intel.

Zeblok Computational, Advantech, and SmartPoint.io Unlock Edge AI's Potential

Powered by Intel® Xeon® Scalable processors, Zeblok Ai-MicroCloud™ deploys on Advantech and SmartPoint.io hardware to enable seamless edge-to-cloud AI services

Service providers are facing the next great opportunity in digital transformation. Soon, they will more easily attract and retain customers and increase profitability by offering next-generation artificial intelligence (AI) services as part of their product portfolio. Likewise, AI-driven insights are an essential part of the modern enterprise's digital assets. When combined with 5G Open RAN and multi-access edge computing (MEC), AI has the potential to transform business across industries.

For service providers, this presents a new opportunity. Those ready to deliver seamless edge-to-cloud AI solutions through a vibrant ecosystem will thrive in the upcoming era of Digital Transformation 3.0.

Challenge: Simplifying and accelerating complex edge AI deployments

The creation of a rich AI offering is an ecosystem play that requires a significant effort from operators and service providers. A rich ecosystem of dedicated partners is necessary to bring the diverse set of technologies and vendors together to deliver end-to-end AI solutions. Whether it's an edge or enterprise AI deployment, delivering end-to-end AI solutions requires seamless integration between cloud service providers (CSPs), communications service providers (CoSPs), managed service providers (MSPs), edge data center and network operators, original equipment manufacturers (OEMs), and the independent software vendors (ISVs) that develop the AI algorithms and AI applications. To glue it all together, middleware is a necessity.

Oftentimes, service providers and enterprises that want to deliver edge AI solutions need to create a variety of ad hoc AI assets for each scenario: public safety, smart city, Industry 4.0, etc. Bespoke integration involves working with different AI vendors and developing microservices for each ISV to deliver their solution, but this approach lacks standardization and scalability. Engineering efforts to integrate each new AI solution could take six to eight months, dramatically increasing costs and time to market.

Additionally, new smart city, mobility, and retail edge AI use cases demand edge computing that can support low-latency requirements. These micro data centers at the edge are a greenfield opportunity for communications infrastructure operators and service providers. Low-latency access to analytics is critical to serving the needs of the emerging market segment. But these advanced AI deployments require powerful performance and flexibility from their hardware to deliver full value.

Solution: A turnkey platform for end-to-end AI

To unlock edge AI opportunities, the Zeblok Ai-MicroCloud™ turnkey AI Platform as a Service paired with Advantech's SKY-8000 Edge Servers helps streamline deployment, shorten time to market, and optimize performance for new AI solutions.

Zeblok Ai-MicroCloud is an AI platform-as-a-service that streamlines AI delivery by providing a single, cohesive, turnkey, cloud-native AI environment—acting as end-to-end AI middleware that unifies the development, testing, training, and deployment of AI/ML solutions, from core to edge. Zeblok Ai-MicroCloud enables an AI- API economy by providing the Zeblok Ai-AppStore for AI algorithms and a mechanism to produce AI inference as a secure Ai-API™, then distribute it from cloud to edge rapidly and cost-effectively.

To support Zeblok Ai-MicroCloud, Advantech SKY-8000 Edge Servers provide a powerful and fully integrated edge computing platform that can be deployed closer to citizens, enterprises, or subscribers. Based on Intel® Xeon® Scalable processors, these servers have been optimized for sustained high computing performance in the harsh edge environments that are inherent to deploying AI inference at edge endpoints. They can be deployed in micro data centers at the very edge, including in IP65 pole-mounted boxes, curbside cabinets, and kiosks. Their high-density and compact design, combined with their ruggedness and advanced reliability, make these servers the platform of choice for 5G, MEC, and AI use cases that demand high performance, high availability, and high efficiency.

To fully unlock AI's potential, the solution has been validated to be deployed on a SmartPoint.io smart kiosk: a data center on the sidewalk that unlocks new edge AI opportunities in urban environments.

