



SIEMON AI-READY SOLUTIONS

Emerging AI Data Center Network Architectures and Applications

From the Traditional Front-end to the AI-Powered Back-end

www.siemon.com

Connecting the World to a Higher Standard

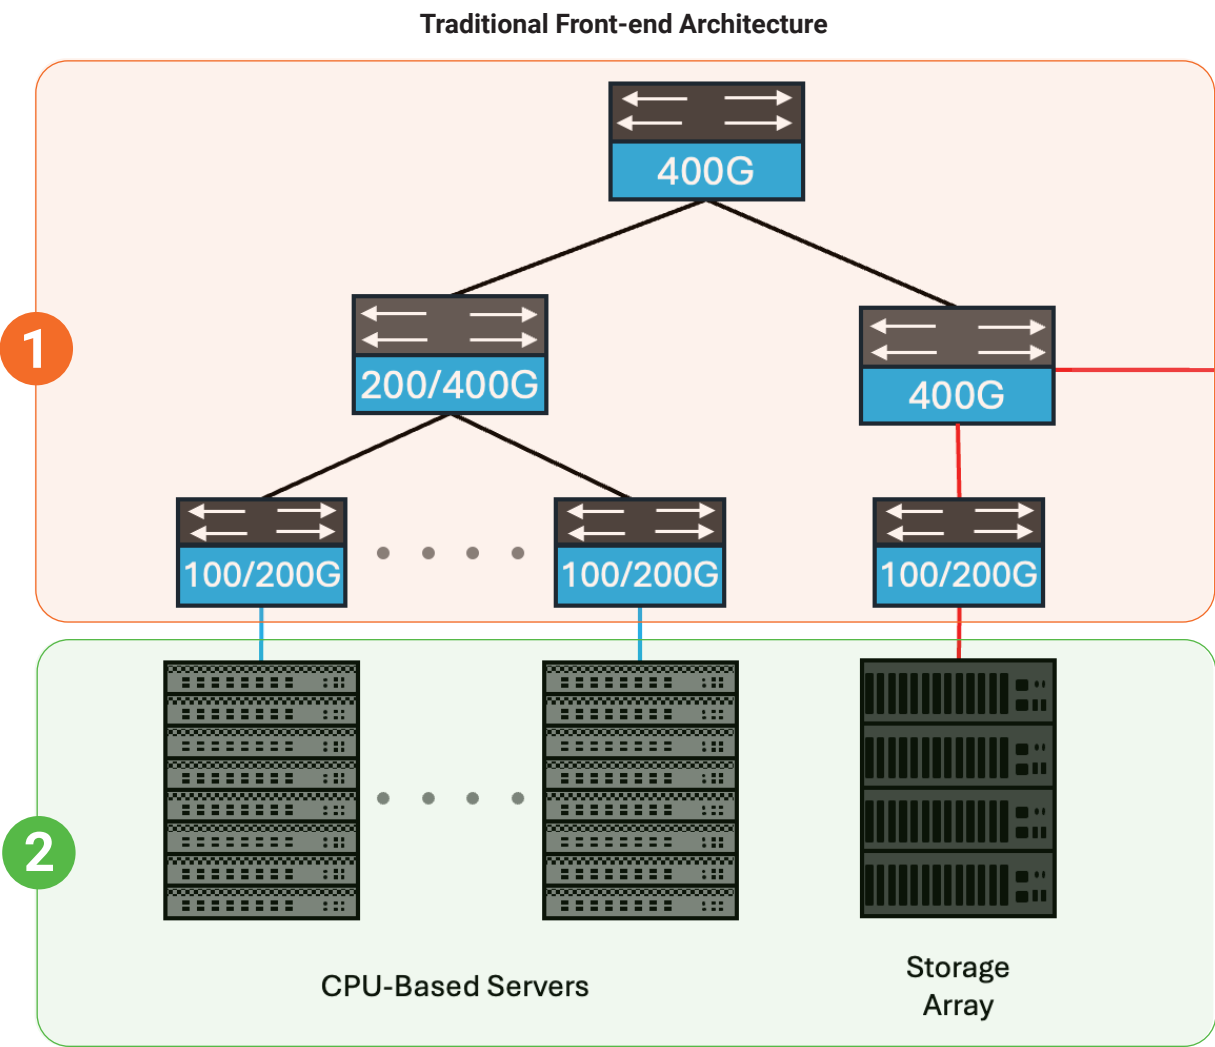
Meeting the Demands of AI & Modern Workloads

