



Case Studies

Case study: 1 OU Open Rack server and dense rack level configurations Samuli Toivola, R&D Squad Group Lead, Nokia





Telecom & openEdge



Open. Together.



Case Study: 1 OU Open Rack server and dense rack level configurations

This study analyzes benefits and challenges in case if density of Open Rack would be increased. Study compares possible server and rack level configurations if Open Rack triplet based server height would be changed from 2 OU to 1 OU.

This study analyzes both single server level challenges related to configuration and cooling limitations as well as rack level challenges related to power distribution, rack weight increase, additional switch needs, cabling increase etc.



MMIT





Case Studies

Open. Together.





Server eve configuration study

Characteristic of typical Open Rack server in Telco NFV use cases

- Mid/high range CPU's in use, Intel Xeon Scalable family SKU level is GOLD.
- Memory need is moderate, configurations typically below 384MB.
- Networking interface need is high, 100Gb/s per server normally implemented with 4x 25Gb/s interfaces.
- HW acceleration need varies from iNIC acceleration (e.g. vSwitch HW offloading) to purpose build accelerator add on cards
- Storage need is varying from compute server having only boot devices to storage unit with server having local high performance NVMe drives + high disc count JBOD combination.
- Server power consumption is 250W..300W in normal load (max. over 400W)





Nokia 2 OU Open Rack Server EMI shielding and cooling

As an enabler to use Open Rack V2 compatible servers in co-located sites Nokia has created easy to maintenance front panel design with full build-in EMI shielding.

Nokia AirFrame system components have regulatory certifications, CE/FCC and Safety certification.

In addition, the server is designed to meet NEBS requirements.

To enable up to 165W TDP processors and possibility to use high performance PCIe accelerator cards cooling fan performance is increased compared to typical Open Rack server designs.

ETSI temperature limits (up to 45 °C / 113 ° F) can be met with CPU SKU's used in typical blueprint telco NFV configurations.





Blueprint configuration of Nokia 2 OU Open Rack Server

- Purley processors 20 cores per socket and 125W TDP (up to 28 cores/165W allowed)
- Memory 192GB (up to 1.5TB)
- Ethernet connectivity 4x 25GbE with dual NIC
- Storage with 2x 2.5" U.2 NVMe/SATA devices and 1x M.2
- Storage interface HBA/RAID options •
- In-band / Out-of-band management HW management connectivity
- Optionally enables up to two 75W accelerators ٠

$\mathbf{D}\mathbf{P}\mathbf{E}\mathbf{\Pi}$ INSPIRED











- Processor TDP and core count must be reduced because of cooling capability limitation
 - 20 cores / socket 125W TDP in 40 °C/ 104 °F
 - 16 cores / socket 105W TDP in 45 °C/ 113 °F
- Memory up to 784GB with low CPU SKU's
- Ethernet connectivity 4x 25GbE with dual NIC is possible \bullet
- Storage needs to be redesigned, with EDSFF drives dual hot pluggable slots are possible
 - Availability of optimal capacity drives is unclear
- Storage interface HBA/RAID cannot be used at same time with dual NIC
- HW accelerator usage is not possible because of limited PCIe slots and cooling capacity
- Energy consumption is about 2%-3% higher because of smaller heat sinks and smaller less effective fans







Rack Eve configuration study



Possible equipping limitations in old Telco Sites

Site Physical Limitations

- Fully equipped Open Rack v2 weight is >800 kg \rightarrow >1200 kg/m2
 - Floor load capacity sometimes limits the configurations. •
- Old sites can also have limitations due to
 - Elevator capacity •
 - Delivery path height (door openings) •

Site Power and cooling limitations

- Old telco sites typically have -48VDC power feed infrastructure with battery rooms
- Site power budgets quite often limit size of installations
 - Limitation can be as low as 4 kW per rack \bullet
- Site cooling capacity quite often limits rack configurations
 - Limitation can be as low as 4 kW per rack

Above restrictions are not valid for new Telco sites used for central data centers





Nokia AirFrame Open Rack v2 variants

- Two independent power zones
 - One 12 VDC busbar per power zone
 - Up to 12,5kW per power zone
- Power shelf options
 - 208/415V AC, 25kW per rack
 - ETSI Negative 48V DC, 17.4kW per rack
- 4 OU for switches by default
- 42 OU rack variant: 36 OU for IT devices
 - Divided in 4 identical blocks for servers/ JBOD
- 36 OU rack variant: 30 OU for IT devices
 - Divided in 3 identical blocks for servers/ JBOD
- Four different block types defined
 - Compute block
 - 3x JBOD block
 - Dual 1x JBOD block
 - Compute + 1x JBOD block

| P | Compute Node | Compute Node | Compute Node | P | 42 41 | |
|------------|-----------------|------------------|--------------------|-----|----------|-------|
| 0 | Compute Node | Compute Node | Compute Node | 0 | 40 39 | Blo |
| 8 * | Compute Node | Compute Node | Compute Node | ~ | 38 37 | ck 4 |
| E | Compute Node | Compute Node | Compute Node | E | 36 35 | |
| 4 8 | | | | R | 34 | |
| 3 | | Power Shelf 2 | | | 33 | |
| 1 Z | Compute Node | Compute Node | Compute Node | z | 31 30 | |
| 9 O | Compute Node | Compute Node | Compute Node | 0 | 29 28 | Blo |
| 7 N | Compute Node | Compute Node | Compute Node | N | 27 26 | ck 3 |
| 5 E | Compute Node | Compute Node | Compute Node | E | 25 24 | |
| 3 2 | Leaf Switch | | | 2 | 23 | |
| 2 | | Cable managemen | 1t | | 22 | |
| | | Rack mechanic | | | | |
| | LeafSwitch | | | - 1 | 21 | |
| | | inagement switch | el agreente anno 1 | | 20 | |
| 3 | Compute Node | Compute Node | Compute Node | | 19 18 | |
| 7 6 | Compute Node | Compute Node | Compute Node | ľ | 17 16 | Blo |
| 5 | Compute Node | Compute Node | Compute Node | E | 15 14 | ck 2 |
| BR | Compute Node | Compute Node | Compute Node | R | 13 12 | |
| 1 | | 1 55. 17 | | | 11 | _ |
| z | Power Shelf 1 | | | z | 10 | |
| 9 | | 200 - N | 16- | ~ | 9 | |
| 8 O | Compute Node | Compute Node | Compute Node | 0 | 8 7 | |
| 5 N | Compute Node | Compute Node | Compute Node | N | 6 | Blo |
| 4 E | Compute Node | Compute Node | Compute Node | E | 4 | ock 1 |
| 2 1 | Compute Node | Compute Node | Compute Node | 1 | 2 | |
| | | | | | 1 | |