Through Zeblok Ai-MicroCloud, the solution creates an open cloud-to-edge ecosystem that allows enterprises to build their own AI assets and deploy a secure Ai-API that can be distributed from core to edge, lowering the overall cost per insight. By employing a standard certification methodology, this solution can be seamlessly scaled to thousands of micro multi-access edge computing (MEC) data centers.
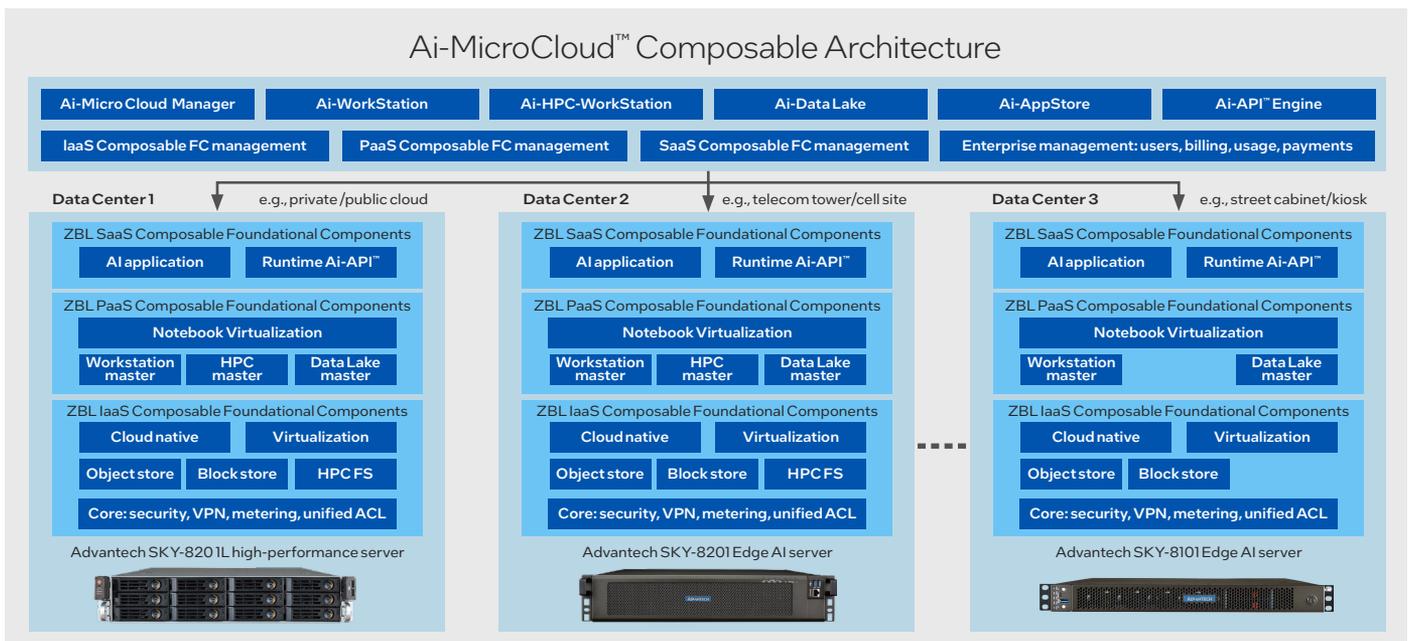
How it works

Zeblok's Ai-MicroCloud cloud-native architecture leverages popular open source frameworks to deliver a comprehensive platform. Advanced extensions improve simple Kubernetes orchestration, compute, and network virtualization, with high performance computing (HPC) orchestration. Key innovations include significant shrink-wrapping of cloud-based technologies as AI middleware, deployable in various environments, including kiosks, MEC hubs, public clouds, on-premises data centers, and different hardware platforms.

Ai-MicroCloud aggregates several composable foundational components, which deliver ML DevOps, Ai-WorkStations, software distribution via Ai-API deployment, and workflows enabling socket-specific optimization of AI models. By integrating developer-friendly model training and distribution workflows, with tools such as the Intel Distribution of OpenVINO™ toolkit and the Intel® oneAPI toolkits, Ai-MicroCloud provides Ai-Optimization-as-a-Service™ for enterprises.

Ai-MicroCloud is the only platform needed to go from development to testing, beta, and production delivery of runtime inferences optimized for heterogeneous chipsets (CPUs, GPUs, and FPGAs)—all secured behind an SSL-protected endpoint, delivered as a service.

