



# Open CloudServer – vNext Generation

Aug 11, 2016  
Mark Shaw

Microsoft Cloud + Enterprise



# Microsoft's Open CloudServer v1, v2, v3

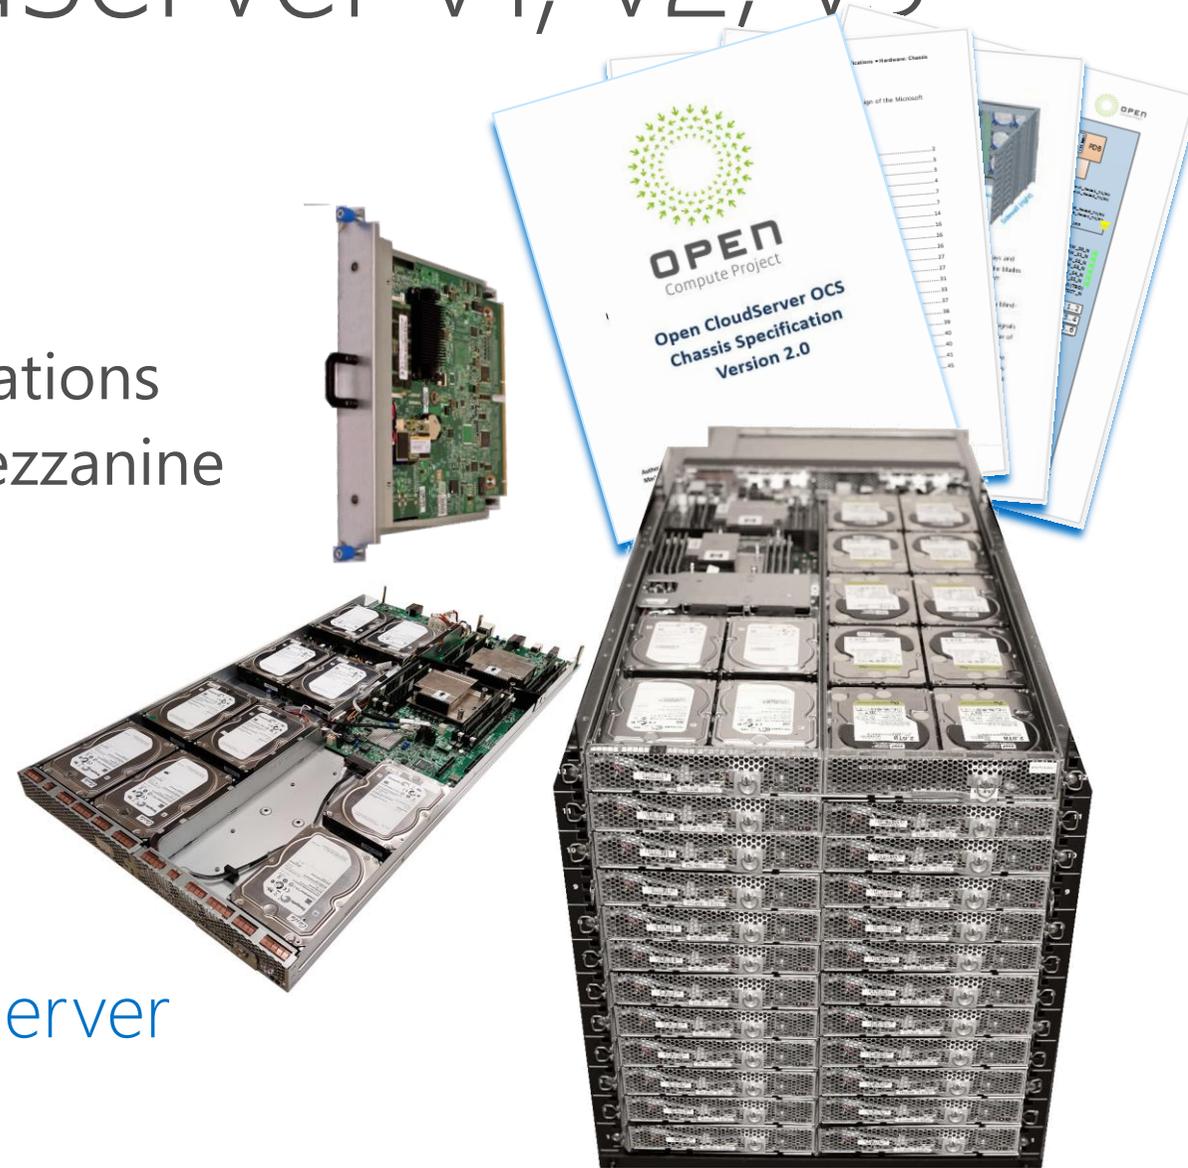
## Three Generations of Contributions

- Specifications, Board Files & Gerbers
- Open Source Code: Management, operations
- Mechanical CAD: Full Chassis, Blade, Mezzanine

