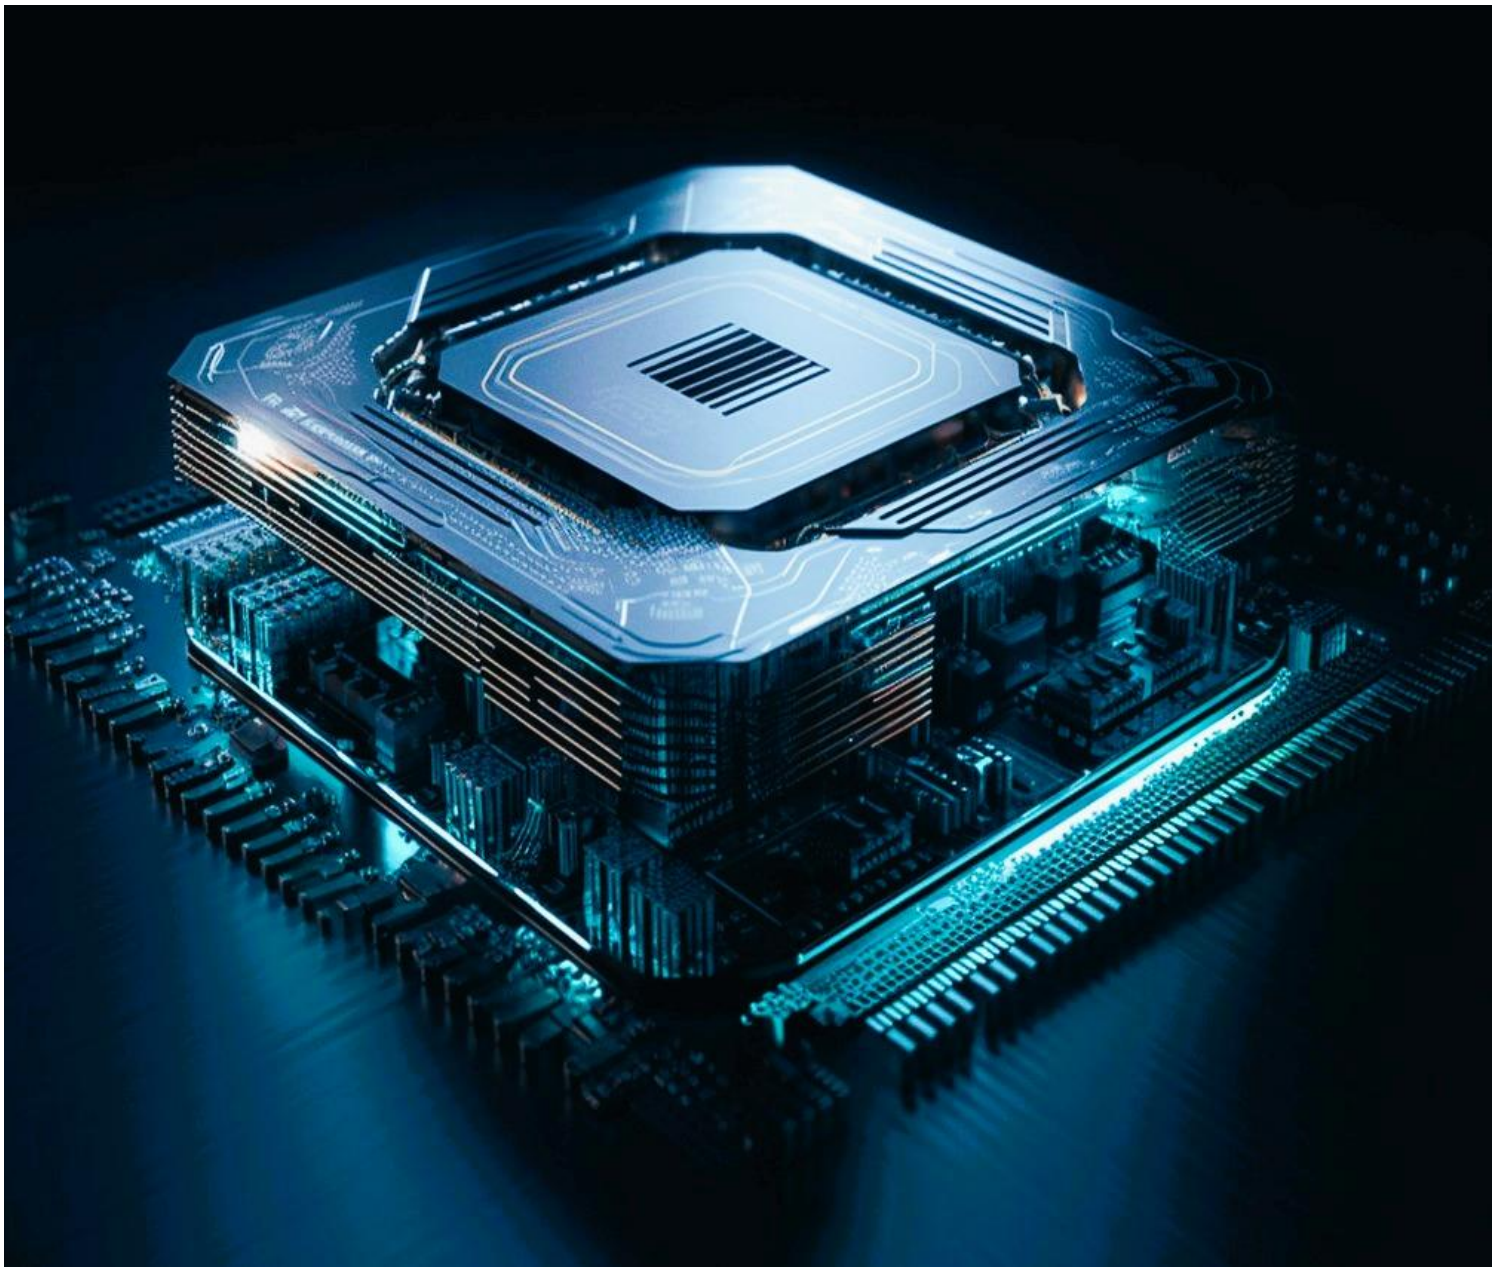

SiSys AI

Silicon and Systems for AI Applications

October 2025

george.apostol@sisysai.com

SiSys AI

Overview of SwitchNoC™

SwitchNoC™ enables next-gen **AI** networking—non-blocking, layout-friendly topology with unified control for predictable QoS.

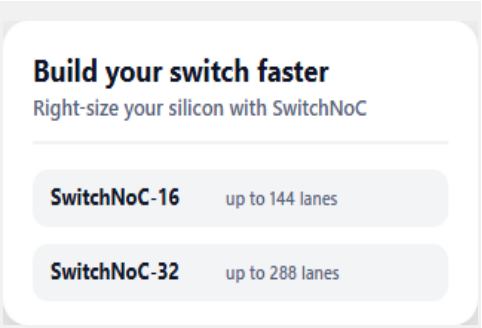

Benefits:

- High bandwidth + low latency: Up to 14.4 TB/s port switching capacity with ~10–30 ns fabric latency—built for chiplet and switch-class throughput.
- Lane-dense scalability: Configurations such as SwitchNoC-16 (144 lanes) and SwitchNoC-32 (288 lanes) are available to match a range of die and system targets, with flexibility to support other lane counts.
- Complete non-blocking design: Port-level non-blocking, virtual channel non-blocking and transaction type non-blocking.
- Faster time-to-market: Designed with a compact, configurable switch fabric and unified control management fabric.
- Better PPA (performance/power/area): Significantly reduced wires/crossings and lower energy per bit
- Protocol flexibility: UALink / PCIe / CXL / Ethernet for AI, HPC, and data center workloads.
- Heterogeneous station: Support bridging between protocols.apt

Physical Cognizance:

- Floorplan Provision: Offers a physically cognizant floorplan reflecting the initial design intent.
- Area-efficient by design: Architecture and implementation target performance, area, and power—the compact fabric helps shrink die footprint.
- Strong Starting Point: Provides an excellent foundation for physical layout teams.

Development History:

- Result of decades of expertise in high-performance switch design @ PLX Technology and Elastics.cloud (Now Broadcom) and HiFn

Key Competitive Advantages:

- Modified 2D-torus topology — Innovative non-blocking topology with shorter wires achieving order-of-magnitude layout efficiency vs. crossbars.
- High bandwidth + low latency: 14.4 TBps switching capacity and 10-30 ns latency.
- Embedded management fabric — mailbox + unified config for centralized or distributed control.
- Power-efficient algorithms — minimized energy per transferred bit.
- Design flexibility — broad configurability for AI/HPC/networking/storage.