Ai-MicroCloud™ Composable Architecture

| Ai-MicroCloud Manager | Ai-WorkStation | Ai-HPC-WorkStation | Ai-Data Lake | Ai-AppStore | Ai-API™ Engine |

| IaaS Composable FC management | PaaS Composable FC management | SaaS Composable FC management | Enterprise management: users, billing, usage, payments |

Data Center 1 — e.g., private /public cloud

ZBL SaaS Composable Foundational Components
- AI application
- Runtime Ai-API™

ZBL PaaS Composable Foundational Components
- Notebook Virtualization
- Workstation master | HPC master | Data Lake master

ZBL IaaS Composable Foundational Components
- Cloud native | Virtualization
- Object store | Block store | HPC FS
- Core: security, VPN, metering, unified ACL

Advantech SKY-8201L high-performance server

Data Center 2 — e.g., telecom tower/cell site

ZBL SaaS Composable Foundational Components
- AI application
- Runtime Ai-API™

ZBL PaaS Composable Foundational Components
- Notebook Virtualization
- Workstation master | HPC master | Data Lake master

ZBL IaaS Composable Foundational Components
- Cloud native | Virtualization
- Object store | Block store | HPC FS
- Core: security, VPN, metering, unified ACL

Advantech SKY-8201 Edge AI server

Data Center 3 — e.g., street cabinet/kiosk

ZBL SaaS Composable Foundational Components
- AI application
- Runtime Ai-API™

ZBL PaaS Composable Foundational Components
- Notebook Virtualization
- Workstation master | Data Lake master

ZBL IaaS Composable Foundational Components
- Cloud native | Virtualization
- Object store | Block store
- Core: security, VPN, metering, unified ACL

Advantech SKY-8101 Edge AI server

Zeblok Ai-MicroCloud offers a composable architecture that makes it simpler to rapidly distribute AI applications from edge to cloud.

Plus, Zeblok Ai-MicroCloud is an easy-to-use solution that can be accessed via a web browser and requires no other specific software or tools. Users can access all familiar tools from

within their Ai-MicroCloud. Ai-MicroCloud has a significant impact on end-to-end pipelines and associated time to market, with engineering time savings of several months.

A streamlined experience for AI delivery

Zeblok Ai-MicroCloud makes it easy to deploy AI environments and workloads. Through a simple web interface, users can:

| Select and configure their AI workstation | Deploy any required microservices | Access the Zeblok Ai-API™ engine |

The power of an open ecosystem

Easy access to reusable AI assets from core to edge is a key accelerator for companies that want to drive Digital Transformation 3.0. The deployment of AI inferencing via an open architecture using Zeblok's Ai-API Engine is a differentiator for network operators and service providers, helping them to include network-wide AI services as part of their portfolio. Currently, public CSPs help deliver AI-APIs through a host of services in their cloud environments, but this also makes the enterprises cloud dependent, resulting in cloud vendor lock-in. CSPs are also exploring edge AI-APIs delivery, further driving cloud vendor lock-in and limiting choices for enterprises that wish to establish new relationships.

Zeblok's Ai-MicroCloud is hardware and cloud vendor agnostic, enabling the creation of an Ai-AppStore that provides enterprises with access to different AI ISVs. This is critical because service providers face the challenge of onboarding many different AI ISVs to drive innovation into their products and services. The challenge is made more difficult by the plethora of ISVs with increasingly specialized AI algorithms. By enabling service providers and enterprises to create their own Ai-AppStore, Zeblok helps them rapidly develop an AI ecosystem that simplifies and accelerates deployment.

Through its ISV network, Zeblok provides proprietary tools for data-driven insights in any industry. Ai-Rover™ for Multi-Dimensional Data is a low-code tool that delivers explainable AI on complex data sets by automatically discovering important correlations and causal relationships through a visual interface, supporting data analysts in the construction of trustable decision-making AI/ML models. Ai-Rover for Time Series Data is a no-code tool that delivers predictive analytics for time series data, enabling anyone—not just data scientists—to create value from business data.

Zeblok's Ai-MicroCloud also includes an automated Ai-API engine that lowers the cost of insight by delivering Ai-APIs to the edge at scale and to on-premises data centers. The solution can be used to automate the deployment of edge AI applications without CSP or hardware lock-in, supporting hybrid cloud strategies.