## Continue to Focus on 19" EIA Rack

- International deployments
- Flexible heights, multiple vendors

Investigating the future of Open CloudServer



# OCS vNext – Design Principles

## One Design across Suppliers

- One motherboard
- One Firmware image
- One hardware design
- One qualification / certification
- Copy-exact for multi-sourcing

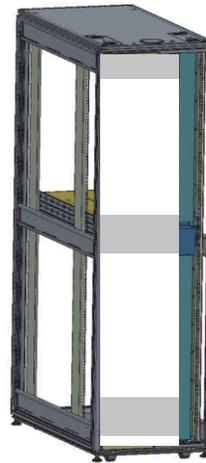


## Cost Reductions

- Simpler Hardware, less metal
- Lower management costs
- Very low cost without servers

## Modular 19" EIA Rack

- Owned and leased data centers world-wide
- Tweaked standard Rack
- 42U & 48U

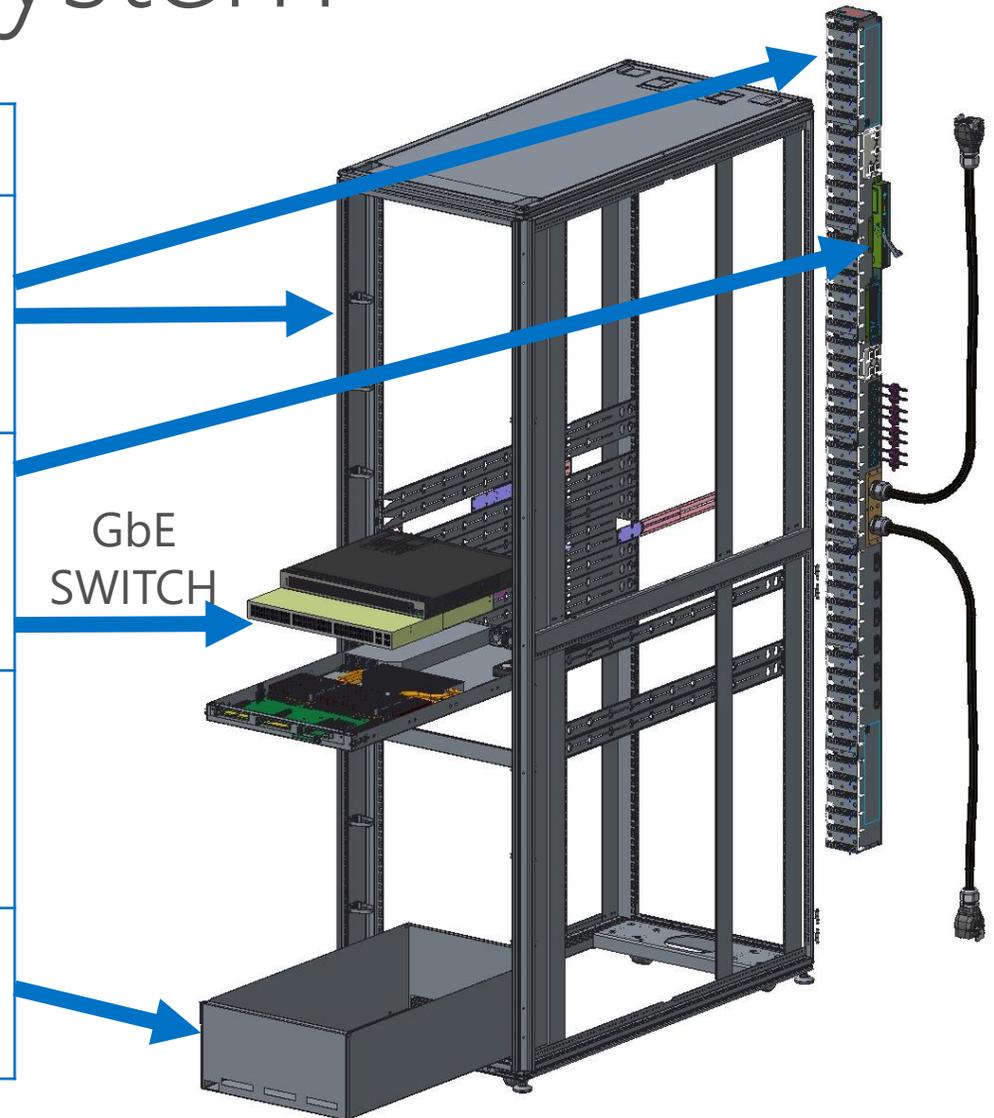


## Supply Chain

- Enable Microsoft's multiple manufacturers
- Single parts locker
- Fast integration & deployment

# OCS vNext – Modular System

Modules	Features
Rack / Enclosure	<ul style="list-style-type: none"><li>• 48U and 42U EIA 19", 1200mm depth</li><li>• 1U granularity, no 12U restrictions</li><li>• Dual A/C power feeds through PMDU</li></ul>
Management	<ul style="list-style-type: none"><li>• Rack Manager integrated into PMDU</li><li>• Same API as OCS v1 (RestFul or Redfish)</li><li>• Ethernet to each server's BMC, no serial</li></ul>
