

Zero Gap AI Total Fabric

Powering the future of AI at the Edge with the Ultimate AI Deployment Model.

AI deployment shouldn't be limited by cloud latency or rigid infrastructure.

Zero Gap AI Total Fabric seamlessly moves AI workloads across on-prem, near-prem, and private 5G, ensuring real-time inferencing and automated workload orchestration for faster, smarter AI deployments—right where they're needed most.



AI elasticity, allowing workloads to scale without disruption



Automated workload balancing across sites, reducing latency and congestion



Private AI networking that eliminates reliance on public cloud networks



Avoid public cloud dependencies, ensuring data privacy and compliance



Use Layer 2 interconnects for direct AI workload communication across multiple locations



Deploy private 5G for ultra-low latency AI applications

Zero Gap AI Total Fabric enables enterprises to deploy AI anywhere with ultra-low latency, private 5G, and scalable edge computing. Whether optimizing smart cities, automation, or retail, this cost-effective, secure AI solution eliminates cloud dependencies and powers real-time decision-making at scale.

Unifying AI, Compute and Networking for Maximum Performance

AI is no longer just about raw processing power. Enterprise AI demands instant intelligence, seamless connectivity, and dynamic workload mobility across infrastructure. Traditional AI deployments—whether in the cloud or on-premises—fail to deliver the speed, flexibility, and scale required for mission-critical AI workloads..



Zero Gap AI Total Fabric is a complete AI deployment solution that integrates on-prem and near-prem edge compute with private AI networking, orchestration, and workload mobility.

Built for organizations that require real-time AI execution without relying on distant cloud data centers, Total Fabric allows enterprises to deploy AI exactly where it is needed—close to operations, on dedicated infrastructure, with full control over performance, security, and cost.



Highlights

- Deploy AI workloads effortlessly across on-prem, near-prem, and private 5G environments for maximum efficiency and flexibility.
- Reduce latency with AI inferencing and decision-making at the edge, enabling ultra-fast responses for critical applications.
- Dynamically move AI workloads to the most optimal location based on compute availability, network conditions, and performance needs.
- Leverage AI-enhanced private 5G to ensure secure, high-bandwidth connectivity for real-time AI applications and data-intensive workloads.
- Keep AI models and data secure with dedicated private AI networking, ensuring full compliance with regulatory and enterprise security standards.

Why Zero Gap AI Total Fabric?

AI infrastructure has traditionally forced a choice between public cloud AI and on-prem AI, each with significant limitations:

The Problem with Public Cloud AI

- AI workloads must be sent to distant cloud regions, increasing network latency and slowing real-time AI processing.
- Data privacy and compliance risks arise when sensitive AI models and training datasets leave private infrastructure.
- Cloud egress fees and unpredictable costs make scaling AI expensive and difficult to control.

The Problem with On-Prem AI

- On-prem AI provides control, but scaling requires massive upfront investments in GPU hardware and networking.
- AI workloads are often trapped in isolated environments, limiting collaboration between sites.
- Expanding on-prem AI across multiple locations is complex and requires dedicated IT resources.

Zero Gap AI Total Fabric combines the best of both worlds by integrating on-prem AI processing, near-prem AI compute, and private networking into a single, orchestrated AI infrastructure.

Zero Gap AI Total Fabric

- AI runs where it is needed most—on-prem for low-latency inference and near-prem for scalable training and high-performance compute.
- AI workloads move seamlessly between locations with private fiber and 5G, ensuring low-latency AI execution without cloud bottlenecks.
- No need to build massive on-prem GPU clusters—Total Fabric provides scalable near-prem compute in 36 US metro areas, reducing infrastructure costs.
- Security and data sovereignty are maintained with dedicated private AI networking that eliminates public cloud dependencies

Zero Gap AI Total Fabric is designed for real-world, enterprise-scale AI deployments.

Smart Cities & Public Safety

- Real-time traffic and infrastructure monitoring with edge AI cameras
- AI-powered public safety analytics, reducing emergency response times
- Secure private fiber and 5G AI networking for seamless data processing

Industrial Automation & Manufacturing

- AI-driven predictive maintenance, reducing downtime and equipment failures

- Machine vision and robotics for automated quality control
- On-prem AI for real-time decision-making, combined with near-prem AI for scalability

Digital Twin & AI-Driven Simulation

- Create real-time digital replicas of physical assets, systems, and operations for AI-powered decision-making
- Simulate, monitor, and predict outcomes using continuously updated real-world data
- Enhance operational efficiency across industries by integrating Digital Twins with AI models at the edge



Learn More About the Zero Gap AI

Find out how your organization can benefit by contacting us at sales@veeva.com or by visiting veeva.com/resources