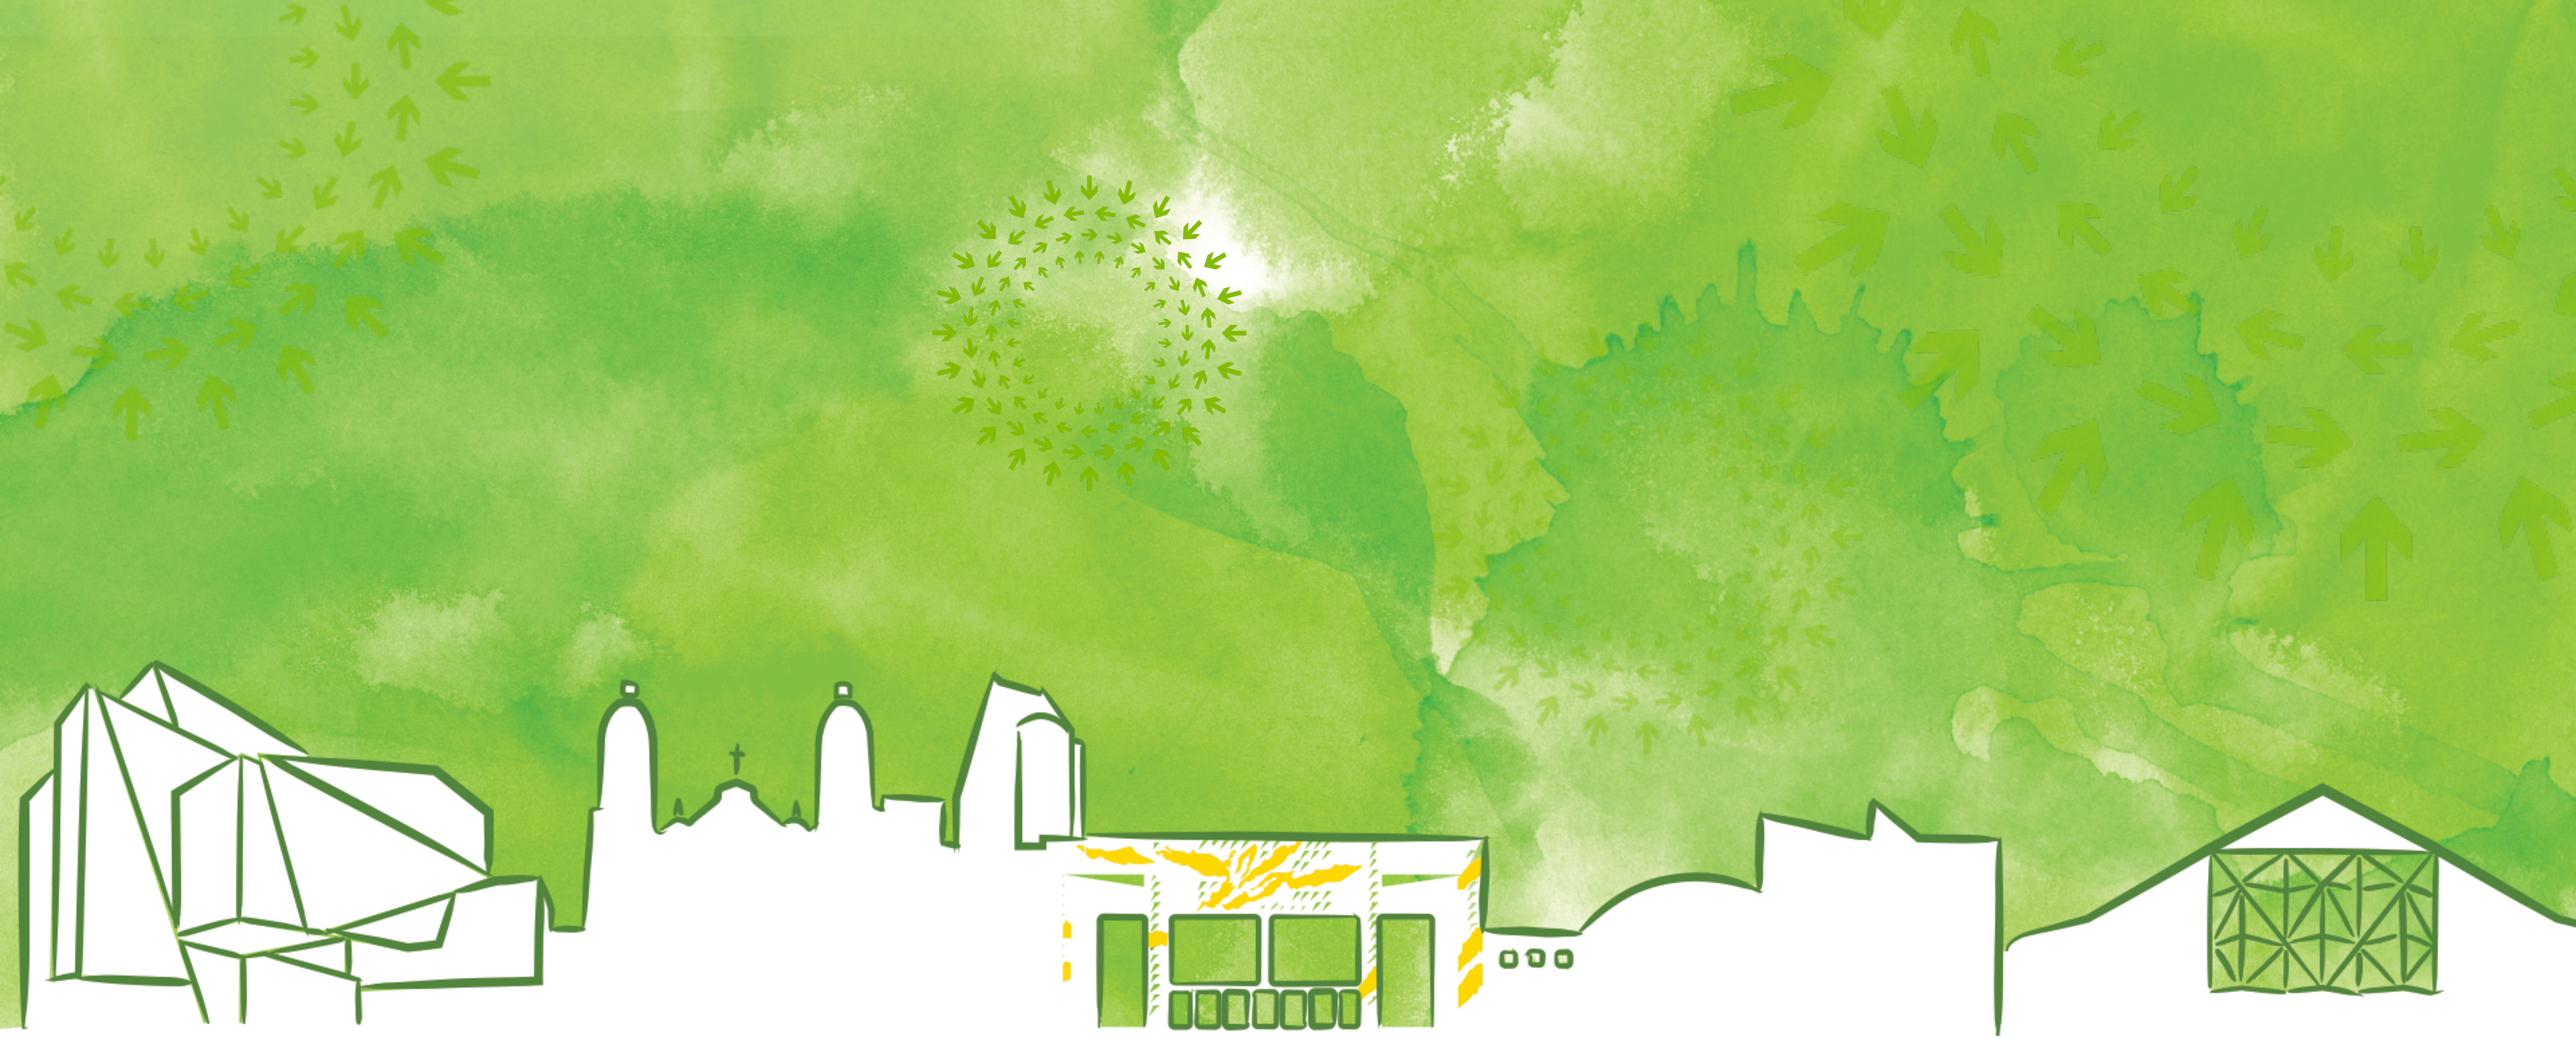

OCP Engineering Workshop 25 September 2017| Dallas, TX

OCP Engineering Workshop – 25 September 2017 – Dallas, TX

Microsoft Project Olympus
September Update

Mark Shaw
Director of Hardware Engineering, Microsoft

PROJECT OLYMPUS – OPEN SOURCE HARDWARE

Next-gen Cloud Hardware

Open sourcing leading edge Hyperscale cloud hardware currently under development at Microsoft

Development Model

New collaboration model with OCP community – co-develop open hardware at cloud speed

Industry Ecosystem

Bootstrap a vibrant ecosystem in OCP for the next generation of datacenter hardware

https://github.com/opencomputeproject/Project_Olympus

PROJECT OLYMPUS – OPEN SOURCE HARDWARE

Next-gen Cloud Hardware

Specs and Hardware on Github upgraded from 33% to 66% complete.