Server Blade	<ul style="list-style-type: none"><li>• Full-width side-by-side CPUs</li><li>• Front cabled I/O using standard PCI-e slots</li><li>• Local high-availability power and fans</li></ul>
Storage	<ul style="list-style-type: none"><li>• High density 88-HDD 4U JBOD</li><li>• Attaches to one, two, or four Servers</li></ul>



# OCS vNext – Server Blade

## Optimized for Performance, I/O flexibility

- Standard PCI-Express Gen3 slots, carrier card for OCP Mezzanine
- CloudSSD M.2 NVMe Flash slots on motherboard
- Side-by-side CPU layout lowers fan power

## Front Cabled I/O

- Blind mate rear power and management
- Latching at the chassis, pull to release
- Cold aisle service

## Embedded PSU and Fans

- PSU Dual-feed, three-phase, N+1 high availability solution
- Six N+2 fans for high-availability, lower per rack CFM



# Management Architecture

## Rack Management

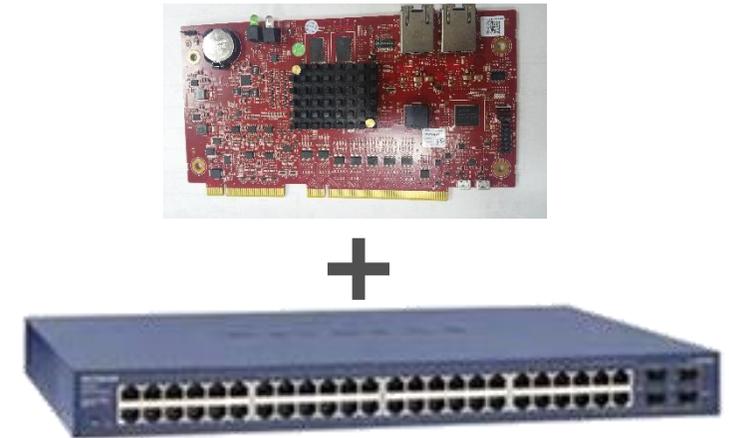
- Restful API I/F – i.e. Redfish via external Ethernet
- Rack Manager (RM) ARM appliance

## Blade Management

- GbE I/F to each blade's BMC
- GbE Network switch – front cabled
- NCSI enabled on motherboard with cable to OpenRack Mezz Carrier
- KVM enabled on motherboard

## Deployments

- Integrated into PMDU for rack scale deployments
- Standalone version for hardware that is not integrated into PMDU