Only Zeblok Ai-MicroCloud provides Ai-Optimization-as-a-Service™, enabling enterprises to leverage Intel® toolkits such as the Intel Distribution of OpenVINO toolkit and Intel oneAPI toolkits to provide socket-specific optimization of code running at the edge as an Ai-API. Ai-MicroCloud is uniquely comprehensive because it allows enterprise customers to create and manage their open ecosystem of CSPs, OEMs, and ISVs to create their own Ai-AppStores—as well as to develop an AI/ML model, train it, test it, optimize it, and deploy it as an Ai-API—all on a single platform, via a simple user interface.

Deploy a flexible edge cloud platform

While Ai-MicroCloud has been designed from the ground up for AI workloads, the compute, network, storage virtualization, and containerized microservices architectures enable enterprise customers, network operators, and service providers to rapidly deploy any application to edge data centers. The combination of Zeblok software and Advantech hardware provides a full-stack edge cloud platform that's based on secure and open architecture to deliver solutions at the best price, with support for hybrid cloud and cloud vendor–agnostic deployment models.

Advantech SKY-8000 Edge Servers

Reliability
- Redundant BIOS and FW
- Redundant PSU and cooling
- High ESD immunity
- MC class-B bare-bones design and dust filter support (selected models)

Serviceability
- Fail-safe BIOS and FW updates via Advantech IPMI Core
- Platform Health Diagnostics (Advantech Server Manager: running on CPU)
- Advantech BMC WebUI (Node Explorer)
- Redfish-compliant OM options

Embedded background
- Product change notifications on key components for firmware and hardware
- Customization options

Performance
- 2nd and 3rd Gen Intel® Xeon® Scalable processors
- NUMA-balanced design
- High I/O and acceleration density
- Workload-optimized Intel® Select Solutions configurations
- In-house electrical, mechanical, and thermal designs

Longevity
- 7–15 years' product life cycle plus extended product warranty
- Fixed BOM for customized product P/N

Environmental
- Extremely wide operating temperature range
- NEBS–3 compliant designs
- IP65 outdoor cabinet configurations
- Chassis mechanical twist strength
- 11- to 31.5-inch depth chassis

Deployed globally in enterprise, industrial, and telecom business and mission-critical applications

Leverage micro data centers at the edge

The availability of an open, network-wide AI infrastructure that can be leveraged by ISVs to quickly deploy new high-bandwidth, low-latency AI applications is a key network differentiator to attract new partners, subscribers, and customers. However, the implementation of such a distributed intelligent cloud requires the deployment of data center–type capabilities at the edge where the redundancy and controlled environment of a data center cannot be replicated. Advantech SKY-8000 Series 5G Intelligent Edge Servers have been designed to deliver high reliability and computing performance for AI services running in micro MEC data centers at the edge.

Advantech SKY-8000 servers combine the high performance and openness of cloud computing with the ruggedness, efficiency, and high reliability required by edge deployments. Their advanced platform management and edge hardware design help service providers and operators minimize costly system downtime, service interruptions, and support interventions.

Advantech SKY-8000 Edge Servers have been certified by Zeblok Computational for hub and edge deployment. They support the latest Intel Xeon Scalable processors and high-density PCI Express payload to integrate powerful AI acceleration in optimized and efficient edge infrastructure solutions. Advantech's carrier-grade platforms have been optimized for the intelligent edge to support a wide range of operating temperature, environmental shock, vibration, and dust conditions. Redundant power supplies, the ability to withstand single fan failures, redundant BIOS and firmware images with fail-safe remote updates, and hot-swappable FRUs improve availability of mission-critical applications that demand absolutely minimal downtime. In addition, their compact design perfectly fits edge environments where limited space is available, including kiosks or IP65 pole-mounted street-side cabinets.

