WHITE PAPER / KB202608

SDVoE network design from 2 to 1,280 endpoints

A topology and sizing guide for 2 to 1,280 endpoints. The worked examples show port and bandwidth arithmetic. Validate actual multicast forwarding, load distribution and failover before ordering.



Three-tier example. A single-switch design stops at the access switch.

Choose the topology

- One switch: count transmitters, receivers, control ports and spares separately. Use a single 10GbE access switch when the endpoints fit and cable reach allows it. For a 48-endpoint island, verify that control and uplinks have separate usable ports.
- Multiple racks: use access leaves with a central spine. Size each direction for the distinct streams that cross each link. A joined multicast group uses one copy on a shared path until replication is needed. Place heavily shared sources to minimize traffic crossing the fabric.
- Large installations: add aggregation only if port density, distance or failure domains require it. Every new tier introduces another bandwidth and multicast-state limit. Confirm how your design transports the media VLAN or routes multicast. Redundant Ethernet links require a supported design such as LACP, MLAG or a routed fabric.

Niveo hardware roles

- Copper access: N10GSM24XP provides 24 x 10GBASE-T and two 100G uplinks. N10GSM48X2Q provides 48 x 10GBASE-T with four 10G SFP+ and two 100G uplinks. Confirm the actual copper channel and PoE demand per device.
- Fiber access: N10GSM8QF48 provides 48 x 10G SFP+ and eight 100G QSFP28 ports. Match optical mode, reach, loss budget and approved modules at both ends.
- Spine or aggregation: N100G32 has 32 x 100G ports. N400G32 has 32 x 400G ports. Confirm the exact port mode, supported breakout, lane mapping, FEC and optics on both ends. Connector shape alone does not establish breakout compatibility.

Size for traffic and failures

- Copper example: 24 x 10G gives a 240 Gb/s aggregate edge line-rate ceiling per direction. Two 100G uplinks provide 200 Gb/s aggregate, or 100 Gb/s after losing one. Twenty 10G ports match 200 Gb/s arithmetically with no headroom. Use measured stream rates, overhead, expected subscriptions and a documented reserve.
- Fiber example: 40 x 10G gives a 400 Gb/s edge ceiling. Four 100G links to each of two spines leave 400 Gb/s aggregate after a spine loss. Eight such access switches consume all 32 x 100G ports on each N100G32 and provide 320 endpoint ports in use. There are no spare spine ports in this example.
- A link bundle does not provide 200 or 400 Gb/s to a single flow. Check multicast distribution on every member and the switch behavior when a path fails. Confirm forwarding tables, multicast replication, buffer capacity and queue drops.

Worked 1,280-endpoint port budget

- Fiber: 32 access switches x 40 endpoints = 1,280. Four 100G links per access switch require 128 x 100G lanes. A 32-port 400G spine has exactly 128 theoretical 100G breakout lanes if all ports and software support 4 x 100G. This consumes every lane and has no spine failover. A second equivalent spine needs a separate set of lanes and four additional uplinks per leaf for the illustrated failover budget.
- Copper: 64 access switches x 20 endpoints = 1,280. Two 100G uplinks each require 128 spine-facing lanes in a direct fabric. In a three-tier illustration, eight aggregation switches each use 16 x 100G down and 16 x 100G up, and the spine again needs 128 x 100G lanes. No spare ports or single-spine failure capacity is included. These are port budgets, not qualified configurations.
- Reserve space for growth, management and redundant paths. If you cannot confirm the breakout or multicast behavior, size the design around native 100G capacity or another supported topology before committing to a BOM.

Configure and commission

1. Confirm 10GbE links and optics. Apply the endpoint vendor's Ethernet frame size end to end and verify tagged trunks separately. A 9,000-byte IP MTU is not the same as a 9,000-byte Ethernet frame limit.
2. Enable IGMP snooping on multicast media VLANs and verify one active querier per VLAN. Use the endpoint-required IGMP version. Keep Fast Leave off on trunks or shared receiver ports. Verify controller discovery before changing unknown multicast handling.
3. Run peak cross-leaf subscriptions and check per-member throughput, errors, queue drops and switching time. Test cable, spine and power failures under load. A/B endpoint protection requires independent paths and endpoint/application support, even when the core has redundant spines.

References

SDVoE Alliance

Niveo Broadcast & Data Centre portfolio

IETF RFC 4541, IGMP and MLD snooping considerations

Documentation review. Validate with the exact firmware and endpoints before deployment.