V1 coming soon

Development Model

Five motherboards and a GPU chassis built and underway

Five of six are from the industry

Industry Ecosystem

Multiple vendors will are enabling Project Olympus

More to come from OCP

https://github.com/opencomputeproject/Project_Olympus

1U SERVER

PROJECT OLYMPUS SERVER SETS THE STANDARD

Intel® Xeon® Scalable Platform

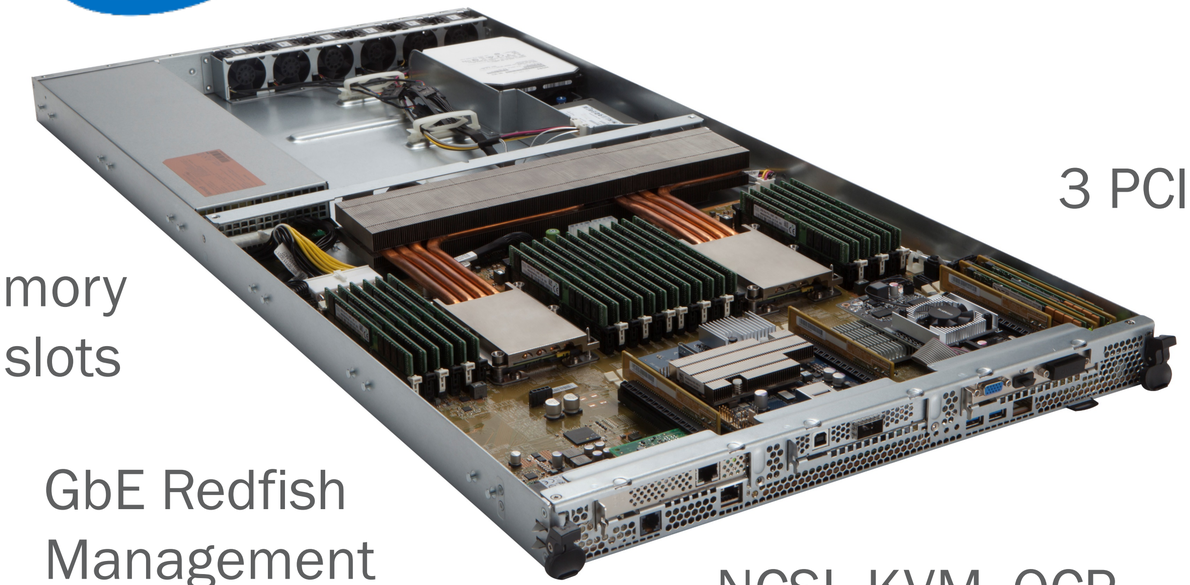

DDR4 memory
24 DIMM slots

3 PCI-E x16 FHHL slots

GbE Redfish
Management

NCSI, KVM, OCP
mezz enabled

Up to 16 M.2
CloudSSD

ECOSYSTEM MOTHERBOARDS

PROJECT OLYMPUS SERVER BUILT BY THE ECOSYSTEM

AMD

EPYC

3 PCI-E x16
FHHL slots

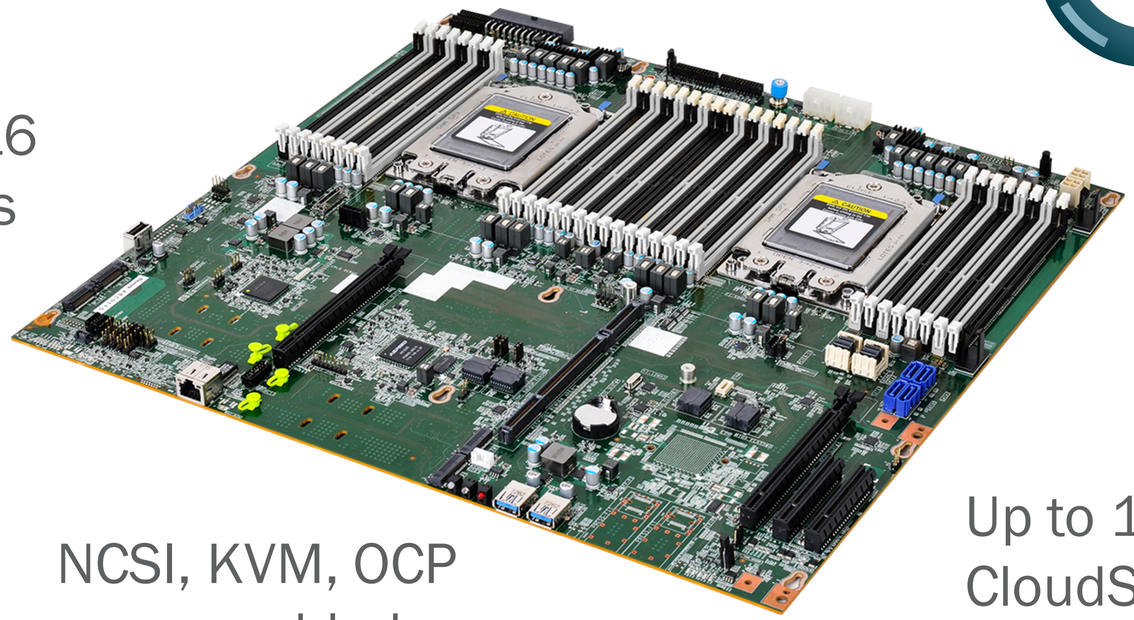

DDR4 memory
32 DIMM slots

GbE Redfish
Management

NCSI, KVM, OCP
mezz enabled

Up to 16 M.2
CloudSSD

ECOSYSTEM MOTHERBOARDS

ARM SERVERS

QUALCOMM

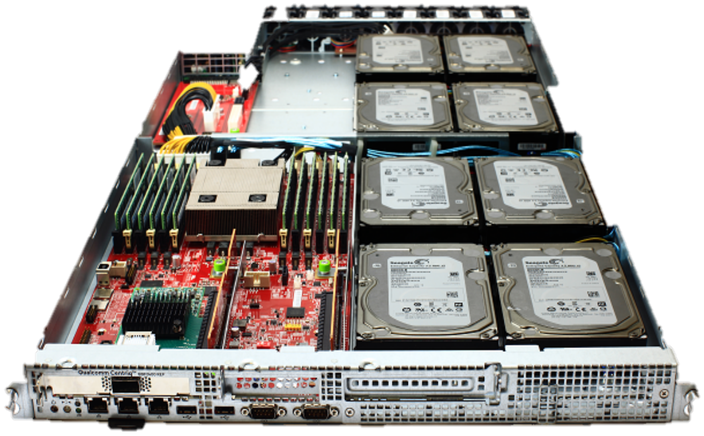

CAVIUM

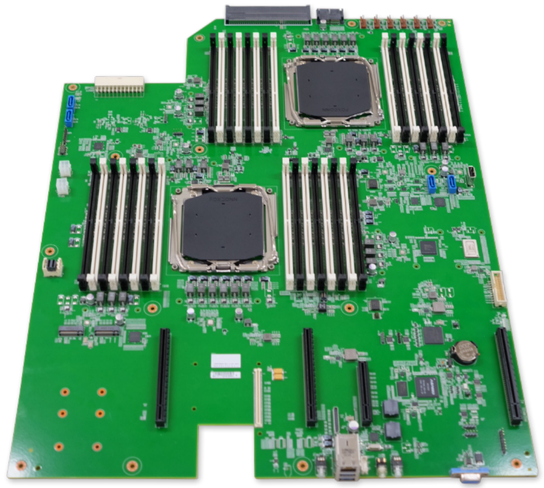

POWER SUPPLY

OPTIMIZED FOR EFFICIENCY, RELIABILITY, HYPERSCALE DATACENTERS

- Three x 340W PSUs Fully Integrated
  - Three-phase balanced AC power
  - 680W N+1 (1000W total)
  - Dual-feed auto-selection (IVS)
- Can operate with a single phase power cord
- Fault Mode Resiliency
  - AC feed failure, automatic fail over
  - PSU failure caps power if necessary
  - N+1 HA => no repair on failure
  - Double fault will be extremely rare

flex

ARTESYN™