Innovate with a first-of-its-kind edge infrastructure solution

Beyond corporations, edge AI is providing cities with the opportunity to build more-efficient and more-attractive public spaces that are both safer and more environmentally friendly. To thrive in the next digital transformation, cities need to integrate edge computing capabilities into their landscape without disturbing their citizens' daily lives. SmartPoint.io offers a first-of-its-kind hyperconverged edge infrastructure that transforms street furniture into smart street furniture, unlocking new applications that leverage high-speed connectivity and edge compute resources.
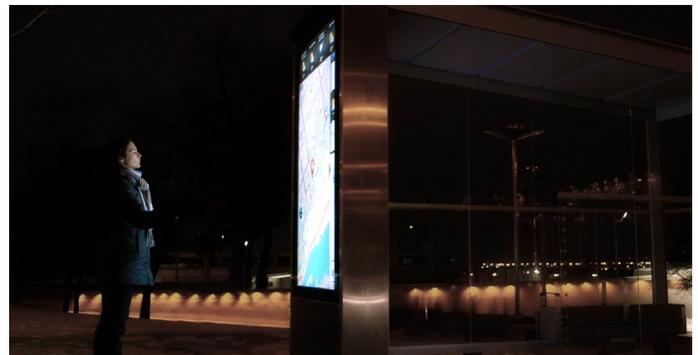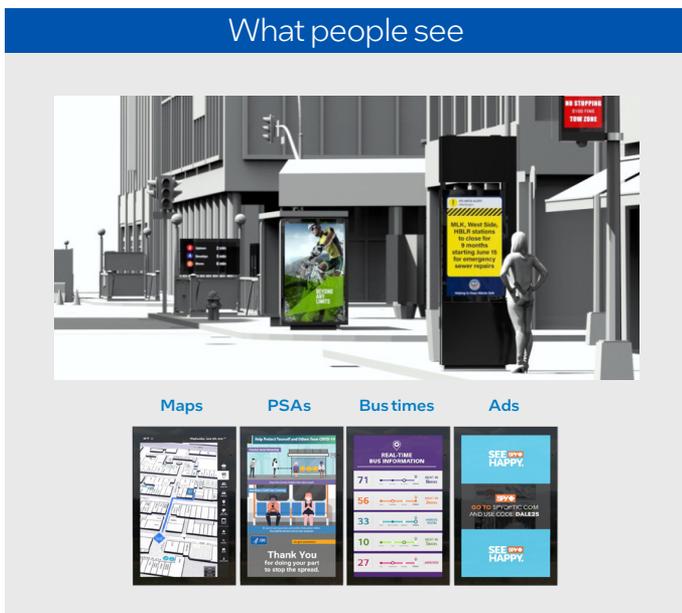
SmartPoint.io deploys what look like giant smartphones on city sidewalks. With their bright display, touchscreen interface, and location-based apps, they are intuitively welcoming for the public to use. The many functions of the SmartPoint.io kiosk hubs include wayfinding, emergency alerts, community news, and transit information. These services can be subsidized by brands that leverage a portion of the screen time for advertising purposes, creating a commercial bridge to a new revenue source that directly contributes to operational value within the installed environment.

Each SmartPoint.io kiosk hub contains a carrier-grade Advantech SKY-8000 Edge Server and high-speed connectivity to enable software applications to run at the physical edge of the network. A multitude of patents make converging all this capability into a single point of presence possible, even as the kiosks are exposed to the harshness of the city streets. Nearly limitless value is created when this edge infrastructure is easily accessible by the services and applications that require proximity. Zeblok Ai-MicroCloud delivers a cloud-native, full-stack edge platform that can be used for both AI and non-AI workloads through integration, centralization, orchestration, and automation.

Powered by Intel Xeon Scalable processors

Intel Xeon Scalable processors are the foundation for the powerful edge platforms that deliver compute agility and scalability. Disruptive by design, they benefit from decades of innovation for the most-in-demand workload requirements and are part of a complete set of network technology from Intel.

Intel Xeon Scalable processors enhance edge server solutions with a balanced architecture that supports AI with built-in acceleration and hardware-based security features. These CPUs are also engineered for modern 5G network workloads, targeting low-latency, high-throughput deterministic performance and high performance per watt.

What people see

Maps   PSAs   Bus times   Ads

What people don't see

Sensors   Wireless communications
Cameras
Network
Compute   Storage

SmartPoint.io consolidates a full stack of edge technologies into welcoming and intuitive kiosk hubs that can operate unobtrusively in busy urban environments.