To support HPC workloads like AI/ML training, back-end networks deploy spine-leaf architecture where leaf switches connect to every spine switch. Within AI pods (clusters) that are purpose-built to perform specific tasks, leaf switches provide high-bandwidth, low-latency interconnections between GPUs. Spine and leaf switches typically connect at 800 Gigabit (G) Ethernet, with 1.6 Terabit (T) on the horizon, while switch-to-server speeds are primarily 400G via Ethernet or InfiniBand™ protocols. To ensure seamless connectivity for general-purpose workloads and the transfer of data to and from the back-end network, the front-end is rapidly migrating to 200/400G Ethernet switch-to-switch links and 25/50/100G switch-to-server links. A high-performance front-end network is essential to accelerate AI inference requests from users and devices for an overall better experience.

There are 4 different types of network interconnection areas:

- 1. Front-end, Switch-to-Switch
- 2. Front-end, Switch-to-Server
- 3. Back-end, Switch-to-Switch
- 4. Back-end, Switch-to-Node

Let's look at each interconnection area in more detail:

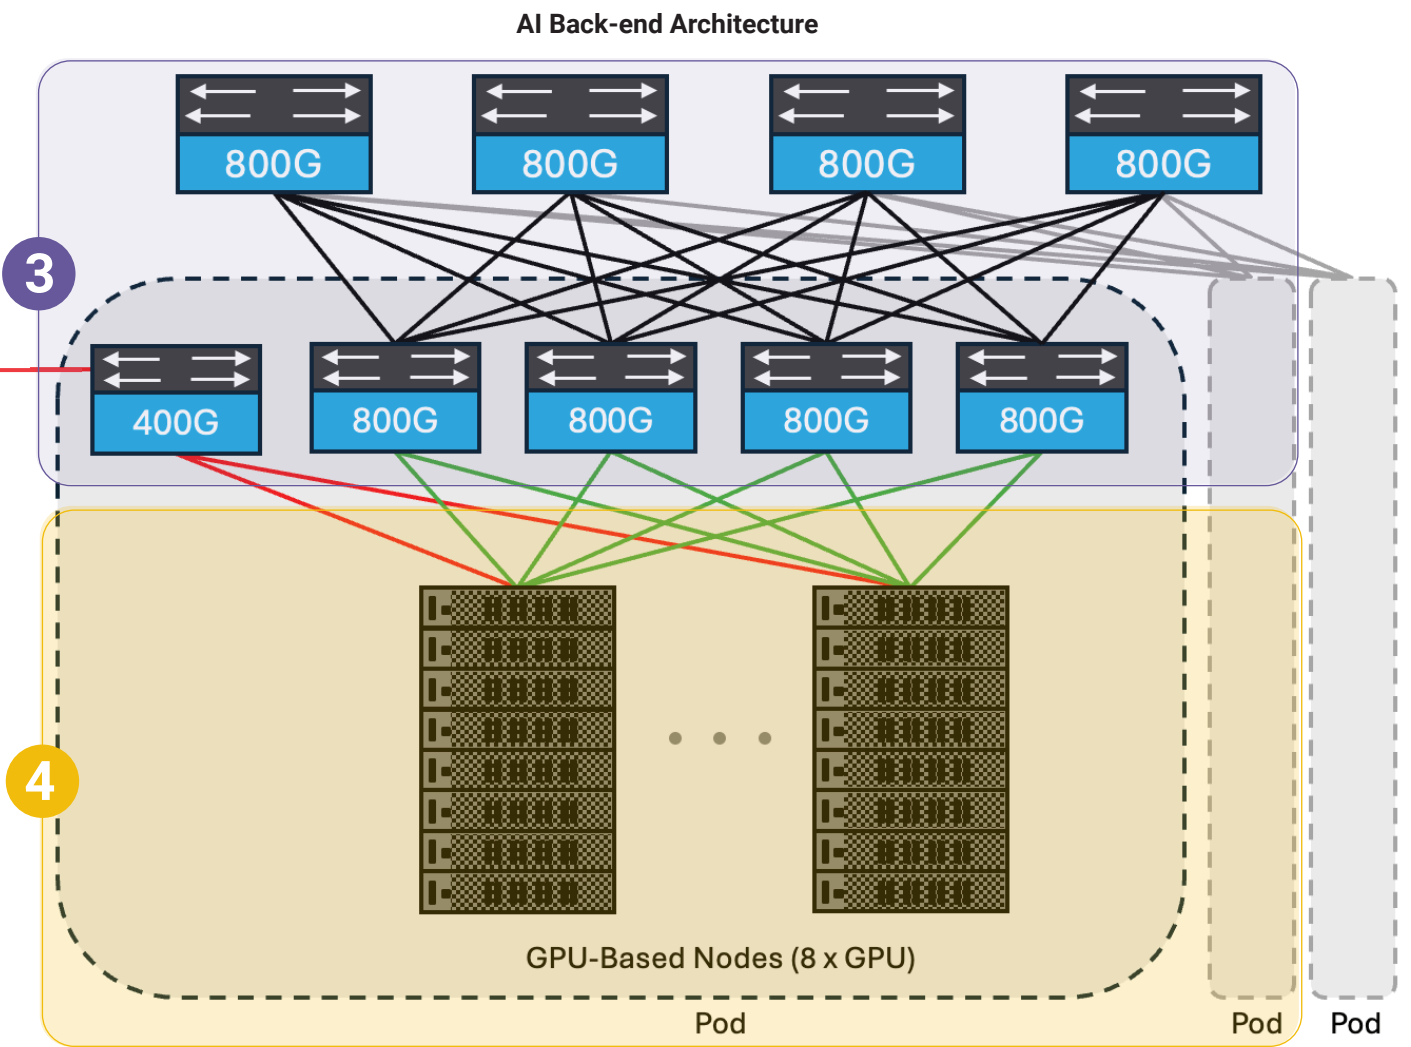


Bridging the AI Divide

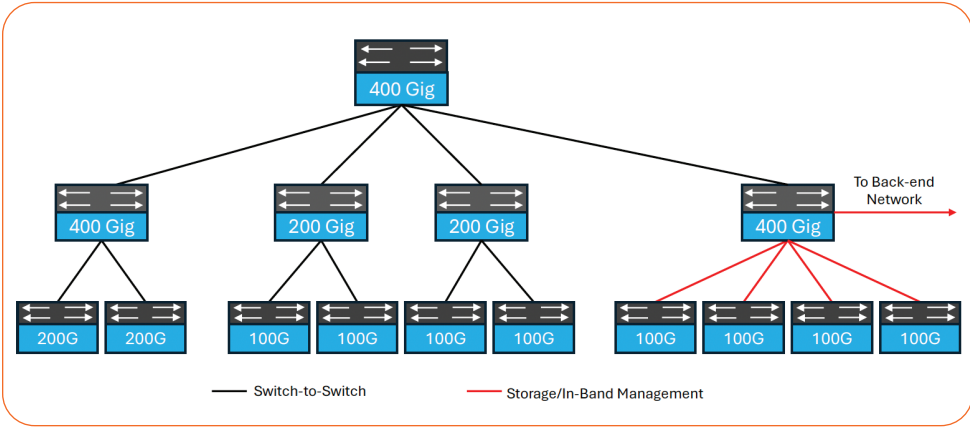
The data center landscape is transforming due to AI and ML. This necessitates a new back-end network in cloud and large enterprise data centers, designed for HPC workloads like AI training. Traditional front-end networks are evolving for general-purpose workloads and external connections.

- **Back-end Network** – A closed **spine-leaf** network for ultra-high-speed, low-latency east-west data transmission across interconnected GPUs for AI training.
- **Front-end Network** – A traditional client-based **spine-leaf** network with Ethernet switching, CPU servers, and storage arrays, leveraging north-south data transmission for internet connection, data storage and movement, and AI inference requests.