EMBEDDED TECHNOLOGIES

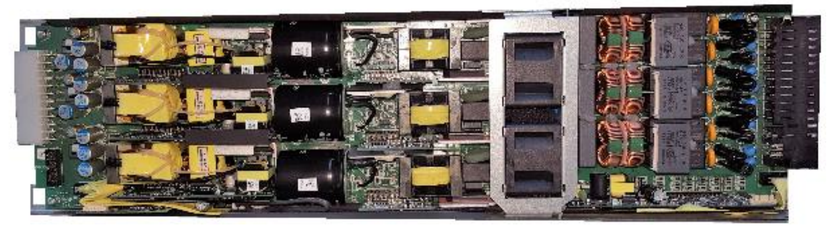

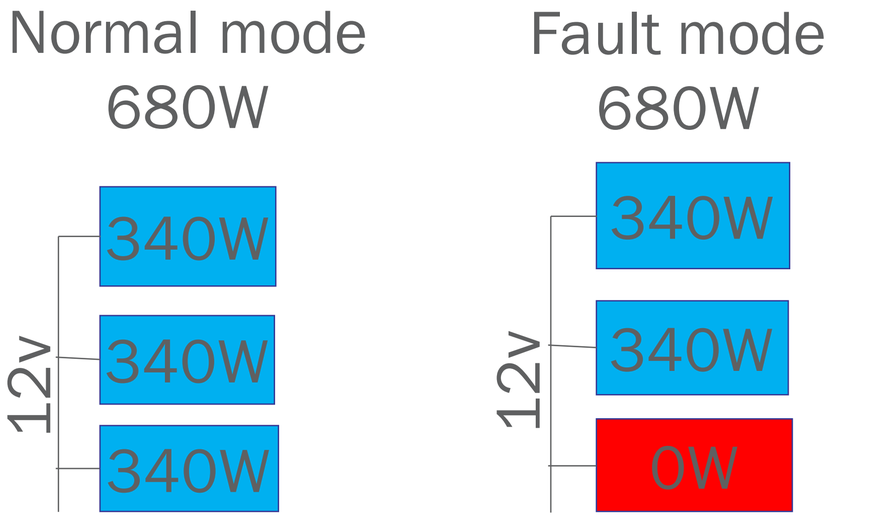

UNIVERSAL POWER

A/C POWER ADAPTED TO YOUR DATA CENTER

- Supports high availability data centers
  - Dual-feed, three-phase A/C
  - Rack power monitoring and capping
- Universal Power Distribution
  - Distributes power and management
  - Integrated Rack Manager
  - Supports blind-mate servers
- A/C power cord adaptation
  - Racks never change, only A/C cord
    - 208VAC 30A 3 phase and 1 phase
    - 208V 50A 3 phase
    - 415 30A, 400V 16A and 32A 3phase

HARTING®

Pushing Performance

GEIST

Future Thinking • Solutions Today

ingrasys®

19” EIA RACK

NO-FRILLS HIGH CAPACITY RACKS

- EIA-310 Standard
  - 600mm wide, 1200mm deep
  - 42U & 48U options
  - 3,000 lb capacity
- Features
  - Integrates with Project Olympus modules
  - Front & rear locking door, sidewall options
  - Three EIA rails for standard equipment
  - Baffles and air blocking panel options

RACK MANAGEMENT

DRIVING UP DATA CENTER UTILIZATION

- Rack Management
  - Restful API I/F or Redfish via Ethernet
  - Rack Manager (RM) ARM CPU
  - Integrated into PDU
- Blade Management flexible to your needs
  - GbE I/F to each blade's BMC
  - NCSI enabled, cable to OCP Mezz Carrier
  - KVM enabled on motherboard
- Standalone, 1U rack mount version
  - For hardware that does not use the PDU