Rack level blueprint configurations with 2 OU servers

42 OU rack maximum configuration for Telco NFV use case is:

- Three Compute blocks
- Compute + 1x JBOD block
- Two Leaf switches
- One HW management switch
- Optionally single Spine switch \rightarrow reduces Compute units
- \rightarrow Up to 45 servers and one JBOD (up to 45 disks)

36 OU rack maximum configuration for Telco NFV use case is:

- Two Compute blocks
- Compute + 1x JBOD block
- Two Leaf switches
- One HW management switch
- Optionally single Spine switch
- \rightarrow Up to 36 servers and one JBOD or 39 servers

42 OU rack



NOKIA



Rack level configuration examples and challenges with 1 OU servers

Four different configuration cases analyzed

Case1: 42 OU rack with DC power: same CPU configuration

- 20 cores/125W TDP in 2 OU vs. 20 cores/ 125W TDP in 1 OU
- 1 OU configuration is DC power shelf limited \rightarrow same amount of servers can fit to • configuration in both 2 OU and 1 OU cases

Case2: 42 OU rack with ETSI (up to 45 °C / 113 ° F) environment limited CPU selections

- 20 cores/125W TDP in 2 OU vs. 16 cores/ 105W TDP in 1 OU •
- 1 OU configuration is DC power shelf limited \rightarrow less cores in rack with 1 OU server ۲
- Additional leaf switch pair required with 1 OU configuration ٠

Case3: 42 OU rack with AC power: max TDP processor configuration

- 28 cores/165W TDP in 2 OU vs. 20 cores/ 125W TDP in 1 OU •
- 1 OU configuration is rack power busbar limited \rightarrow less cores in rack with 1 OU server •
- Additional leaf switch pair required with 1 OU configuration •
- Rack weigh increase and additional cabling can be issues in Telco environments ٠

Case4: 36 OU rack with AC power: max TDP processor configuration

- 28 cores/165W TDP in 2 OU vs. 20 cores/ 125W TDP in 1 OU •
- 1 OU configuration is rack power busbar limited •
- Less servers and cores in rack with 2 OU server ۲
- Additional leaf switch pair required with 1 OU configuration ٠
 - © 2019 Nokia 13

| Case 1: | | | |
|------------------|--------------|-------------|-------------|
| - Ambient 40C | | 2 OU Server | 1 OU Server |
| - 42 OU Rack | Servers/rack | 48 | 48 |
| - Telco 48VDC | Cores/socket | 20 | 20 |
| - 8,7kW per zone | Cores/rack | 1920 | 1920 |

| Case 2: | | | |
|------------------|--------------|-------------|-------------|
| - Ambient 45C | | 2 OU Server | 1 OU Server |
| - 42 OU Rack | Servers/rack | 48 | 54 |
| - Telco 48VDC | Cores/socket | 20 | 16 |
| - 8,7kW per zone | Cores/rack | 1920 | 1728 |

| Case 3: | | | |
|-------------------|--------------|----------------------|------|
| - Ambient 40C | | 2 OU Server 1 OU Ser | |
| - 42 OU Rack | Servers/rack | 48 | 60 |
| - AC power | Cores/socket | 28 | 20 |
| - 12,5kW per zone | Cores/rack | 2688 | 2400 |

| Case 4: | | 1 | |
|-------------------|--------------|-----------------------|------|
| - Ambient 40C | | 2 OU Server 1 OU Serv | |
| - 36 OU Rack | Servers/rack | 36 | 54 |
| - AC power | Cores/socket | 28 | 20 |
| - 12,5kW per zone | Cores/rack | 2016 | 2160 |
| | | | |







Summary

Summary

Technically 1 OU variant of Open Rack server is possible

- New form factor SSD drives (EDSFF) enables sufficient local storage capacity
- 1 OU Server level cooling capability limits usage of highest performance CPU's in dual socket design \bullet
- There are also limitations in amount PCIe extension slots for storage interface or accelerators
- There is minor energy efficiency penalty because of smaller fans and heat sinks

Rack level configuration calculations shows only minor core count/ computing capacity differences between 1 OU and 2 OU server configurations

- With telco -48VDC power supply capacity limits server amount to same level
- With AC supply amount of servers can be significantly increased using 1 OU Open Rack servers
 - However effect to rack level core count is limited because 2 OU server can have higher core count per server
 - With 1 OU configurations additional Ethernet switches and cabling are needed \bullet
 - Rack level weight would be higher with 1 OU server and higher amount of servers, this can be issue with traditional Telco CO sites

Conclusion: 1 OU server benefits with current dual socket server design are limited. It will be option in future designs with new form factor NVME drives.







Thank You!



Come and visit us at Nokia booth B15

Open. Together.



OCP Global Summit | March 14–15, 2019