As today's data centers embrace the AI revolution, Siemon understands the criticality of infrastructure that seamlessly connects high-performance back-end networks with the agile front-end, while supporting the unique connectivity needs of each environment. Whether you're housing back-end and front-end networks in a single rack for edge AI, building an entry-level AI experimental playground with just a few racks, or expanding AI training at scale with large complex rail-optimized pods, Siemon has everything you need to ensure superior front-end and back-end network performance – no matter what your topology or where you are in your AI journey.



Front-end Network | Switch-to-Switch Interconnections




- Traditional front-end networks use 200G and 400G Ethernet for switch-to-switch connections.
- These connections often run row-to-row, but can use EoR or MoR topology in smaller data centers.
- Breakout configurations (2 x 200G, 2 x 100G, 4 x 100G) can maximize port utilization and space efficiency for switch connections.

Siemon Solutions for Front-end Switch-to-Switch:

[DACs](#) for within a cabinet (3m typical) | [AOCs](#) for within a row (up to 30m typical)
[BASE-8](#) MTP cabling, enclosures, panels and modules (+10m typical)

Front-end Network | Switch-to-Server Interconnections

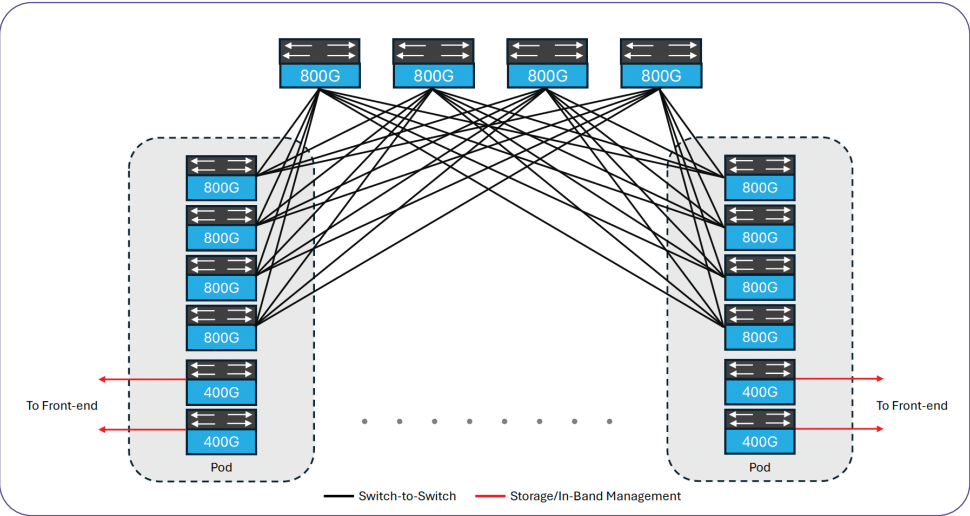


- Front-end networks' access switches are migrating to 100G and 200G Ethernet.
- This supports 25, 50, or 100G server links.
- Breakout configurations (2 x 50G, 4 x 25G, 2 x 100G) maximize port utilization and space efficiency.
- Front-end switch-to-server connections use EoR/MoR or ToR topologies.
- Topology choice depends on data center size and layout.

Siemon Solutions for Front-end Switch-to-Server:

[DACs](#) for within a cabinet (3m typical) | [LC and MTP](#) jumpers within a cabinet (3m typical)
[Cat6/6A](#) Skinny Patch Copper Cords for 10G (3m typical) | [AOCs](#) for within a row (up to 30m typical)
[BASE-8](#) MTP cabling, enclosures, panels and modules (+10m typical)