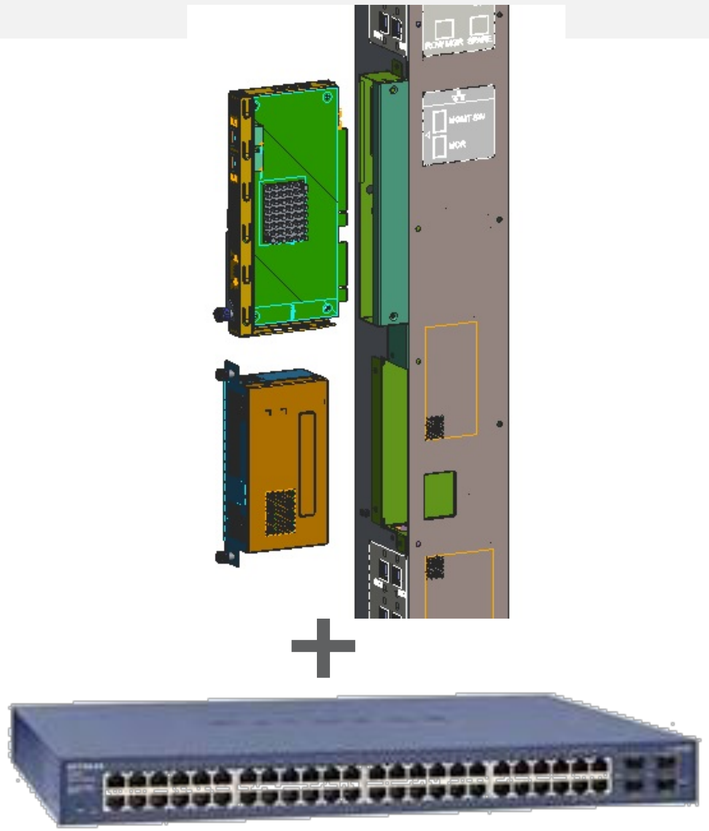

STANDALONE RACK MANAGER

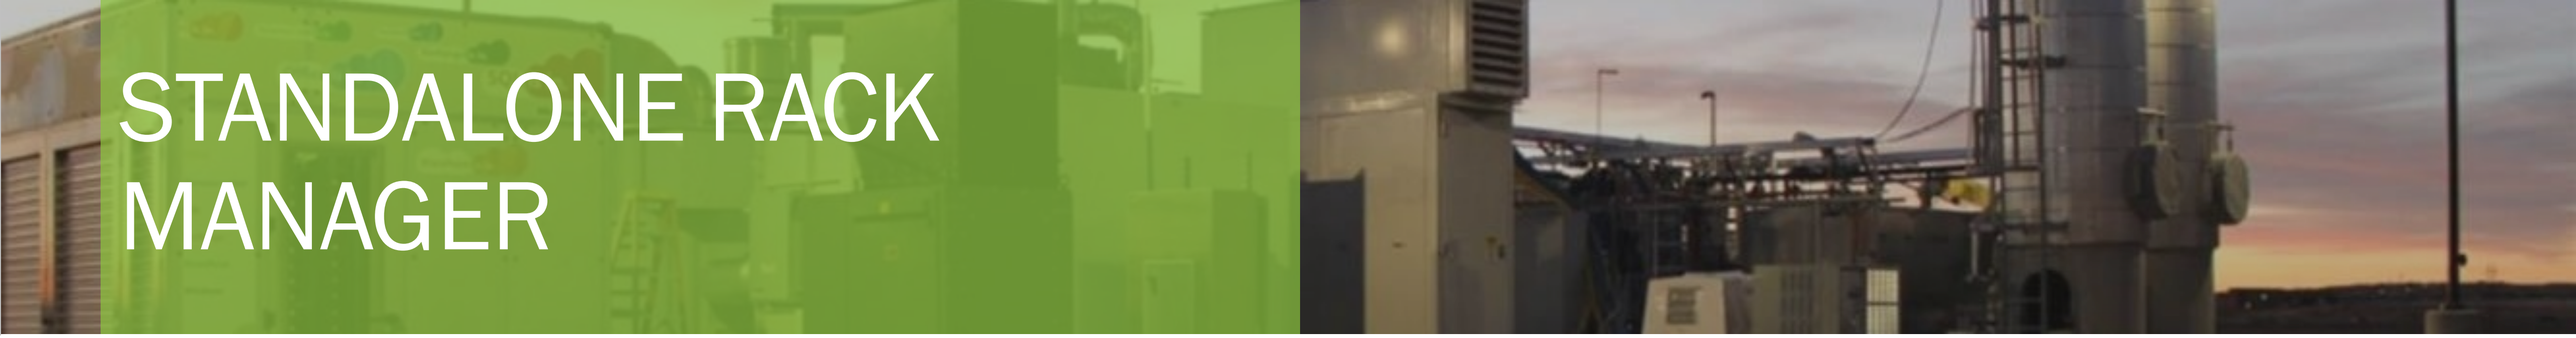

- 1U Rack Manager Assembly
  - Reuse of Rack Manager
  - Redundant AC Input
- Supports Non-WCS Rack Management
- Support Row Management
  - Power cycle Rack Managers
- Single SW image