Advanced Features:

- **Per-port VOQs and VIQs** for throughput and fairness.
- **Virtual Channels** to segment traffic and enforce priorities across congestion domains.
- **Multicast/Broadcast** with efficient design to support for sync-heavy AI flows.
- **Protocol-agnostic datapath** design to support UALink/PCIe/CXL/Ethernet.
- **Port-Level Bifurcation** dynamically configurable in SwitchNoC-16.
- **Link-level Bifurcation** naturally supported.
- **Deterministic QoS** with end-to-end enforcement.
- **A Lightweight Control Network** to enable intelligent system management, configuration, and diagnostics.

Available Now

Custom Turn-key ASIC with SwitchNoC™

SiSys AI

SwitchNoC™ Common Features

Physical Optimization:
- Modified 2D-torus switch topology optimized for on-die placement and routing.
- Eliminates cross-chip long wires to reduce latency and routing congestion.
- Substantial wire-count reduction versus a traditional crossbar.

Datapath Interfaces:
- AXI-Stream–like interfaces to connect cleanly with high-speed link controllers.
- Additional metadata connects to the multicast/broadcast data and credit system of the high-speed link controllers.

Integrated VOQs and VIQs:
- Virtual Output Queues (VOQs) and Virtual Input Queues (VIQs) are embedded into the NoC structure. These queues help prevent head-of-line blocking and facilitate clean separation of traffic classes or destinations.
- Support intelligent packet cut-through to improve latency.

Multicast and Broadcast Support:
- SwitchNoC™ natively supports efficient multicast and broadcast mechanisms. This capability enables scalable collective communication across endpoints, particularly valuable for AI, cache coherency, and control message propagation.

End-to-end Quality of Service (QoS):
- SwitchNoC™ features deterministic QoS enforcement and end-to-end bandwidth guarantees across the fabric.
- Each end-port can independently configure unique bandwidth weights to every other destination port, enabling fine-grained, per-destination traffic prioritization and QoS enforcement.
- Eventually connects to the credit system of the end-port controller.
- Prevents oversubscription and traffic congestion.

Embedded Management Interfaces:
- A lightweight management interface is embedded into the switch to allow an on-chip CPU or firmware agent to configure, monitor, and manage the NoC fabric in real time. This enables fine-grained traffic and performance tuning.
- Mailbox mechanism to facilitate communication between an on-chip CPU or firmware agent to communicate with each of the port hosts.
- A unified switch configuration interface provides streamlined access to control registers and runtime parameters, enabling centralized or distributed intelligent switch management.

Data Protection and Integrity:
- Built-in ECC and SECDED mechanisms are available for packets/flits, metadata and routing tags.

Configurability and Customization:
- SwitchNoC™ dynamically adapts to diverse chip design intents—enabling customization of port count, transaction types, packet/flit structure, performance characteristics, etc. to meet the needs of specific applications.
- Optimized for physical design, SwitchNoC™ supports fine-tuned adjustments based on layout, routing, and floor planning needs.

Verification Model:
- A plug-in SystemVerilog-based testbenches, protocol checkers, and packet/flit checkers can be directly used in a SoC environment.
- Stand-along SystemVerilog-based testbenches with generators, drivers, monitors and scoreboards, directed/random test suites can be provided as well.

SiSys AI

SwitchNoC™ Product Versions and Features

SwitchNoC-16:

SwitchNoC™ Common Features with Additions:

1. Per XP Port-Level Bifurcation Support:
   - 8 Port x 2 Lane
   - 4 Port x 4 Lane,
   - 2 Port x 8 Lane,
   - 1 Port x 16 Lane
2. Supported Protocol: UALink, PCIe, CXL, Ethernet and other high-speed port
3. Heterogeneous Station Support: Yes
4. Total Switching Capacity at Port Level: Up to 7.2 TBps (configurable based on port protocol)
5. Total NoC Switching Capacity: Up to 9 TBps (configurable based on port protocol)

