

Storage Network Throughput Optimization - Official Technical Overview & Hardware Datasheet

EXECUTIVE SUMMARY

Storage Network Throughput Optimization represents a paradigm shift in how modern data centers and telecom operators architect their storage area networks. As enterprises and service providers scale their digital infrastructure to support AI/ML workloads, real-time analytics, and cloud-native applications, the storage network has emerged as the critical bottleneck in overall system performance. This document provides a comprehensive technical overview of a purpose-built hardware platform engineered to deliver deterministic, ultra-high-throughput storage networking with carrier-grade reliability and sub-microsecond latency.

The platform addresses the fundamental challenges of modern storage networking: the need for wire-speed packet processing at 400GE and beyond, lossless transport for NVMe over Fabrics (NVMe-oF) and RoCEv2 protocols, and intelligent traffic management that ensures fair bandwidth allocation across tenants and workloads. By leveraging a combination of advanced ASIC-based forwarding engines, deep packet buffers, and hardware-accelerated congestion management, the solution delivers up to 51.2 Tbps of non-blocking throughput per chassis with consistent performance under all traffic conditions.

Key performance differentiators include hardware-accelerated RoCEv2 congestion control, dynamic load balancing across multiple paths, and support for industry-standard protocols including NVMe-oF, iSCSI, FCoE, and RDMA over Converged Ethernet. The platform's modular architecture enables seamless scaling from 100GE edge deployments to 400GE/800GE core storage fabrics without disruption to existing services.



ARCHITECTURE & CHASSIS DESIGN

The system is built upon a modular, non-blocking switching fabric architecture that ensures every port can simultaneously operate at full line rate without any internal contention. The chassis design incorporates a high-speed backplane with a total capacity of 102.4 Tbps, supporting full-duplex operation across all

line cards and fabric modules.

The architecture separates the control plane, data plane, and management plane to ensure that control protocol processing never interferes with critical storage traffic forwarding. The data plane is implemented entirely in custom ASIC silicon, delivering deterministic forwarding latency of under 800 nanoseconds for storage-optimized traffic profiles.

A key architectural innovation is the Dynamic Buffer Allocation (DBA) mechanism, which provides up to 128 MB of shared packet buffer per line card. This large buffer pool, combined with intelligent threshold-based allocation, enables the platform to absorb microbursts that are characteristic of storage traffic patterns without dropping a single packet. The platform supports both shared buffer and dedicated per-port buffer configurations, allowing operators to fine-tune behavior for specific workload profiles.

The chassis supports front-to-back and back-to-front airflow configurations, accommodating diverse data center thermal architectures. All major components, including power supplies, fan trays, fabric modules, and line cards, are hot-swappable and support online insertion and removal (OIR) without disrupting traffic forwarding.

HARDWARE FEATURES

STORAGE PROTOCOL ACCELERATION:

Hardware-accelerated NVMe-oF target and initiator offload, iSCSI and FCoE termination, RoCEv2 with hardware-based congestion control (DCQCN), and lossless Ethernet support via Priority Flow Control (PFC) and Enhanced Transmission Selection (ETS).

INTELLIGENT TRAFFIC MANAGEMENT:

Dynamic load balancing across Equal-Cost Multi-Path (ECMP) and Weighted-Cost Multi-Path (WCMP) groups, Application-Layer Traffic Optimization (ALTO) integration, and hardware-based flow hashing with configurable entropy fields. Support for up to 256,000 queues per line card with hierarchical quality of service (HQoS).

TELEMETRY AND VISIBILITY:

Streaming telemetry at 100ms intervals via gNMI/gRPC, hardware-based In-band Network Telemetry (INT) for per-packet path and latency visibility, and support for sFlow and IPFIX at line rate. Dedicated out-of-band management ports with hardware-based security isolation.

RELIABILITY AND AVAILABILITY:

Hitless software upgrades (ISSU), hardware-based Fast Reroute (FRR) with sub-50ms convergence, and redundant control modules with stateful failover. The platform supports Bidirectional Forwarding Detection (BFD) at scale, with hardware-offloaded BFD sessions for up to 16,000 neighbors.

