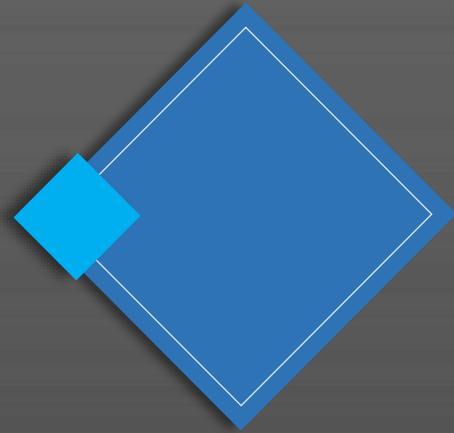




CLOUD COMPUTING AND AI TOTAL SOLUTIONS PROVIDER

Open Hardware Trends and Use Cases in the Hyperscale Datacenter

Dolly Wu, VP and GM
March 20, 2018



- Scaling Seamlessly Across the Globe
- Resource Utilization
- Intensive Computing for AI in the Cloud
- Processing Huge Data Sets
- Massive Data Storage

Open Hardware Projects for the Hyperscale Datacenter

6

International Open
Datacenter Platforms

Inspur is the only server vendor in the world to support all open source hardware projects.



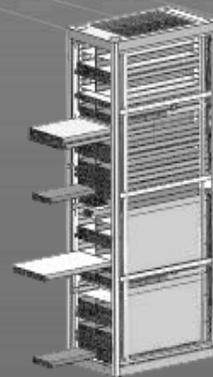
ODCC Rack



OCP Rack



Project Olympus Rack



Open19 Rack



InCloud Rack
with Intel® RSD



ODCC
Solution Provider



OCP Platinum Member
Solution Provider



Lead Vendor



Intel® Rack Scale
Design



Microsoft
Project Olympus

Inspur is a Key Member in
Open Platform Communities.



GOLD MEMBER

Inspur's Goal for being OCP Platinum Member and Solution Provider



Contribute building blocks to OCP Community



Standard Rackmount

inspur



OCP adoption for Cloud Datacenters in different industries



Deploy common server building blocks into all hardware platforms

Inspur Contribution to OCP: San Jose Motherboard

inspur

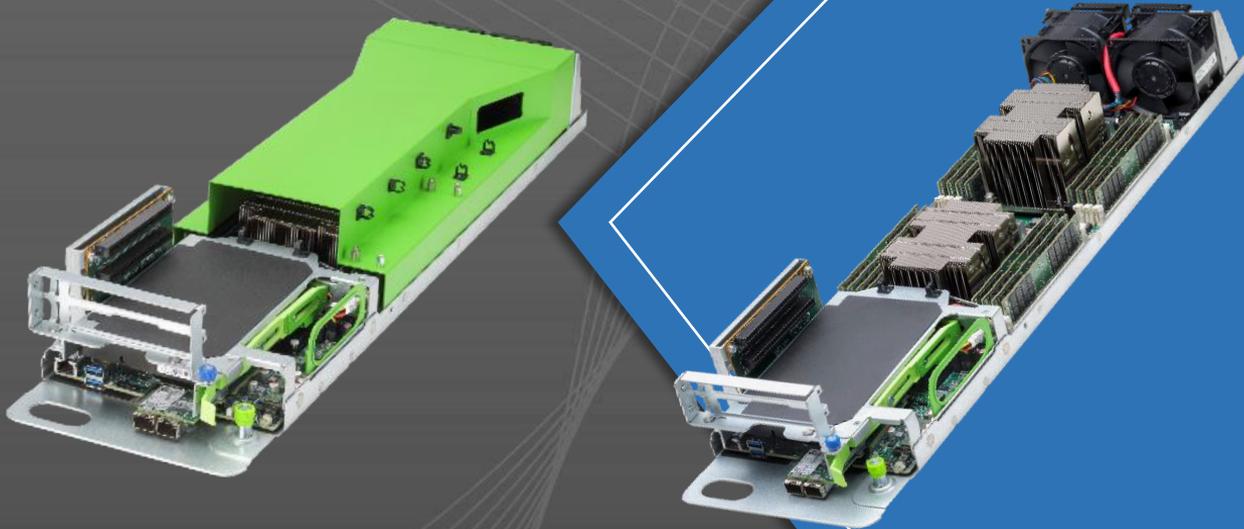


The World's First OCP-Accepted™ Motherboard Based on Intel® Xeon® Scalable Processor Family (Skylake/Purley)

Inspur Contribution to OCP:

inspur

INSPUR ON5263M5 The World's First OCP-Accepted™ Node Based on Intel® Xeon® Scalable Processor Family (Skylake/Purley).



