WHITEPAPER

PENGUIN SOLUTIONS

Open Compute Project: Benefits and Innovation

OVERVIEW

This White Paper will explain the mission and principles of the Open Compute Project® (OCP), how it works to overcome current design limitations for computing hardware and software, and the benefits of OCP over EIA standards. It will also detail how Penguin Solutions™ has HPC and AI solutions that align with OCP principles.

WHAT IS THE OPEN COMPUTE PROJECT?

The Open Compute Project (OCP) is a collaboration among leaders in hyperscale data center operators, colocation and telecom providers, and enterprise IT users to create new open solutions for the industry. OEMs and engineers work together to redesign hardware to support the evolving demands of the compute infrastructure.[1]

OCP works with vendors to develop and commercialize innovation in a broad range of areas that impact the IT cloud ecosystem, including:

- **Artificial intelligence and machine learning**
- **Optics**
- **Sustainable data centers**
- **Advanced power management**
- **Cooling techniques**
- **Composable silicon**

THE OCP MISSION AND VISION

The OCP's mission is to spur rapid innovation to enable the mainstream delivery of the most efficient designs for sustainable and scalable computing. They look for projects that will meet at least three of their five core tenets:[2]

EFFICIENCY

Some examples of efficient design from OCP contributions include:

- **Thermal efficiency**
- **Reduction in infrastructure costs**
- **Reduction in latencies**
- **Platform performance, such as per W**

PENGUIN™ SOLUTIONS

IMPACT

OCP looks for projects that will create meaningful and lasting impact, such as building on top of OCP solutions to improve the value and realize efficiency gains. Contributions for members also help enable a more robust supply chain with alternate solutions.

OPENNESS

OCP contributions should be open to benefit stakeholders throughout the industry. While not every project can be 100% open source, every project should strive to comply with a set of already existing open interfaces.

SCALABILITY

Technology should be designed for large-scale deployments with an emphasis on future scalability. This includes features such as:

- **Maintenance**
- **Remote management**
- **Upgradability**
- **Error reporting**

SUSTAINABILITY

Contributions also must be sustainable with transparency of the environmental impact. This includes the responsible use of natural resources and reducing energy, water, and material costs.

LEARN MORE

See related solutions at https://www.penguinsolutions.com/computing/products/ocp-hpc-and-ai-systems/

For pricing on your specific design needs, contact a representative by email at

sales@penguincomputing.com or call 1-415-954-2800.