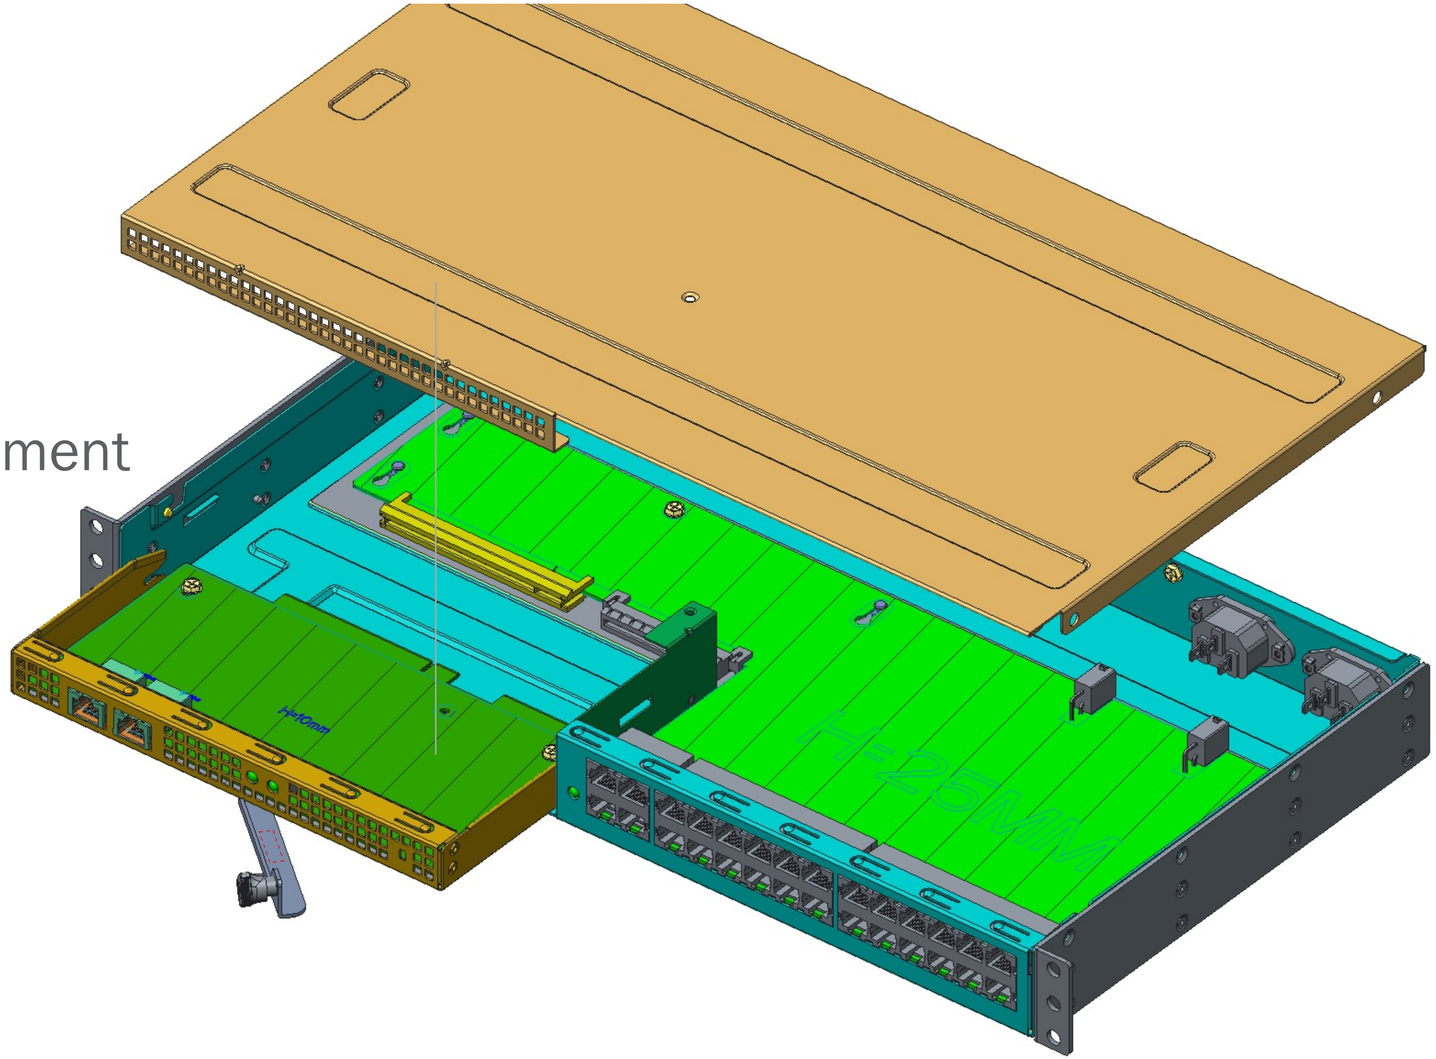

ARTIFICIAL INTELLIGENCE CLOUD COMPUTING

PROJECT OLYMPUS HYPERSCALE GPU ACCELERATOR - HGX-1

- 4U Flexible and Configurable Chassis
  - Eight NVIDIA SXM2 high-performance GPU
  - Eight PCI-E based GPU cards
  - Four additional x16 I/O slots
- Scalability
  - Up to 32 GPUs with four chassis communicating across PCI-Express fabric
- Flexible Topologies
  - Four PCI-E switches, eight x16 cable ports
  - Config Peer-to-Peer Bandwidth via cables