Back-end Network | Switch-to-Switch Interconnections

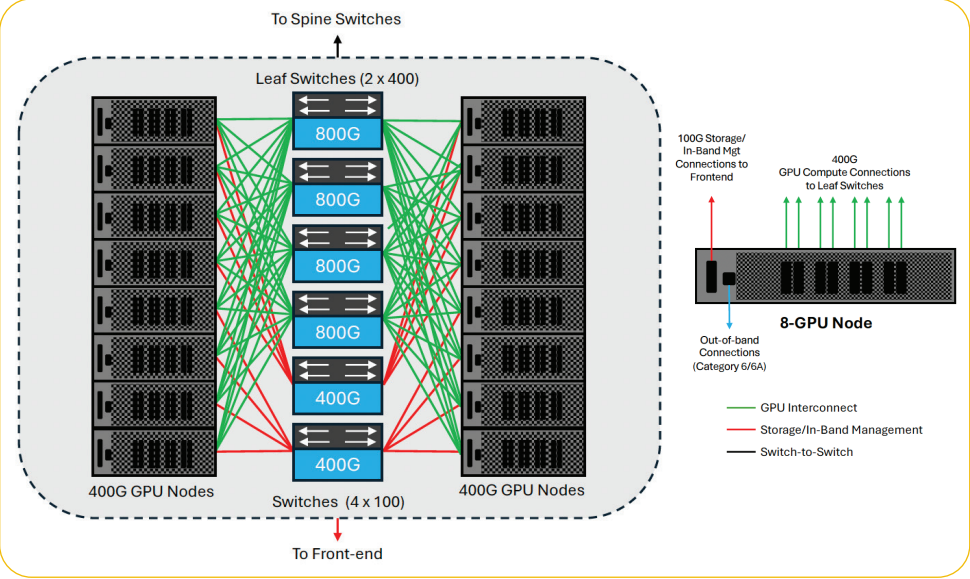


- HPC back-end networks utilize 800G Ethernet spine-leaf connections, with 1.6T Ethernet emerging.
- These connections support row-to-row or within-row deployments via EoR or MoR topologies.
- ToR topology suits small edge/entry-level AI deployments where spine and leaf switches share a cabinet.

Siemon Solutions for Back-end Switch-to-Switch:

[DACs](#) for within a cabinet (3m typical)
[ACCs/AECs](#) for within a row (up to 3-5m typical) | [AOCs](#) for within a row (up to 30m typical)
[BASE-8](#) MTP MM/APC and SM cabling, enclosures, panels and modules (+10m typical)

Back-end Network | Switch-to-Node Interconnections



- AI pods feature 400G GPU connections within server nodes, each linked to a distinct leaf switch for optimized GPU communication.
- This is commonly implemented with high-port-count 800G leaf switches using 2 x 400G breakout configurations.
- Each server node also requires CPU-to-front-end network connections (typically 100G) for storage and management.
- Lower-speed copper connections are used for out-of-band management.


Siemon Solutions for Back-end Switch-to-Node:

[DACs](#) for within a cabinet (3m typical) | [ACCs/AECs](#) for within a cabinet/POD (up to 3-5m typical)
[AOCs](#) for within a row (up to 30m typical) | [Cat6/6A](#) Skinny Patch copper cords for 10G (3m typical)
[BASE-8](#) MTP MM/APC point-to-point cables (3-50m typical)
[BASE-8](#) MTP cabling, enclosures, panels and modules (+10m typical)

Siemon has AI-Ready Solution that Supports Every Topology

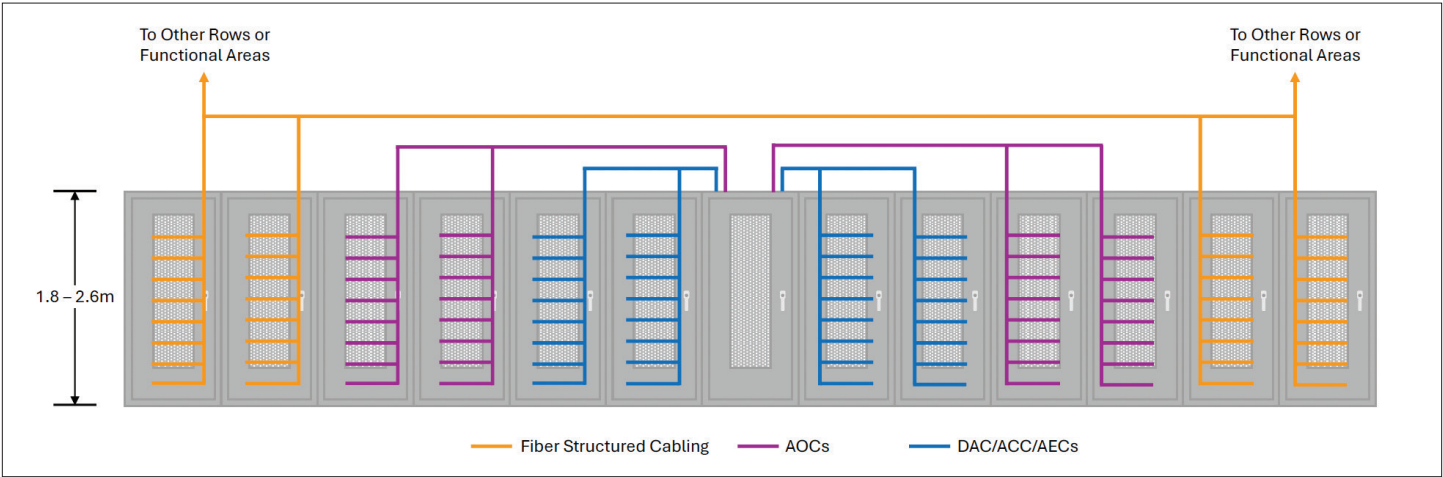
Data center cabling topologies vary. Back-end and front-end networks may use point-to-point or structured cabling over different distances. Links can connect row-to-row, cabinet-to-cabinet (End-of-Row/Middle-of-Row), or within-cabinet (Top-of-Rack). Switch-to-switch and switch-to-server links can use breakout configurations, maximizing port utilization and space

