



OPEN
Compute Project

WHITE PAPER: NEXT GENERATION OPEN EDGE SERVER DESIGNS

EDITOR(S):

Samuli Toivola, Nokia Solutions and Networks OY

Executive Summary

This document describes the main characteristics of future Open edge server design targets and key requirements. It provides server motherboard and barebone key design rules, connectivity and storage targets, power and cooling enhancements targets and chassis level dimensioning principles.

Table of Contents

Introduction	4
1 Next Generation Open Edge server development	5
Environmental and regulatory compliancy	5
Commodity development - enablers for future evolution	6
Server motherboard design requirements	7
Sled level power and cooling targets for 1U and 2U server sleds	9
Chassis level configuration rules	10
2 Conclusion	12
3 Glossary	13
4 References	13
5 License	14
6 About Open Compute Foundation	14
7 Appendix A. [Title]	15

Introduction

This document describes the main characteristics of future Open edge server design targets and key requirements. It provides server motherboard and barebone key design rules, connectivity and storage targets, power and cooling enhancements targets and chassis level dimensioning principles.

Key focus in server development is taking maximum advantage of processor and other semiconductor technology development and at the same time maintain 100% compatibility with other Open Edge ecosystem components like 2U/3U chassis, power supplies, RMC module and other sled types like Open Edge switch sled.

This document describes design targets for 1U and 2U server sleds for following key areas:

1. Top level environmental requirements for Open Edge solution
2. Commodity development – enablers for future evolution
3. Server motherboard and riser block diagram proposals
4. Sled level power and cooling targets for 1U and 2U sleds
5. Chassis level configuration rules

1 Next Generation Open Edge server development

This document describes design targets for 1U and 2U server sleds for following key areas:

1. Top level environmental and regulatory requirements for Open Edge solution
2. Commodity development – enablers for future evolution
3. Server motherboard design requirements
4. Sled level power and cooling targets for 1U and 2U sleds
5. Chassis level configuration rules

Environmental and regulatory compliancy

The key environmental and regulatory requirements for Open edge server sled will remain same as earlier product releases. Main standard compliancy requirements are listed in this chapter.

Operating conditions

- Operating temperature range: -5 C ...+45 C [ETSI EN300 019-1-3 Class 3.2]
- Short term operating temperature: -5 C to +55 C [GR-63-CORE]
- Operating humidity: 5 % to 95 %

EMC

- EN300386 (v1.6.1)
- FCC CFR47 15 (class A), CISPR 22/32 (class A) CISPR 24
- TEC/EMI/TEL-001/01/FEB-09 and TEC/IR/SWN-2MB/07/MAR-10
- GR-1089-CORE, and more

Safety

- IEC 62368-1:2014
- GR-1089-CORE (electrical safety, grounding and bonding)

Seismic tolerance

- GR-63-CORE (Zone 4)

Acoustic noise

- GR-63-CORE (equipment room criteria)

Fire resistance

- GR-63-CORE (shelf level criteria)

Commodity development – enablers for future evolution

The key technology evolution topics for server development in few coming years are listed in this chapter.

Processor development forecast

- Rapidly increasing core count, from 28 cores Cascade Lake (2019 design) about doubles or even more in 2022/2023 platforms
- Fast TDP increase because of increased core count and faster IO
- PCIe development PCIe Gen3 → Gen4 → Gen5
 - 4x speed
 - Lane count increase at the same time → Not that much design compromises needed

DDR Memory performance evolution

- From DDR4 with 2933MT/s (2019) to DDR5 with 5200MT/s (2022)
- Trend to increase memory channels

OCP mezzanine/NIC evolution

- New OCP NIC v3 gives hot-swappable option for NIC cards
- Increased NIC card TDP and cooling solution allows more powerful NIC components
- Increased area in faceplate allows bigger amount of physical interfaces

Storage evolution

- SSD/NVMe technology as an only option for local storage
- With dense E1.S EDSFF devices it is possible maintain required amount of hot-swappable local storage devices in small physical footprint
- M.2 is still valid option for internal storage devices

Server level power increase is challenge

- All this commodity development leads higher power density
- More focus to server cooling solution selections and system level configurations are needed

Server motherboard design requirements

Server sled main component design targets are listed in this chapter. There should be available two server form factors for different use cases:

- A 1U server sled enables a denser in terms of core count and system memory but has less space for I/O and storage.
- A 2U server provides more storage and I/O options due to increased height. Also, full length, dual width processing units, such as HW accelerators and GPGUs, should be supported in the 2U sled.