NVIDIA

ingrasys®

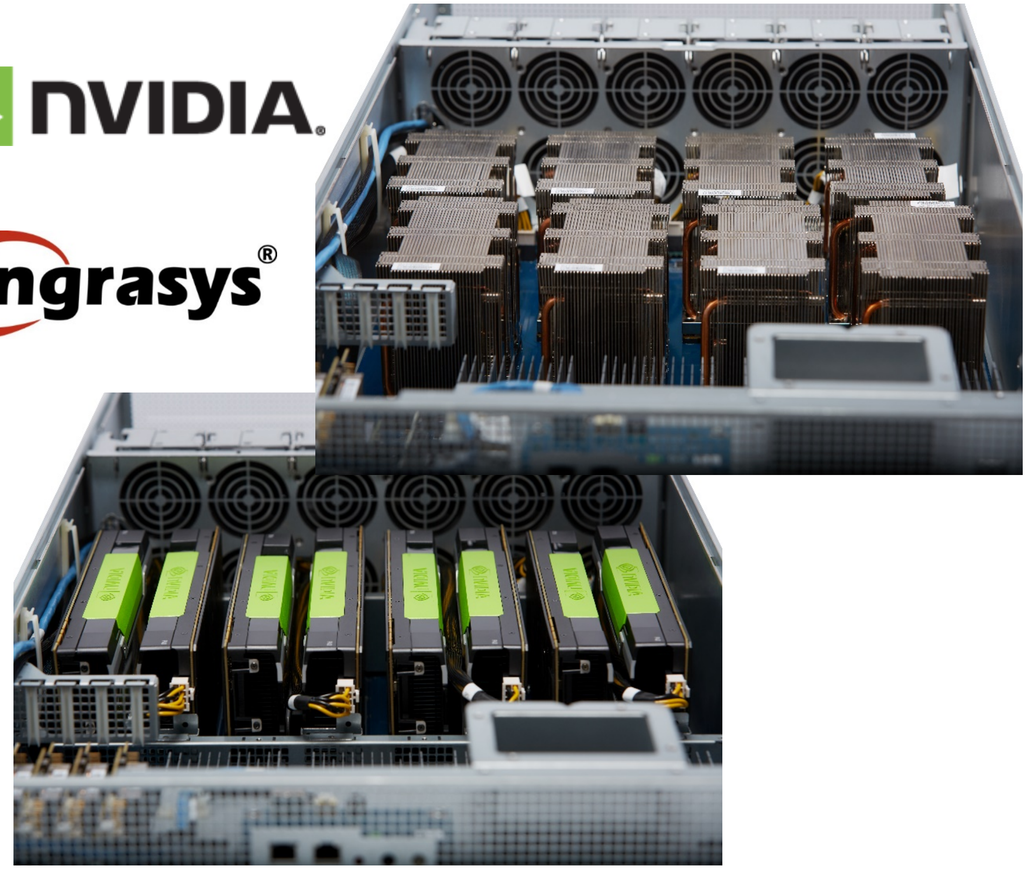

HDD STORAGE

HIGH DENSITY, HIGH RELIABILITY

- 4U JBOD – 88 HDDs / chassis
  - Cold-aisle hot repair service
- Robust Feature Set
  - Minimized Rotational Vibration, Acoustic Noise
  - BMC/Fans ensure <51C Temp
  - Individual HDD on/off to minimize NTF
- Partitionable with Project Olympus Servers
  - 88 HDDs, 1.2PB, on one server
  - 44 HDDs, 600TB, on two servers
  - 22 HDDs, 300TB, on four servers

Quanta

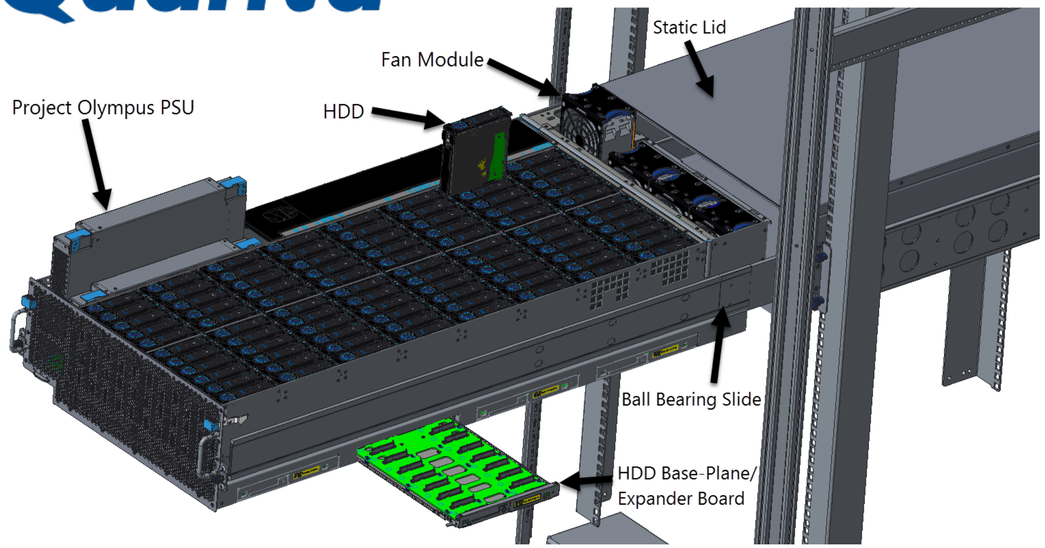

ARTESYN™
EMBEDDED TECHNOLOGIES

FLASH STORAGE

HIGH PERFORMANCE

- 1U JBOF – 64 NVMe M.2s / chassis
  - 64TB – 256TB capacity
  - 16 hot-plug M.2 carriers, 4 M.2’s each
  - Cold-aisle hot repair service
- Robust Feature Set
  - OpenBMC manages chassis, fans, power
  - Individual power domain per carrier