PENGUIN™ SOLUTIONS

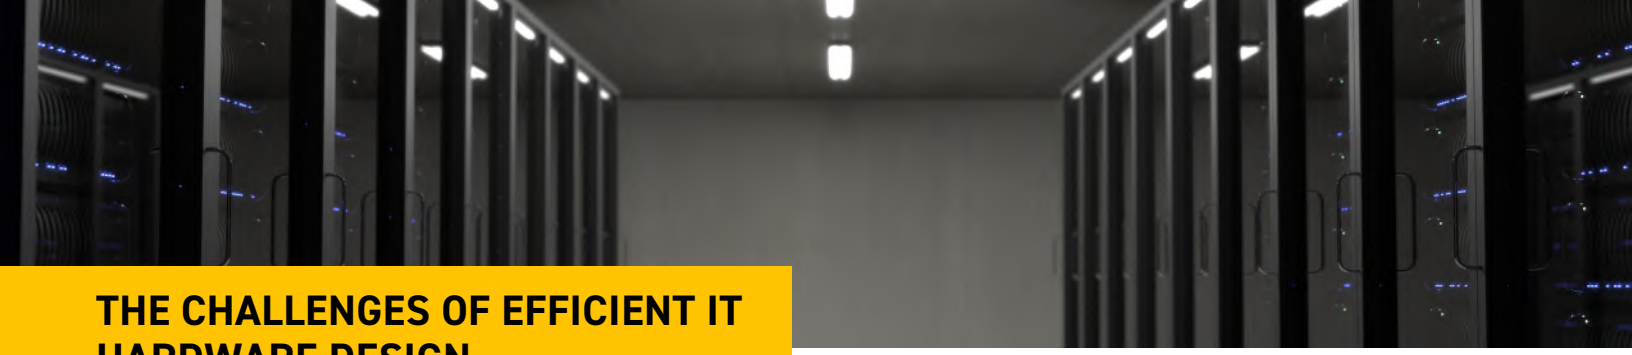

THE CHALLENGES OF EFFICIENT IT HARDWARE DESIGN

Data center and cloud services continue to expand at a significant rate. By 2025, Gartner reports that 70% of organizations will deploy infrastructure automation, with nearly 40% of on-prem computing and storage moving to an as-a-service delivery.[3] This exceptional demand requires new solutions to address compute power, scalability, form factor, and power density challenges.

While manufacturers strive for efficient design, there are existing limitations. As engineers try to wedge more circuits into exceptionally small sections of microchips, the power density produces significant heat. This can be inefficient and harmful to components.

Common design challenges for data centers and HPC computing include the increasing power requirements for CPUs, ASICs, and AI processors requiring high constant (TDC) and peak current. With mass amounts of data — and a constantly growing volume of data — moving through edge data centers, efficiency demands more processing power while the space to do so often remains the same.

Load lines are even tight as high-performance processors shift to smaller nodes to increase processing power requiring exceptionally tight voltage tolerances to maintain high power density without loss of power or performance.

For data centers, space and energy efficiency is crucial. Data centers already account for about 2% of the total electricity use in the country, consuming between 10 and 50 times the energy per floor space of a typical office space.[4] Increases in power density also increase the challenges of heat removal within these spaces.

Projects like OCP, the National Data Center Energy Efficiency Program, and the Better Building Alliance, work to overcome these challenges.[5,6]

PENGUIN™ SOLUTIONS

HOW OCP DESIGNS BENEFIT THE INDUSTRY

Here are some of the ways OCP designs benefit the industry as a whole.

OPEN HARDWARE DESIGN

One of the primary solutions proposed by OCP is an open hardware design that can be shared freely, allowing different companies to collaborate on more efficient designs. Open-source hardware standards allow for greater compatibility between components, allowing the use of best-of-breed solutions rather than relying on proprietary hardware.

One example is the Open Rack design, providing a standardized platform for data center hardware to allow for more efficient cooling, power distribution, and maintenance. Modular OCP server design removes front bezels to create front serviceable units with more efficient cooling. Panel mount 12V or 48X power connectors click into the power busbar directly on the back of the open rack. Larger heat sinks and more chassis space provide more efficient heat removal. Such design changes result in significant benefits, such as:[7]

- **Four times faster completion of required hardware tasks**
- **65% more servers handled per operator**
- **61% increase in employee productivity**
- **38% less time required to resolve unplanned downtime**

Another example is the importance of OCP Mezzanine for service design. Mezzanine PCIe expansion cards optionally include NC-SI management, reducing cables required, and providing better optioning of compute elements networking.[8] OCP multi-sled shelves support computer element expansion without adding additional expensive powered chassis.

EFFICIENT POWER USAGE

Reducing power consumption and heat transfer is another significant goal, both when systems are in use and when idle. For example, the OCP's Power Shelf is a high-density 48V DC power distribution unit that can provide up to 15kW of power to IT equipment.

The team at Penguin Solutions has worked with multiple rack and power supply vendors to extend and optimize these initial OCP capabilities for the highly demanding HPC environment.

Princeton University researchers say the power density in today's microprocessors rivals the interior of a nuclear reactor.[9] In a paper they presented at the Open Computer Project annual conference, researchers demonstrated a new design structure, rethinking power delivery and using capacitors instead of magnetics to process power. Researchers designed systems vertically rather than horizontally, delivering a ten times smaller system than the best off-the-shelf products without sacrificing speed or efficiency.[10]

Traditional EIA power distribution units (PDUs) create design challenges as well, including power quantization due to low AC voltage 208V and small breaker limitations. This makes it difficult to distribute and aggregate power within each rack since power is available and consumed in discrete steps, making it difficult to optimize. A simple power supply failure may lead to a catastrophic failure of the power within a rack.