# JBOD – High Density Storage

## 4U JBOD – 88 HDDs / chassis

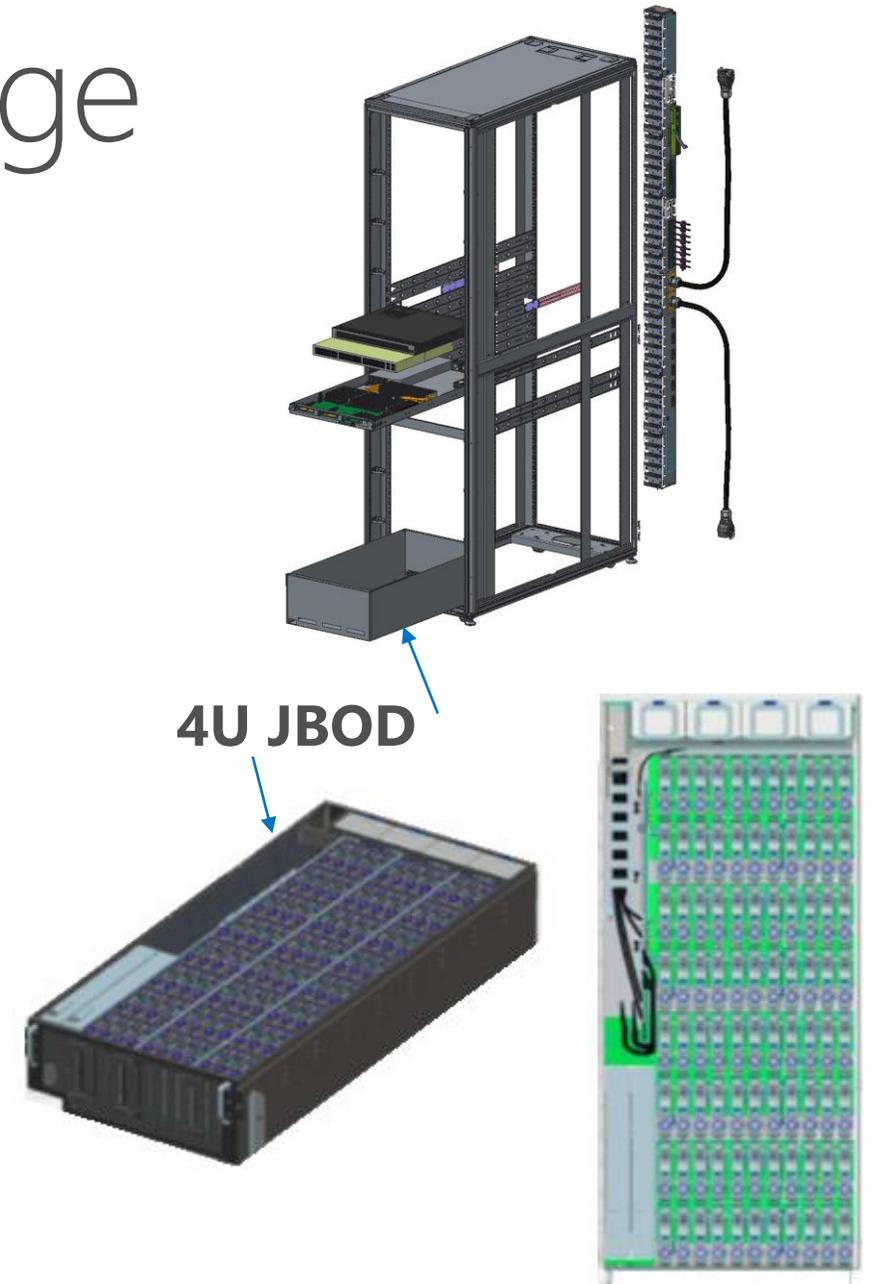
- Drawer design – slide out for repair
- Hot-plug HDDs, hot-serviceable expanders
- Front cabled to compute blade head-node

## Robust Feature Set

- BMC gathers HDD temps and status info
- Individual HDD on/off to minimize NTF

## Configuration

- One OCS vNext head nodes, 88 HDD each
- Two OCS vNext head nodes, 44 HDD each
- Four OCS vNext head nodes, 22 HDD each



# Additional Storage Investigations

## Low Density Storage

- 8-12 HDDs per 2U server, no hot-plug

## Flash Storage

- 48 – 96 NVME Cloud SSD M.2's
- 48TB – 192TB capacity

# OCS vNext Motherboard Flexibility

OCS vNext EIA 19" Racks  
*Microsoft Enabled*



EIA 19" 1U/2U Chassis  
Open Rack  
*Community feedback?*