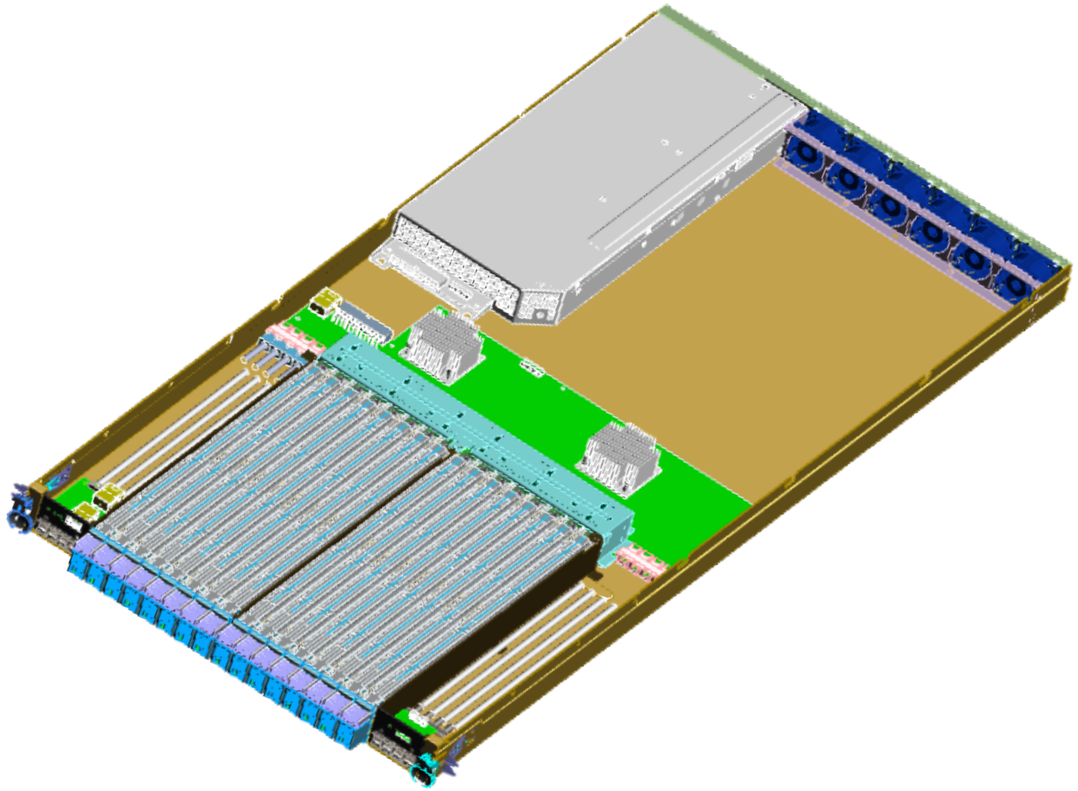

HOW TO DOWNLOAD

COMMUNITY PARTICIPATION

https://github.com/opencomputeproject/Project_Olympus

V1.0 coming soon

- LICENSE.md
- Project Olympus_Universal_PDU.pdf
- Project_Olympus_2UServer_Mechanical.pdf
- Project_Olympus_Chassis_Mechanical.pdf
- Project_Olympus_Rack_Manager.pdf
- Project_Olympus_Server_Mechanical.pdf
- Project_Olympus_StandAlone_Rack_Manager.pdf
- Project_Olympus_Universal_Motherboard.pdf
- Project_Olympus_Universal_Power_Cord.pdf
- Project_Olympus_Universal_Power_Monitoring_Distribution_Unit.pdf
- README.md

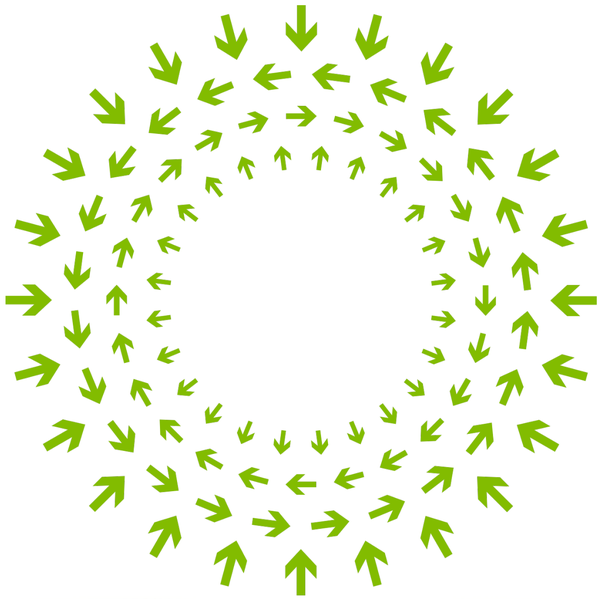

OPEN

Compute Project