OCP 48V DC busbars allow large shared redundant power supplies to distribute power at the scale of the consuming compute elements, minimizing quantization challenges.

IMMERSION COOLING

Mitigating heat from increased power density is a challenge for component design and data center operators. The OCP has developed an open immersion cooling system that can be used with various hardware designs. One example is the OCP's immersion cooling solution for Microsoft's Project Natick, which involved submerging an entire data center in the ocean to reduce energy usage.[11]

PENGUIN™ SOLUTIONS

Network Linux is an open-source software platform that can be used to run network switches.

MODULAR HARDWARE DESIGN

The OCP has also proposed the use of modular hardware designs, which can be scaled up or down depending on the user's needs. This allows for more efficient use of hardware and reduces waste. An example of this is the Open CloudServer design, a modular hardware platform that can be used for various cloud-based services.[12]

OPEN SOFTWARE DESIGN

Besides open hardware design, the OCP has also proposed using open software design, which allows for more collaboration and innovation in the IT hardware industry. One example of this is the Open Network Linux operating system, which is an open-source software platform that can be used to run network switches.[13]

CONTINUED INNOVATION

Since OCP principles encourage collaboration among developers, vendors, and users, this leads to continuous innovation, creating new hardware designs and technology that are increasingly more efficient, reliable, and cost-effective for users.

PENGUIN™ SOLUTIONS

THE DIFFERENCES BETWEEN OCP AND EIA

The scope of the OCP is focused on open-source designs for data center and IT hardware, emphasizing efficiency, scalability, and sustainability. The Energy Information Administration (EIA), on the other hand, provides standards for a broader range of electronic and electrical equipment, including consumer electronics, telecommunications equipment, and more.[14]

Key differences between OCP and EIA include scope, collaborators, design standards, and efficiency.

Scope:

The scope of the OCP is focused on open-source designs for data center and IT hardware, with an emphasis on efficiency, scalability, and sustainability. Conversely, the EIA provides standards for a broader range of electronic and electrical equipment, including consumer electronics, telecommunications equipment, and more.

Community-driven vs. Industry-driven:

The OCP is a community-driven organization that relies on collaboration and contributions from its members to develop and refine hardware designs. The EIA, on the other hand, is an industry-driven organization that includes member companies from various sectors of the electronics industry.

Open vs. Proprietary:

The OCP focuses on open-source designs that members can freely share and modify. The EIA, on the other hand, provides standards for proprietary technologies that are developed and owned by member companies.

Emphasis on Efficiency:

The OCP flexibility of the hardware design focuses on reducing energy loss and increasing density and power capacity while improving cooling efficiency. The EIA also considers energy efficiency in its standards but may not have the same level of focus as the OCP.

PENGUIN™ SOLUTIONS

THE BENEFITS OF OCP OVER EIA

We believe there are clear benefits for all stakeholders in deploying OCP principles to create greater efficiency and flexibility in compute, power, networking, and cooling designs. Here are some of the key reasons why engineers today prefer OCP over EIA designs.

POWER DISTRIBUTION

PDUs on EIA racks have three phases, often with two breakers per phase. These tend to leave some watts on each phase or breaker, effectively "stranding" power that cannot be used. OCP design eliminates this by putting all of the power onto the busbars. 48V DC busbars increase the efficiency of power, avoiding stranded power, and reducing energy usage with a simple, high-capacity internal power distribution design.

SERVICE AND INSTALL DESIGN

Most components are cabled and can be serviced on the cold aisle, making installation and maintenance easier. The exception would be liquid-cooled solutions.

Power, rack, and network infrastructure can also remain in place while compute or GPU sled are periodically refreshed.