CPU complex

- Design shall support high TDP CPU's for minimum 300W
- PCIe technology upgrade to PCIe Gen 5 and eventually to Gen 6 in high throughput use cases:
 - PCIe slots
 - OCP NIC slots
- PCIe gen 4 can still be used for Storage interfaces
- PCIe gen 3 can be used for motherboard legacy resources to optimize product cost

DDR Memory

- To maintain best possible memory speed 1DPC design is preferred
- DIMM slot for available memory channels should be populated

- Motherboard design shall support highest memory speeds of selected processor technology

PCIe slots/OCP NIC slots

- All PCIe slots should have x16 connectivity
- In 2U server design cabled PCIe signals are preferred to lower motherboard and 1U server sled cost
- PCB material cost level shall be optimized using short routing for PCIe lanes

Storage devices minimum requirements

- Dual E1.S EDSFF devices for 1U server sled
- Six E1.S EDSFF devices for 1U server sled
- Dual M.2 for internal operating system storage devices

Proposed high level server sled block diagram shown in in *Figure 1*

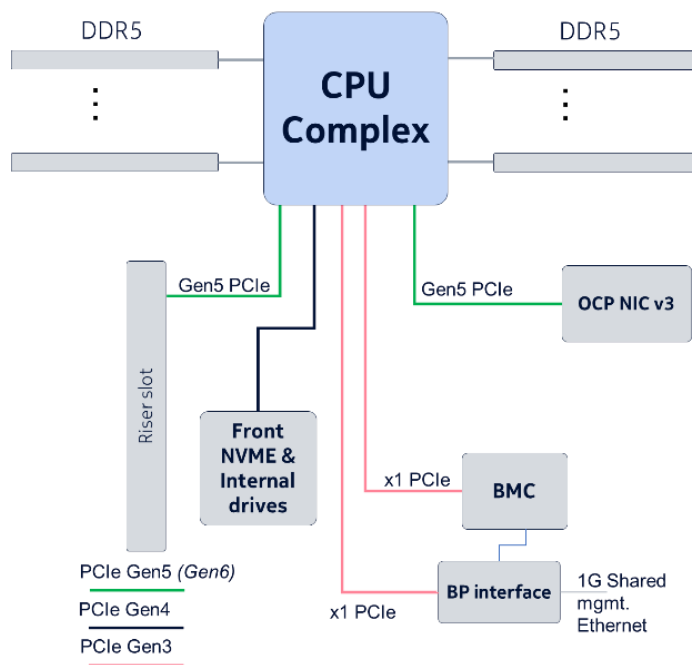


Figure 1 Simplified Block Diagram of Next Gen Open Edge Server sled

Sled level power and cooling targets for 1U and 2U server sleds

Proposed main component placement of an Open edge server sled are illustrated in Figure 2.

This proposal focuses locations of motherboard main resources like CPU, DIMMs, OCP NIC and E1.S SSDs. For CPU heat sink and extended heat sink locations are illustrated in drawing. Design target should be maximizing heat sink capability in quite limited physical footprint. To allow this also PCIe cabling connector location for 2U sled should be selected so that extended heat sink size can be maximized.

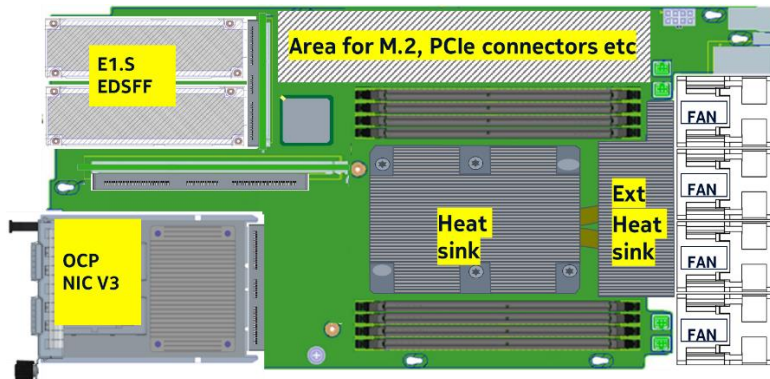


Figure 2 Proposed Open edge Server sled main component placement (1U server sled)

Here are listed key minimum cooling requirements to meet future needs of Open Edge server:

- Airflow targets with fans at maximum speed
 - 1U server sled 100CFM
 - 2U server sled 160CFM
- Both Front to Rear and Rear to Front airflow options shall be supported
- Power budget for server
 - 1U server sled 500W
 - 2U server sled 800W

Chassis level configuration rules

Open edge chassis mechanical design is simple, consisting of only few components. The power distribution board, backplane assembly and mounting brackets are attached to a steel enclosure.