Ai-MicroCloud™ and SKY-8000 Edge AI use cases

- **Ai-MicroCloud certification for hub:** Ai-MicroCloud Manager certified on Advantech SKY-8000 Edge Servers can be deployed across the network in hubs such as regional data centers or central offices or in micro data centers at the edge in communications towers, cell sites, or IP65 pole-mounted or street-side cabinets.

- **Ai-MicroCloud certification for kiosk:** Ai-MicroCloud Manager certified on Advantech SKY-8000 Edge Servers can run in kiosks to offer MEC capabilities for AI inferences at the very edge.

- **AI ISV certification:** ISVs can deploy their certified solution as an Ai-API to each kiosk or edge endpoint, leveraging MEC by the curbside. They can then retrain the AI/ML model underlying their product and optimize the revised model using the Intel® Distribution of OpenVINO™ toolkit and Intel® oneAPI toolkits to optimize their code and quickly deploy the revised Ai-API to each kiosk from their Ai-MicroCloud Manager.

- Enterprise customers delivering smart solutions will oversee all aspects of the solution from their Ai-MicroCloud Manager.

Conclusion

The combination of AI, 5G, and MEC is driving the next digital transformation. Zeblok Computational and Advantech have joined forces to help infrastructure and service providers build a more attractive and innovative network and portfolio for Industry 4.0, smart city, telecom, fintech, healthcare, manufacturing, autonomous vehicle operators, and retail customers by creating an Ai-API economy where AI solutions can be distributed from cloud to edge rapidly and cost-effectively.

By leveraging the Zeblok Ai-MicroCloud AI Platform as a Service and Advantech SKY-8000 Series of Edge Servers, operators and service providers can create a vibrant AI ecosystem and offer enterprise customers an Ai-AppStore where they can quickly onboard new ISVs to deliver edge AI and integrate AI solutions into mission-critical processes, faster and with lower costs. The integration of such an open and scalable platform, leveraging the SmartPoint.io kiosk solution, unlocks limitless value by providing an AI-enabled data center on the sidewalk that can deliver value to citizens, government agencies, and businesses.

Learn more

Intel Distribution of OpenVINO toolkit
The Intel Distribution of OpenVINO toolkit helps accelerate the development and deployment of machine learning solutions.

Learn more›

Intel oneAPI AI Analytics toolkit (AI Kit)
The Intel oneAPI AI Analytics toolkit (AI Kit) gives data scientists, AI developers, and researchers familiar Python tools and frameworks to accelerate end-to-end data science and analytics pipelines on Intel architectures.

Learn more›

Zeblok Ai-MicroCloud
Zeblok Ai-MicroCloud solves the problem of scaling at the edge, making it easy to deploy AI inference to thousands of edge locations with end-to-end life-cycle management.

Learn more›

Advantech SKY 8000 Edge Servers
Advantech SKY 8000 Series Servers optimize throughput and offload required by 5G and Edge AI workloads, to meet high availability and carrier-grade network needs.

Learn more›

About Zeblok Computational

Zeblok Computational's mission is to lower the cost per AI inference insight, promote the Ai-API economy, and enable cost-effective scaling at the edge. Zeblok's Ai-MicroCloud is the most straightforward way to efficiently pipeline data—including the critical data comprehension step—and then quickly and affordably develop, train, and deploy pragmatic AI into mission-critical enterprise business processes from cloud to edge.

computational.zeblok.com

About Advantech

Advantech's corporate vision is to enable an intelligent planet. The company is a global leader in the fields of IoT intelligent systems and embedded platforms. To embrace the trends of IoT, big data, and artificial intelligence, Advantech promotes IoT hardware and software solutions with the Edge Intelligence WISE-PaaS core to assist business partners and clients in connecting their industrial chains. Advantech is also working with business partners to cocreate business ecosystems that accelerate the goal of industrial intelligence.

advantech.com

About SmartPoint.io

SmartPoint.io offers a scalable platform designed to meet the growing demand on digital infrastructure by consolidating resource-intensive workloads at the edge, optimizing network traffic, and reducing latency for real-time decision-making and execution. The SmartPoint.io offering is a data center on the sidewalk.

smartpoint.io

intel + **ZEBLOK COMPUTATIONAL** + **ADVANTECH** + **SMARTPOINT.IO**