RACK DENSITY

High rack density with OCP allows data centers to significantly reduce floor space, allowing for greater scalability, and faster deployment. Racks are available in flexible configurations, including:

- **1 OU full width**
- **1 OU 1/3 width**
- **1 OU 1/3 width**
- **2 OU 1/3 width**
- **2 OU full width**

INCREASED MEAN TIME BEFORE FAILURE (MTBF)

Nodes using the OCP design do not require small point power supplies, which eliminates a common point of failure. Sled-based architecture also provides faster time to repair (TTR) and easier servicing from the cool aisle-focused design strategy.

PENGUIN™ SOLUTIONS

BETTER VALUE

Optimized OCP systems provide better value than EIA alternatives in several ways, such as:

- **EIA design typically requires the use of proprietary hardware, which can drive higher initial costs. OCP allows for greater flexibility (and scalability), reducing upfront hardware costs.**
- **OCP designs are more energy efficient. Improved thermal management can reduce cooling costs for data centers.**
- **EIA hardware may require more energy and cooling.**
- **EIA hardware may also require specialized maintenance and support, which can increase operational costs.**

LOWER TOTAL COST OF OPERATION

OCP designs produce a lower total cost of operation (TCO).

Optimizing cooling enables peak efficiency of CPU and GPU elements. This increased density and scale reduces data center facilities with several flexible solutions, including:

- **Higher power connections are available to 277V/480V 60A.**
- **Water cooling options reduce the need for computer room air conditioning (CRAC) and computer room air handling (CRAH).**
- **Rear-door cooling options reduce data center air cooling limits per rack.**
- **Oil Immersion options enable fan-less compute.**