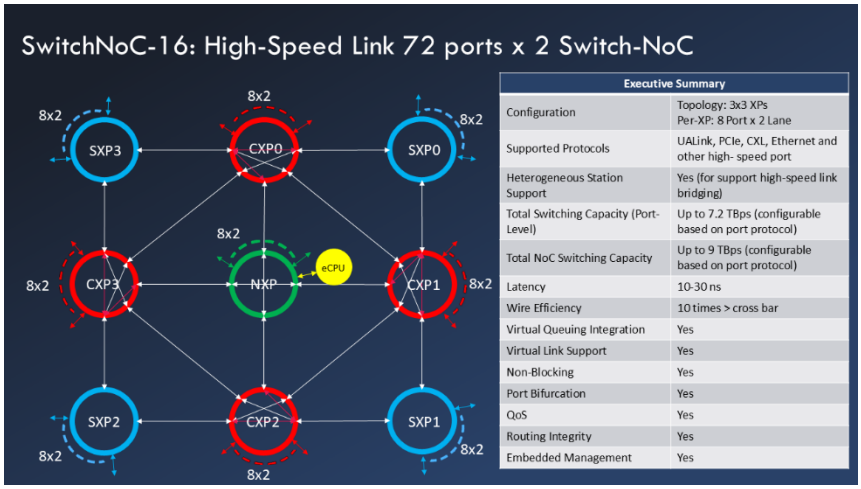

Executive Summary	
Configuration	Topology: 3x3 XPs Per-XP: 8 Port x 2 Lane
Supported Protocols	UALink, PCIe, CXL, Ethernet and other high- speed port
Heterogeneous Station Support	Yes (for support high-speed link bridging)
Total Switching Capacity (Port-Level)	Up to 7.2 TBps (configurable based on port protocol)
Total NoC Switching Capacity	Up to 9 TBps (configurable based on port protocol)
Latency	10-30 ns
Wire Efficiency	10 times > cross bar
Virtual Queuing Integration	Yes
Virtual Link Support	Yes
Non-Blocking	Yes
Port Bifurcation	Yes
QoS	Yes
Routing Integrity	Yes
Embedded Management	Yes

SwitchNoC-32:

SwitchNoC™ Common Features with Additions:

1. Per XP Support:
   - 4 UALink Stations, with each station supports link-level bifurcation
     - 8 port x 1 lane
     - 4 port x 2 lane
     - 2 Port x 4 lane
2. Supported Protocol: UALink
3. Total Switching Capacity at Port Level: Up to 14.4 TBps
4. Total NoC Switching Capacity: Up to 18 TBps

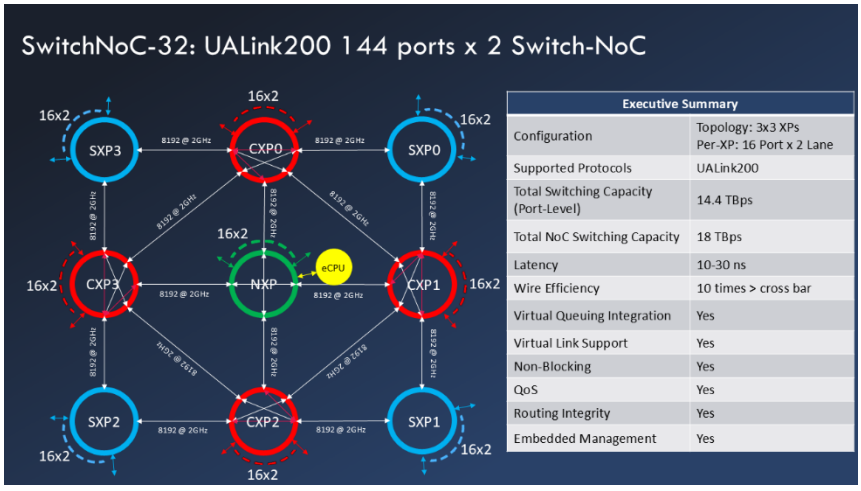

Executive Summary	
Configuration	Topology: 3x3 XPs Per-XP: 16 Port x 2 Lane
Supported Protocols	UALink200
Total Switching Capacity (Port-Level)	14.4 TBps
Total NoC Switching Capacity	18 TBps
Latency	10-30 ns
Wire Efficiency	10 times > cross bar
Virtual Queuing Integration	Yes
Virtual Link Support	Yes
Non-Blocking	Yes
QoS	Yes
Routing Integrity	Yes
Embedded Management	Yes

SiSys AI

2025