Siemon provides a complete range of high-performance cabling solutions for any distance, topology, and breakout configuration. This includes passive and active direct attach copper (DAC) assemblies, active optical cables (AOCs), and fiber and copper structured cabling systems.



Key topology and cabling considerations include:

- **Space and Layout** – Data center layout, space, and distances between areas influence topology and cabling choices.
- **Flexibility and Scalability** – Structured cabling with patch panels enables flexible cross-connects, simplifying operations and upgrades. This streamlines integrating new pods in expanding AI networks.
- **Power and Cooling** – Co-locating switches and servers (ToR) increases rack power densities, requiring robust infrastructure.
- **Latency** – Point-to-point DACs offer lower latency and power consumption in ToR and MoR topologies.
- **Loss Performance** – For extended structured cabling runs (row-to-row and EoR/MoR), ultra-low-loss BASE-8 multimode and singlemode cabling maximize performance.



Direct Attach Copper Cables

Siemon’s passive and active DACs, AECs, and ACCs have flexible cable construction and support up to 800G for short reach switch-to-server or switch-to-switch connections in both Ethernet and InfiniBand™ versions. They are available in various form factors, half-meter increments from 0.5m to 5m, and offered in straight-thru or multiple breakout and color options.


go.siemon.com/DACs



Active Optical Cables

Siemon’s multimode fiber AOC cable assemblies support point-to-point connections up to typically 30 meters, ideal for EoR/MoR topologies. Multiple breakout options are available. Their energy-efficient design consumes less power than transceiver assemblies with a small cable diameter for improved airflow and reduced cooling costs.

go.siemon.com/AOCs



BASE-8 MTP Cabling

Siemon’s wide range of singlemode and multimode plug and play trunks and assemblies include BASE-8 MTP-to-MTP trunks in a variety of fiber counts up to 144F, MTP conversion cords, and MTP Pro jumpers.

go.siemon.com/BASE-8



Ultra-High-Density LightStack® Fiber Optic Cabling System

Siemon’s LightStack and LightStack 8 ultra-high-density fiber plug and play system delivers superior singlemode and multimode fiber density, port access, and cable management in a modern platform for advanced data centers. The system offers 1U and 4U enclosures with integrated cable management, supporting up to 144 LC and 864 MTP fibers per rack unit. A full range of fiber trunks, jumpers, and assemblies are available.

go.siemon.com/LightStackEnclosures



High-Density LightVerse® Fiber Optic Cabling System

Siemon’s LightVerse Fiber Optic Cabling System offers high-density enclosures, ultra-low-loss MTP-to-LC modules, MTP and LC adapter plates, and MTP and LC jumpers and assemblies. These support singlemode and multimode fiber links in front-end and back-end networks, designed for 800G and beyond.

go.siemon.com/LightVerseEnclosures



SIEMON. YOUR PARTNER FOR EVERY STEP OF YOUR AI JOURNEY.



For more information visit:
www.siemon.com



Find your local Siemon distributor:
go.siemon.com/distributor



24/7 Customer Support:
customer_service@siemon.com