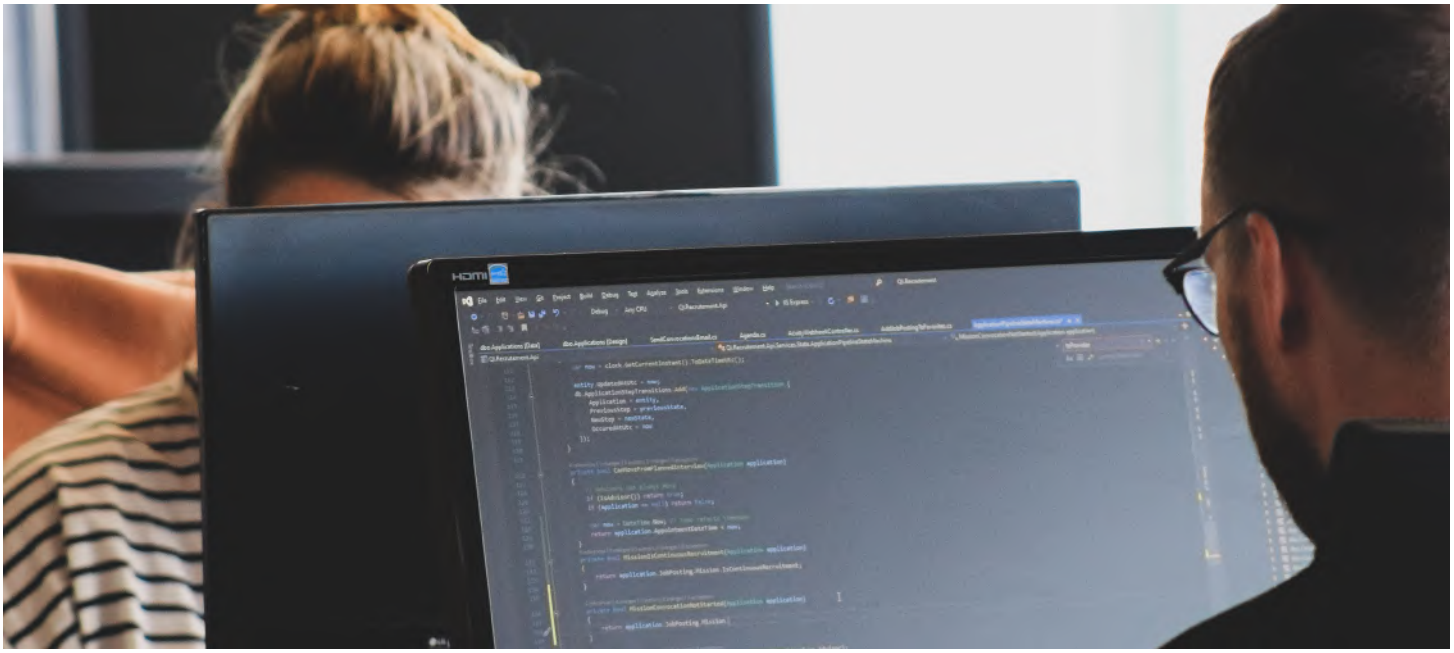

PENGUIN™ SOLUTIONS

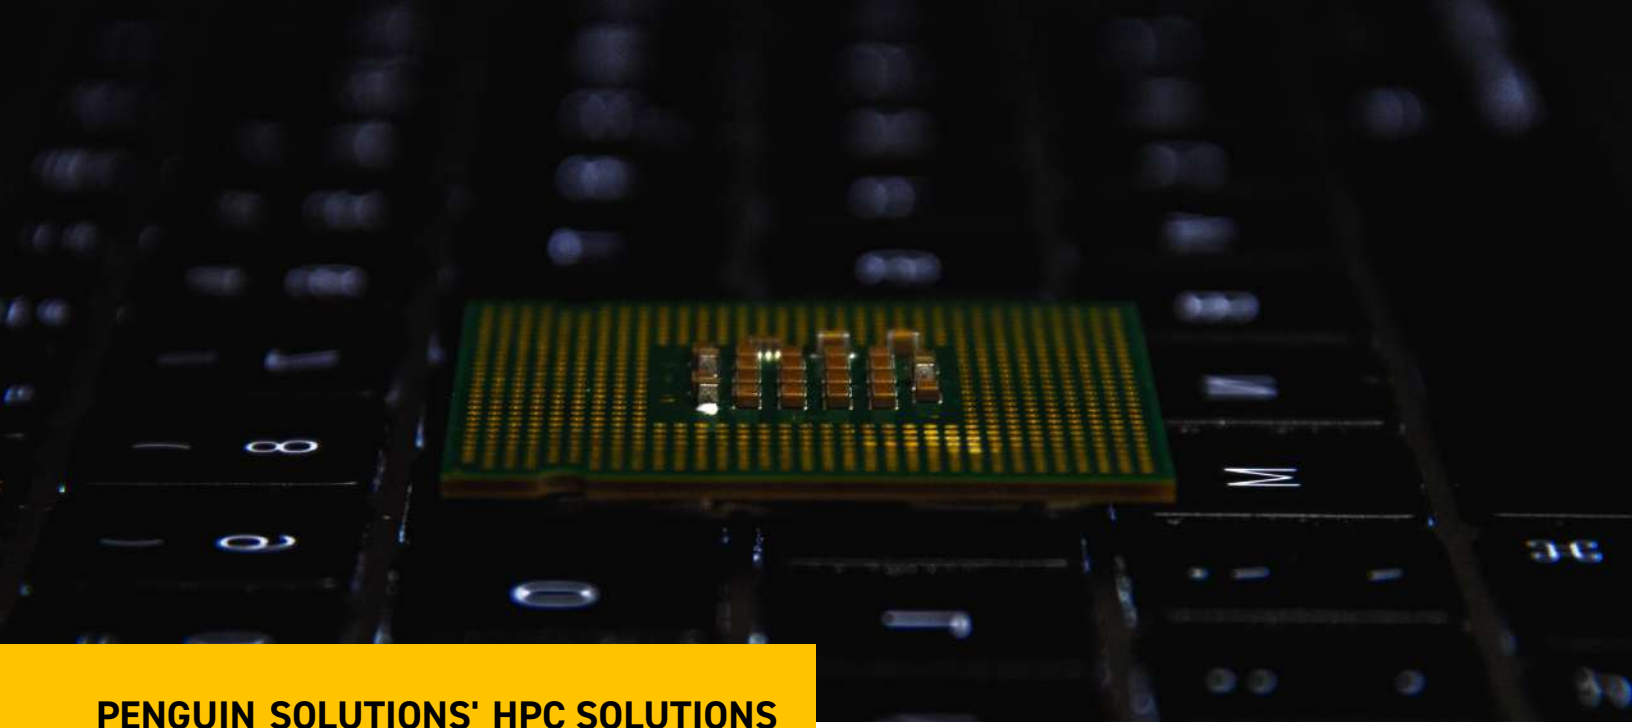

PENGUIN SOLUTIONS' HPC SOLUTIONS ALIGN WITH OCP PRINCIPLES

Penguin Solutions' HPC solutions align with OCP principles with a focus on energy efficiency, scalability, sustainability, and open-source solutions.

Energy Efficiency:

Efficient cooling technologies and power management tools reduce energy consumption and improve system performance.

Scalability:

Using modular hardware architectures and open-source software that can be easily customized and expanded over time.

Sustainability:

Penguin Solutions' HPC solutions are designed to be as energy-efficient as possible, reducing the environmental impact of their turn-key systems.

Open Source:

Penguin Solutions uses a range of open-source technologies, including the Slurm workload manager, OpenMPI message passing interface, and various Linux operating systems. This allows users to take advantage of a wide range of community-driven tools and technologies, which can help reduce costs and improve performance.

PENGUIN SOLUTIONS™

HOW OCP DESIGNS BENEFIT THE INDUSTRY

Penguin Solutions occupies a unique space in the data center and high power computing industry. Few vendors can integrate all of the required components to provide a single supplier support solution that aligns with OCP principles.

As an early adopter and supporter of OCP technology, Penguin Solutions designs OCP systems and rack scalable units customized for customer applications, ranging from a single rack to acre-sized data centers. This allows an organization to deploy compute resources at an industrial scale using a flexible design aligned with their unique needs.