Inspur will also contribute server management software to OCP. For more information, please go to Inspur Executive Track on 3/20 at 4pm, Room #210E/F. Live demo also available at Inspur booth # A16.

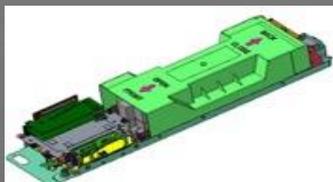
Inspur OCP Server and JBOD

inspur



2 OU 3*Node

- 2*Skylake CPU
- Support FPGA
- 16*DIMM GPU

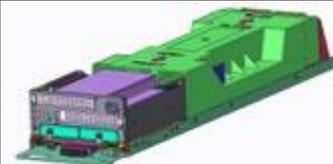


SKU#1 2 OU

1x3.5"HDD + 1x M.2

Expansion Slot:

- 1x FHHL(x16) : CPU0
- 1x FHHL(x8) : CPU0

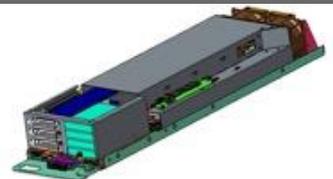


SKU#2 2 OU

2x2.5"HDD or NVMe
2x M.2

Expansion Slot:

- 1x FHFL(x16) : CPU0
- 1x FHHL(x16) : CPU1

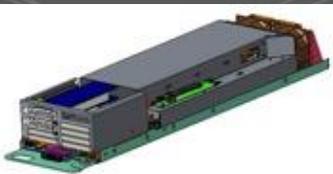


SKU#3 2 OU

4x2.5"HDD or NVMe
1x M.2

Expansion Slot:

- 3x HHHL(x8) : CPU0

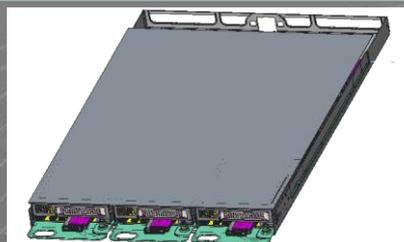


SKU#4 2 OU

6xSSD
1x M.2

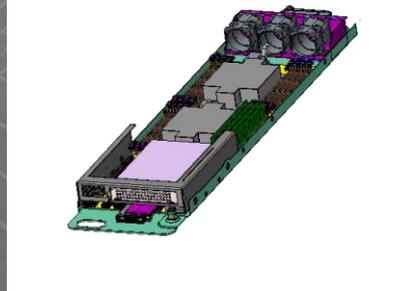
Expansion Slot:

- 2x HHHL(x16) : CPU0



1 OU 3*Node

- 2*Skylake CPU
- Support FPGA
- 16*DIMM GPU



SKU#1 1 OU

2x M.2

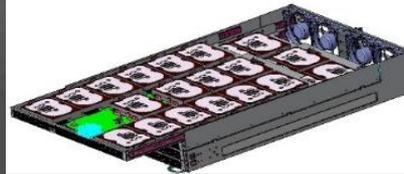
Expansion Slot:

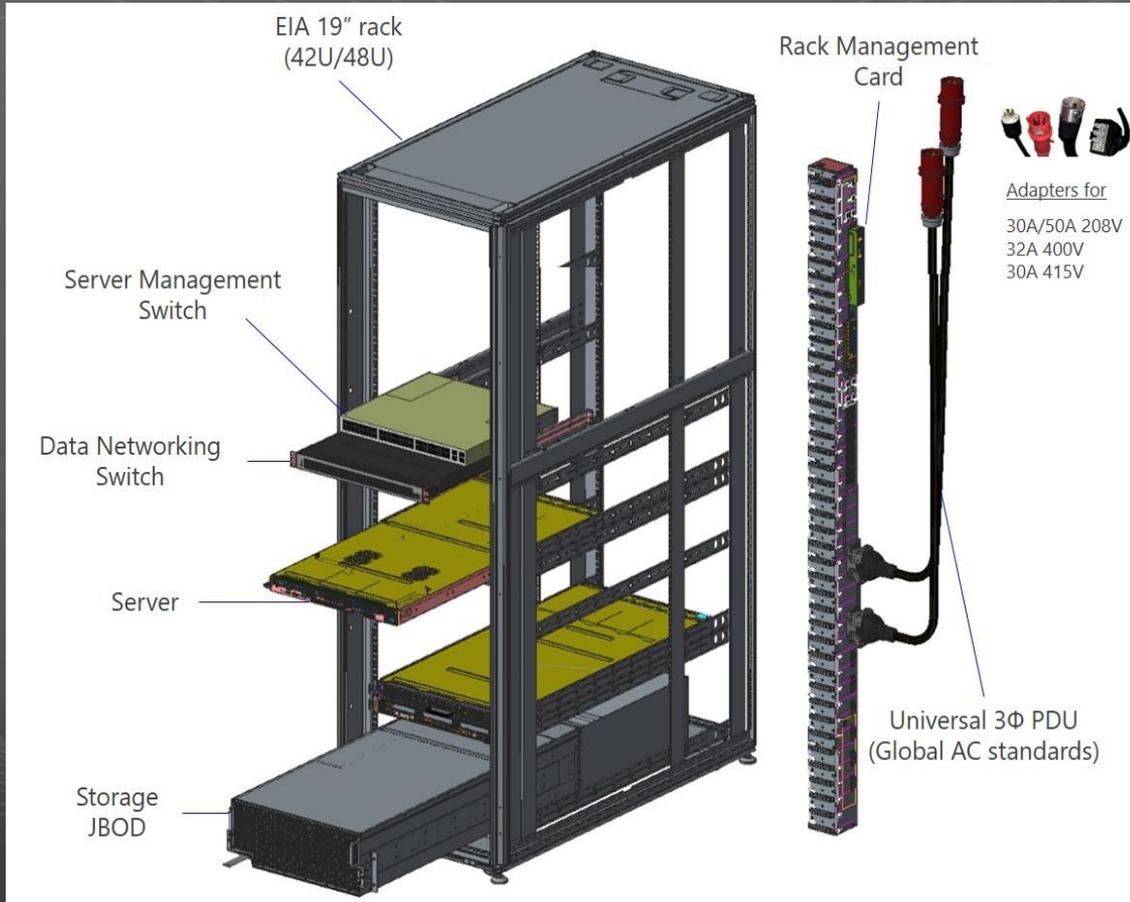
- 1x FHHL(x16): CPU0



3 OU JBOD

- 64*3.5"HDD
- Support M.2 or U.2
- Support Tri-mode





Important Differentiator: Universal PDU with adapters that work for every country – easy for global scaling

ODCC: Inspur SR Rack Scale System

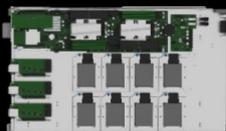


Inspur is the leader in ODCC Deployments with over 60% of the Market Share.

Storage Pooling



SR Server node



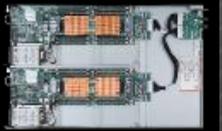
I/O Expansion



Balanced Node



Cold Storage



High Density Compute node



Centralized Power
Centralized Cooling



Power Consumption

-15%

Integration Delivery
Plug And Play



Speed Up Delivery

20x

Cold Aisle Operation
Centralized Management



Improve Operational Efficiency

3x

TCO

-12%

Failure Rate

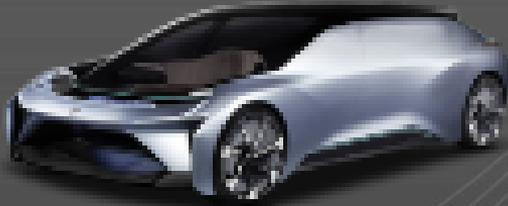
-50%

- 85% of Baidu rack server products
- 60% of Alibaba rack server products
- 40% of 12306 railway system Phase II
- 60% of Qihoo's rack server products are Inspur Rackscale SR

Case Study: Baidu



Inspur is the largest server provider for Baidu.



Jointly developed and released the world's first 4U16GPU for Baidu Driverless Car Project.

- Baidu is the world leader in voice, image and autonomous driving technology. Baidu Brain TTS is currently being requested 250 million times per day and delivers a speech recognition rate of 97% and a face recognition rate up to 99.7%.



- Inspur set a new record in 2016 of deploying 10,000 nodes per day using Inspur ODCC SR Rackscale servers for Baidu datacenters. In comparison, conventional rackmount servers deployment is around 300~500 servers per day.



**Singles Day 11/11/2017
eCommerce Revenue
\$25.3B for 1 day.**

Jack Ma, Chairman of Alibaba, celebrated Singles Day 2017 with actress Nicole Kidman.

- ▶ **Inspur is the No.1 Server supplier for Alibaba**
- ▶ **Inspur ODCC SR Rackscale infrastructure and AI Solutions supported the massive amount of data analytics, images search processing and floating point computation with single precision.**

Technology Resource Sharing from Hyperscale

Inspur JDM Model: Full product life-cycle customized to fulfill all types of customer needs



Hyperscale

All



JDM

Introduce JDM building blocks for Hyperscale to All CSPs, Enterprise, Traditional Industry





Delivery cycle shortened from 15 days to 5-7 days
Production efficiency increased by 30%
Reduced human workforce by 60%

Thank You

Inspur Systems, Inc.:

www.inspursystems.com

Dolly Wu, VP/GM, dollywu@inspur.com