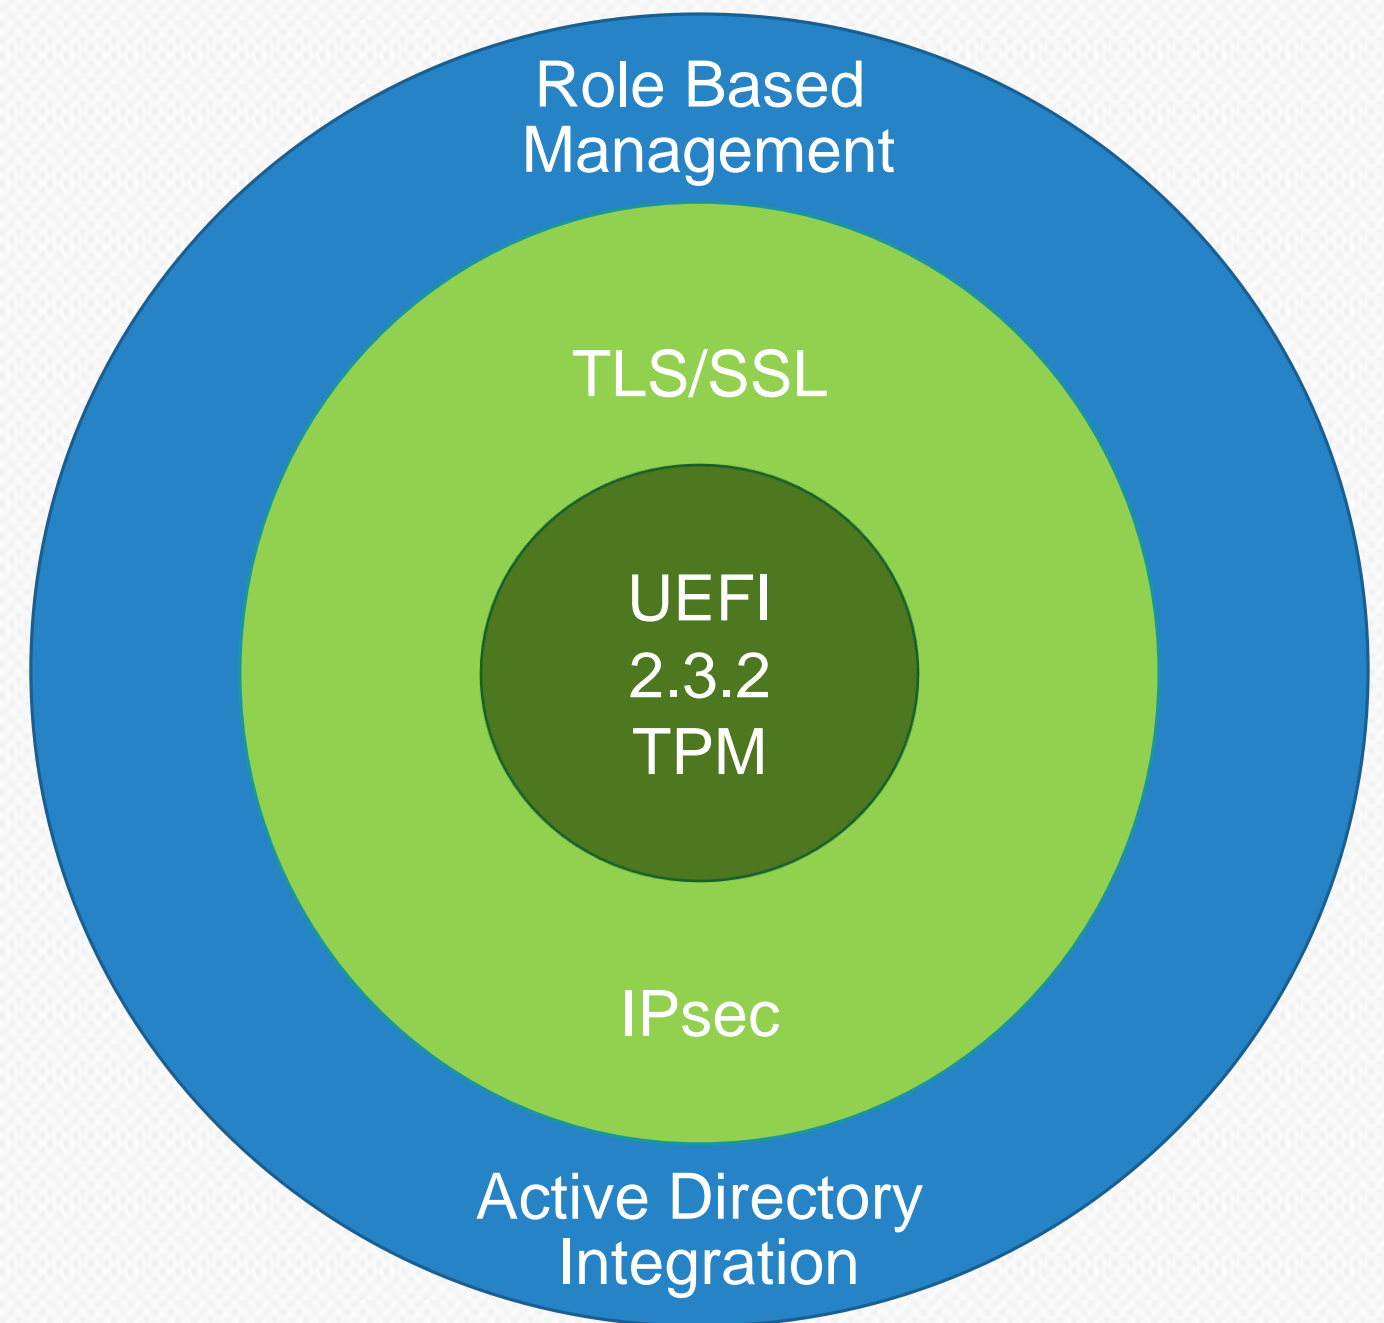
- Secure BIOS and Boot

Chassis manager interfaces

- TLS (SSL) and IPsec for communication encryption

User Management


- Active Directory integration and authentication



Key features


Network and storage cabling via backplane architecture

- Architectural flexibility for multiple network types
- One-time cable install at factory
- No cable handling needed during operations
- **Saves 1,100 miles of cable** for a deployment of one million servers



Interoperability optimizes supply chain

1. Chassis installed into racks



2. Networking integrated and tested



3. Deployment at datacenter



Supplier A

Supplier B

Supplier C



More information: Technical breakouts

Technical Workshop	Presenter
Hardware Overview	Mark Shaw, Director
Management Software Overview	Badriddine Khessib, Director
Blade Overview – Compute & Storage	Martin Goldstein, Principal Systems Architect
Chassis Manager Hardware Overview	Bryan Kelly, Sr. Platform SW Engineer

Visit the Microsoft booth for live demos by the subject matter experts



Microsoft cloud server spec: OCP contribution

Source Code

Chassis management
source code through
Open Source

```
/// <summary>  
/// Gets Fan speed in RPM  
/// </summary>  
/// <param name="fanId">target fan Id</param>  
/// <returns>Fan speed in RPM</returns>  
internal FanSpeedResponse GetFanSpeed(byte fanId)  
{
```



Specifications

Chassis, Blade,
Chassis Manager, Mezzanines,
Management APIs



Mechanical CAD Models

Chassis, Blade, Chassis Manager,
Mezzanines



Board Files & Gerbers

Chassis Manager, Tray Backplane,
Power Distribution Backplane





Q & A

Microsoft datacenter resources

Microsoft Datacenters Web Site & Team Blogs

- www.microsoft.com/datacenters

Windows Azure

- <http://www.windowsazure.com>

Office 365

- <http://www.office365.com>





© 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.