TUNDRA AP

The Penguin Computing Tundra® platform combines the capital and operating savings of OCP-based hardware while utilizing the latest and most-efficient processes for HPC and AI applications.

Tundra AP is the latest generation supercomputing platform, combining the processing power of Intel® Xeon® scalable 9200 series processors and Penguin Computing's Relion® XO1122eAP server in an OCP form factor that delivers a high density of CPU cores per rack.

The Tundra AP platform:

- **Leverages power of Intel® Xeon® Scalable 9200 Series processors in an OCP form factor**
- **Delivers higher node density**
- **Enables greater power capacity**
- **Improves serviceability**
- **Provides room-neutral, integrated, direct-to-chip cooling**
- **Lowers total cost of ownership**

TUNDRA EXTREME SCALE (ES) FOR HPC

The second-generation Tundra platform is a complete HPC system that is dense enough for the most challenging projects and flexible enough for nearly any HPC architecture.

Tundra ES for HPC supports:

- **A diverse array of technologies, including a graphics processing unit (GPU)-accelerated computing on the latest NVIDIA® graphics accelerators**
- **Server formats from 10U-40U with a capacity for more than 100 nodes per rack.**
- **The latest AMD EPYC™ processors or Intel® Xeon® Scalable processors, high-speed software-defined networking (SDN), and localized storage for flexibility and performance.**

