

# Systems Engineering Technical Reference Manual: Server Swap Space Configuration

## PRODUCT IDENTIFICATION


Document Number: TN-SWAP-2026-001

Revision: 1.0

Classification: PUBLIC TECHNICAL WHITEPAPER

Platform Applicability: All x86\_64 & ARM64 Carrier-Grade Server Appliances

This document defines the architectural standards, performance parameters, and engineering best practices for implementing server swap space (paging file) configuration within carrier-grade Linux environments. Swap space serves as an overflow reservoir for volatile memory, providing system stability under transient workload spikes, crash dump facilitation, and enhanced memory overcommit capabilities. Proper configuration directly impacts mean time between failures (MTBF) and quality of service (QoS) metrics.



## SYSTEM HARDWARE TOPOLOGY

The server memory hierarchy is structured as a three-tier capacity model:

TIER 1 – L1/L2/L3 CACHE (CPU-integrated, nanosecond access)

TIER 2 – DRAM MAIN MEMORY (DIMM slots, microsecond access)

TIER 3 – SWAP SPACE (NVMe/SAS/SATA persistent storage, millisecond access)

Swap space acts as the final arbitrage layer before out-of-memory (OOM) termination. The Linux kernel memory manager (kswapd) asynchronously migrates inactive anonymous memory pages to the pre-allocated swap device. Page-in operations occur on demand when applications reference previously paged-out memory regions.

CRITICAL PARAMETER – SWAPPINESS (vm.swappiness): Kernel weighting factor (0–200) that biases page cache reclamation versus swap reclamation. Factory default: 60. Carrier-grade profile: 10–20 for low-latency transaction systems; 80–120 for memory-constrained batch processing.

## DATA & CONTROL PLANE CAPABILITIES

CONTROL PLANE (Memory Management Subsystem):

- Page table walk latency optimization
- Transparent Huge Page (THP) interaction with swap cache
- Swap slot allocation and deallocation via radix tree index

DATA PLANE (I/O Path):

- Block I/O scheduler selection (mq-deadline, kyber, BFQ)
- Swap device queue depth and NVMe polling mode
- Direct I/O bypass for swap read-ahead predictions

## COMPONENT BREAKDOWN

SWAP SPACE IMPLEMENTATION TYPES:

TYPE A – DEDICATED SWAP PARTITION (RECOMMENDED)

Block-level contiguous allocation with no filesystem overhead. Predictable performance and crash dump compatibility (kexec, kdump). Partition type hex code: 82 (Linux swap).

**TYPE B – SWAP FILE ON PERSISTENT FILESYSTEM**

Supported on ext4, XFS, and btrfs. Requires hole-punching and block mapping. Higher fragmentation potential; not suitable for high-frequency paging workloads.

**TYPE C – STRIPED SWAP ACROSS MULTIPLE DEVICES**

Kernel device-mapper (dm-linear) or md RAID0 striping. Increases aggregate I/O bandwidth but introduces single-device failure risk. Use with redundant NVMe namespaces.

<b>Parameter</b>	<b>SMB-Class (1RU)</b>	<b>Enterprise-Class (2RU)</b>	<b>Carrier-Class (4RU)</b>
Minimum DRAM	16 GB	64 GB	256 GB
Recommended Swap Size (General)	8 GB	32 GB	128 GB
Recommended Swap Size	32 GB	128 GB	512 GB

(HPC/Database)			
Swap Device Interface	NVMe Gen3 x2	NVMe Gen4 x4	NVMe Gen5 x8 / SCM
Max Swap Devices per System	8	16	32
Optimal Swappiness Setting	15	20	10
kdump Reserved Swap (MB)	2048	4096	16384
Peak Swap Throughput (MB/s)	3,500	7,000	14,000+

OPERATIONAL SPECS MATRIX

| Metric | SMB-Class (1RU) | Enterprise-Class (2RU) | Carrier-Class (4RU) |

|-----|-----|-----|-----|

| Minimum DRAM | 16 GB | 64 GB | 256 GB |

| Recommended Swap Size (General) | 8 GB | 32 GB | 128 GB |


| Recommended Swap Size (HPC/Database) | 32 GB | 128 GB | 512 GB |  
| Swap Device Interface | NVMe Gen3 x2 | NVMe Gen4 x4 | NVMe Gen5 x8 /  
SCM |  
Max Swap Devices per System	8	16	32
Optimal Swappiness Setting	15	20	10
kdump Reserved Swap (MB)	2048	4096	16384
Peak Swap Throughput (MB/s)	3,500	7,000	14,000+

## REGULATORY COMPLIANCE

The swap space provisioning methodology described herein satisfies requirements set forth by:

- Common Criteria (CC) for memory separation (FDP\_ACF.1)
- NEBS GR-63-CORE: Thermal and vibration resilience for persistent storage
- DISA STIG: Partition allocation and swap encryption mandates
- ISO/IEC 27001: Annex A.8 – Asset management (swap sanitization)
- PCI DSS v4.0: Requirement 3.2 – Sensitive data not written to unprotected swap

Additionally, all swap devices MUST be encrypted using LUKS2 (dm-crypt) with AES-256-XTS to prevent cold-boot forensic recovery of ephemeral secrets. Key management via TPM 2.0 (sealed to PCR 7, 11, 14).



## DEPLOYMENT VALIDATION PROCEDURE

Verification checklist after swap configuration:

1. ``swapon --show`` → Verify active swap spaces
2. ``cat /proc/swaps`` → Confirm priority and size
3. ``sysctl vm.swappiness`` → Validate kernel parameter
4. ``echo 3 > /proc/sys/vm/drop_caches && stress --vm 2 --vm-bytes 90%`` →

Test memory pressure behavior

5. ``dmesg | grep -i "swap"`` → Audit kernel messages for errors

Fault remediation: Swap-related OOM events are logged to ``/var/log/kern.log`` with signature ``Out of memory: Killed process``. Post-event analysis includes ``vmstat 1``, ``sar -S 1``, and ``/proc/meminfo`` delta tracking.

## LIFECYCLE ASSURANCE (MTBF)

Estimated swap device endurance ratings at standard workload (10% memory paging per hour):

- Enterprise NVMe (3 DWPD): 8 years MTBF
- Read-Intensive NVMe (0.3 DWPD): 5 years MTBF
- SATA SSD (Consumer grade): NOT RECOMMENDED – causes IO stall amplification

Firmware-level atomic write protection for swap space is enabled via ``nvme set-feature /dev/nvme0 -f 0x0C -v 1`` (Volatile Write Cache disabled).

END OF DOCUMENT