Each sled in chassis configuration can be independently selected. Configuration can consist of 1U and 2U servers, switches and other type of sleds. Every chassis configuration shall have RMC and PSU's (100-240VAC and -48VDC options are available).

The different configuration options of an Open edge Chassis are illustrated in Figure 3.




Figure 3 Open edge chassis configuration options

The Open edge chassis supports flexible configuration of two dimensions of sleds, 1U and 2U. One 2U sled can be installed in place of two 1U sleds. 1U sleds in the top row have support brackets on the inner sides of the chassis. When a 2U sled is installed, the support brackets are removed. Removal and installation is tool-less.

- Supported configurations in 3U chassis are

- 5x 1U
- 3x 1U + 1x 2U
- 1x 1U + 2x 2U
- Supported configurations in 2U chassis are
 - 3x 1U
 - 1x 1U + 1x 2U
- A support bracket for 1U sled is removed when installing a 2U sled (tool-less)
- Principle about equipping of different sled types is presented in Figure 4




Figure 4 Open edge chassis sled equipping options

Sled level maximum power will increase because of higher power CPU's, PCIe bus evolution and introduction of DDR5 memory technologies. However, there are wide variety of Open Edge use cases and power level of single server configuration can be from 200W up to 500W(1U) or up to 800W(2U).

Use of the top server SKUs means also increase in chassis level power consumption and requires that dimensioning of configuration is considering chassis and PSU maximum power levels.

System level configuration rules and flexible use of 2U and 3U chassis enables continuation with existing chassis and power distribution designs. Table below shows few workable configurations using 3U and 2U chassis variants with 2kW Open Edge PSU variant.

Table 1 Configuration examples with sled average power levels

Configuration	1U sled power	2U sled power	Total Power	Notes
3U Chassis 1x 1U + 2x 2U	500W	750W	<i>2kW</i>	Minor limitation for average power needed
3U Chassis 4x 1U + 1 filler sled	500W	NA	<i>2kW</i>	1 sled place left empty
3U Chassis 5x 1U	400W	NA	<i>2kW</i>	Average power limited to 400W
2U chassis 1x1U + 1x2U	500W	800W	<i>1300W</i>	
2U Chassis 3x1U	500W	NA	<i>1500W</i>	

2 Conclusion

Key focus in future server sled designs shall be taking maximum advantage of processor and other semiconductor technology development and at the same time maintain 100% compatibility with other Open Edge ecosystem components like 2U/3U chassis, power supplies, RMC module and other sled types like Open Edge switch sled.

Sled level maximum power will increase because of higher power CPU's, PCIe bus evolution and introduction of DDR5 memory technologies. In new generations of server design shall tolerate high power components and be able to cool those in harsh environmental conditions.

However, there are wide variety of Open Edge use cases and power level of single server configuration can be from 200W up to 500W(1U) or up to 800W(2U). System level configuration rules and flexible use of 2U and 3U chassis enables continuation with existing chassis and power distribution designs.

In coming generations OCP NIC V3 enables higher networking bandwidths and dense EDSFF E1.S allows high performance storage in limited physical real estate.

3 Glossary

OCP: Open Compute Project

RMC: Rack Management Controller

4 References

[1] Nokia Open edge chassis specification, 2018

<https://www.opencompute.org/documents/open-edge-chassis-ocp-contribution-v1-3-pdf>

[2] Nokia Open edge RMC specification, 2019

<https://www.opencompute.org/documents/open-edge-rmc-ocp-contribution-v0-2-2-pdf>

[3] Nokia Open edge server specification, 2018

<https://www.opencompute.org/documents/open-edge-server-ocp-contribution-v1-0-pdf>

[4] Nokia Open edge switch specification, 2020

<https://www.opencompute.org/documents/open-edge-switch-specification-v0-5-2-pdf>

5 License

© Copyright Nokia Solutions and Networks Oy 2021 All rights reserved.
This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

6 About Open Compute Foundation

The Open Compute Project Foundation is a 501(c)(6) organization which was founded in 2011 by Facebook, Intel, and Rackspace. Our mission is to apply the benefits of open source to hardware and rapidly increase the pace of innovation in, near and around the data center and beyond. The Open Compute Project (OCP) is a collaborative community focused on redesigning hardware technology to efficiently support the growing demands on compute infrastructure. For more information about OCP, please visit us at <http://www.opencompute.org>

7 Appendix A. [Title]