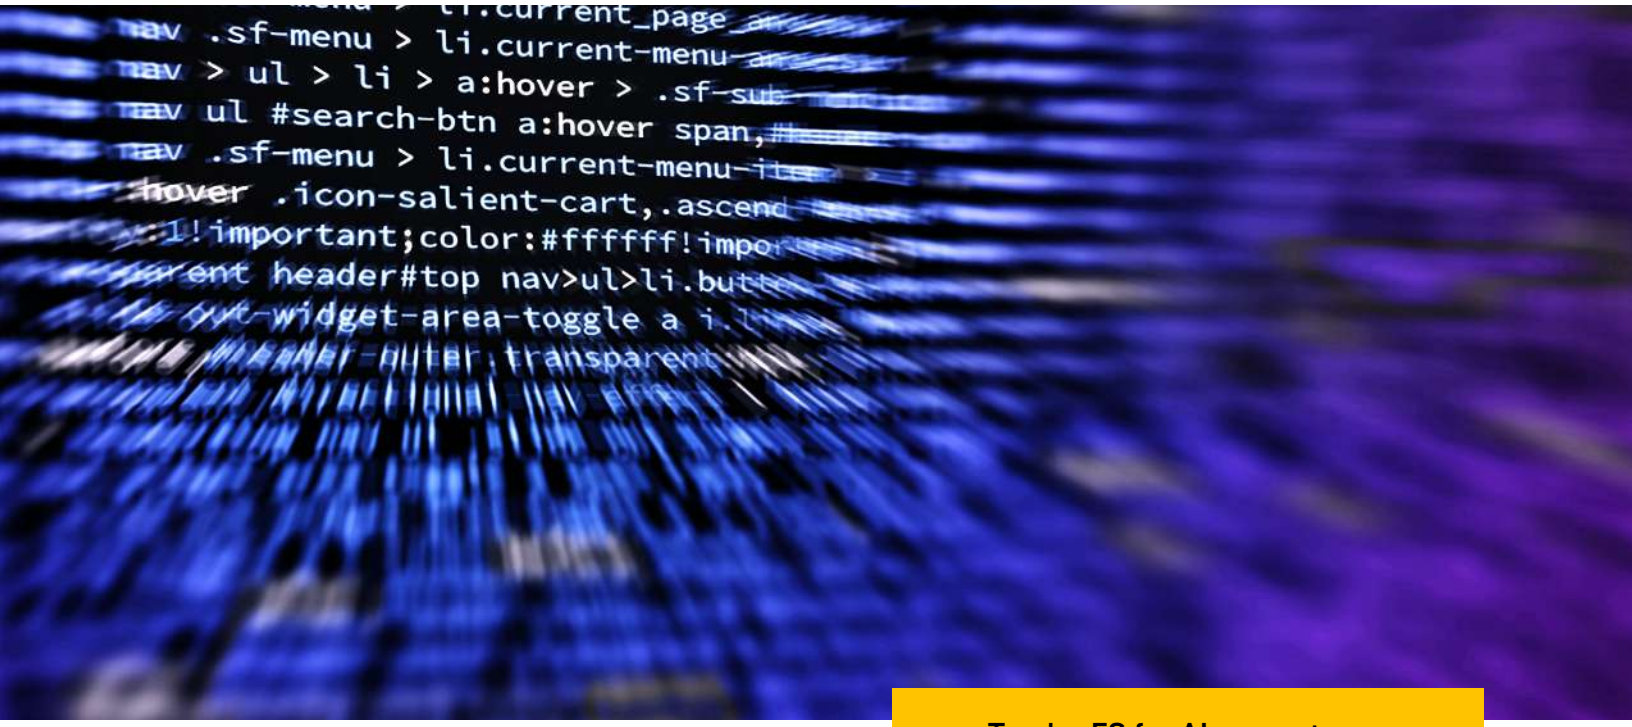

Tundra ES for AI supports:

Tundra Extreme Scale (ES) for AI

Tundra ES for AI is a reference design to support evolving AI developments leveraging OCP's low TCO and enabling mass scale-up as needed. Optimized to support the technology required for inference workloads, Tundra ES for AI is dense enough to enable more high-value technology per rack than a traditional solution and a more diverse compute array than most OCP designs.

- **An exceptionally diverse array of technologies, including the NVIDIA T4 with Turning Tensor Cores for inference.**
- **Server formats from 10U-40U with a capacity for over 100 nodes per rack.**
- **Comes with the latest AMD EPYC processors or Intel Xeon Scalable processors, high-speed software-defined networking (SDN), and localized storage for flexibility and performance.**

PENGUIN™ SOLUTIONS

OCP SERVERS AND STORAGE

The Tundra ES architecture provides a flexible infrastructure for workload-optimized clusters and virtualized server farms.