COMPLIANCE & STANDARDS

The platform complies with all relevant industry standards and regulatory requirements for telecom and data center equipment:

SAFETY AND EMC:

- IEC 60950-1 / IEC 62368-1 (Product Safety)
- EN 55032 / EN 55035 (EMC Emissions and Immunity)
- FCC Part 15 Subpart B (Class A)
- VCCI Class A
- BSMI CNS 13438

TELECOM STANDARDS:

- ITU-T G.8032 (Ethernet Ring Protection)
- IEEE 802.1Q / 802.1ad (VLAN and QinQ)
- IEEE 802.1AX (Link Aggregation)
- IEEE 802.1BR (Bridge Port Extension)

- IEEE 802.3ba / 802.3bs / 802.3cd / 802.3cm / 802.3db (Ethernet)
- RoCEv2 (RDMA over Converged Ethernet)
- NVMe over Fabrics (NVMe-oF) Specification 1.1

ENVIRONMENTAL:

- NEBS Level 3 (GR-63-CORE, GR-1089-CORE)
- ETSI EN 300 019 (Environmental Conditions)
- RoHS 2 (EU Directive 2011/65/EU)
- WEEE (EU Directive 2012/19/EU)

TECHNICAL SPECIFICATIONS

The following table summarizes the key technical parameters of the Storage Network Throughput Optimization platform:

Parameter	Specification
Form Factor	1RU / 2RU / 4RU / 8RU Modular Chassis Options
Switching Capacity	Up to 51.2 Tbps per chassis (non-blocking)
Forwarding Latency	< 800 nanoseconds (storage-optimized profile)

Packet Buffer	Up to 128 MB per line card (shared pool)
Power Supply	1+1 Redundant AC/DC, hot-swappable
Cooling	Front-to-back or back-to-front airflow, redundant fan trays
Management Ports	2 x 10GE SFP+ (out-of-band), 1 x Console, 1 x USB
Protocol Support	NVMe-oF, RoCEv2, iSCSI, FCoE, RDMA, PFC, ETS
Telemetry	gNMI/gRPC, INT, sFlow, IPFIX at line rate
Redundancy	Hitless ISSU, hardware FRR, redundant control modules
Operating Temperature	0°C to 40°C (32°F to 104°F), NEBS Level 3
MTBF	> 300,000 hours (Telcordia SR-332)

ORDERING OPTIONS

The platform is available in multiple configurations to address diverse

deployment scenarios, from compact top-of-rack storage switches to high-density core storage fabric chassis. All configurations include the base system software with support for NVMe-oF, RoCEv2, and advanced traffic management features.

BASE SYSTEM CONFIGURATIONS:

- SNT-1000-32C: 1RU fixed-form factor, 32 x 100GE QSFP28 ports, single AC power supply
- SNT-1000-64C: 2RU fixed-form factor, 64 x 100GE QSFP28 ports, redundant AC power supplies
- SNT-4000-128C: 4RU modular chassis, 128 x 100GE or 32 x 400GE ports, redundant fabric modules
- SNT-8000-256C: 8RU modular chassis, 256 x 100GE or 64 x 400GE ports, full redundancy

LINE CARD OPTIONS:

- SNT-LC-32Q: 32-port 100GE QSFP28 line card with hardware NVMe-oF offload
- SNT-LC-16D: 16-port 400GE QSFP-DD line card with hardware RoCEv2 acceleration
- SNT-LC-8E: 8-port 800GE OSFP line card with integrated storage protocol acceleration

- SNT-LC-48X: 48-port 25GE SFP28 line card for edge storage aggregation

SOFTWARE LICENSES:

- SNT-SW-BASE: Base system software with standard storage protocols
- SNT-SW-ADV: Advanced traffic management and telemetry package
- SNT-SW-SEC: Security and encryption acceleration package