A variety of compute, application-optimized, and OCP servers and storage options are available, including a dual-processor, high-performance x86 server, an Intel Xeon Phi server, and a 4-drive storage server optimized for Swift or Ceph storage architectures.

Based on Intel Xeon Scalable processors with up to 16 cores per processor, servers deliver maximum power efficiency, serviceability, and cost efficiency.

Servers and storage for Tundra and Open Rack are available in a variety of 10U and 20U vanity-free form factors.

10U Servers

- **Relion® XO1122eAP**
- **Relion XO1132e**
- **Relion XO1132g**
- **Altus® XO1232g**

20U Servers

- **Relion XO2132**

PENGUIN™ SOLUTIONS

DATA CENTER SOLUTIONS

Penguin Solutions also offers several data center solutions. Whether you plan to deploy to your own data center, incorporate a hosted environment, grow your edge footprint, or leverage cloud computing with fully-managed, private Infrastructure-as-a-Service, we can help.

From a single rack to infrastructure spanning multiple data centers, our experienced architects and engineers will work closely with you to design and implement a robust, efficient data center solution.

Data center services include:

- **Design consulting: Penguin Solutions can design the ideal data center, hosted, or hybrid solution, creating a customized solution for your unique infrastructure and operational requirements.**
- **Project management: Penguin Solutions can provide expert project coordination and oversight for managing the most complex projects. Professional services include expert configuration and deployment to get you up and running immediately.**
- **Hosting: Penguin Solutions can deliver a wide variety of co-location hosting options, leveraging strong partnerships with world-class data center providers**
- **Fully-managed hosting: Penguin Solutions also offers a combined hosting and managed service as a complete as-a-Service solution for your organization.**

PENGUIN™ SOLUTIONS

Conclusion

Open compute standards developed through collaboration between users and suppliers offer best-in-class customer value. Customers receive a reliable, cost-effective system that can be deployed faster than constrained EIA alternatives.

OCP design solutions enable greater innovation and flexibility not envisioned by legacy EIA fixed-function servers. This allows you to build a reliable, sustainable, and scalable platform you can depend on for an internal HPC cloud.

You can trust the experts at Penguin Solutions to design a customized IT infrastructure for your HPC and AI needs that embodies the OCP design principles for optimal performance. From project discovery to ongoing maintenance and future planning, Penguin Solutions takes a holistic and agile approach to customer service. Whether you need help planning a new project from start to finish or augmenting your existing staff to meet a specific need, the Penguin Solutions services team is ready to implement and support the solution that fits your needs and budget.

Call the OCP experts at Penguin Solutions today at +1-415-954-2800 or contact us online to discuss your needs.

Sources

1. https://www.opencompute.org/
2. https://146a55aca6f00848c565-a7635525d40ac1c70300198708936b4e.ssl.cf1.rackcdn.com/images/bf648bb75091907147e76846cad590f402660d2e.pdf
3. https://www.gartner.com/en/articles/4-predictions-for-i-o-leaders-on-the-path-to-digital-infrastructure
4. https://www.energy.gov/eere/buildings/data-centers-and-servers
5. https://www.energystar.gov/ia/partners/prod_development/downloads/NDCFactSheet.pdf?18d4-b6e9
6. https://betterbuildingssolutioncenter.energy.gov/alliance
7. https://www.opencompute.org/files/Experiences-from-OCP-gear-deployments-with-Telcos.pdf
8. https://www.opencompute.org/documents/facebook-ocp-mezzanine-20-specification
9. https://ece.princeton.edu/news/vertical-thinking-broke-bottleneck-powering-high-performance-computers
10. https://www.princeton.edu/~minjie/files/baek-LEGO-2021.pdf
11. https://natick.research.microsoft.com/
12. https://www.penguinsolutions.com/computing/products/servers/ocp-servers-storage/
13. https://www.opencompute.org/projects/open-network-linux-onl
14. https://www.eia.gov/