White Paper
July 2026

Best Practices for RAID Management

Protecting Mission-Critical Data in Government and Enterprise Environments

Co-authors:
Ray Leventhal, general manager CPR Tools
Joe Bruckler, senior engineer CPR Tools
Roger Hutchison, technology officer CD ROM Inc.

Executive Summary

RAID (Redundant Array of Independent Disks) storage systems are widely deployed throughout government agencies, defense contractors, healthcare organizations, financial institutions and commercial enterprises because they provide enhanced storage performance, increased capacity, and varying levels of fault tolerance. However, despite their reliability advantages, RAID systems remain vulnerable to hardware failure, controller corruption, firmware issues, environmental conditions, operator error, ransomware and catastrophic multi-drive events.

When RAID failures occur, the operational consequences can be severe. Agencies and organizations may experience extended downtime, interruption of mission-critical services, loss of sensitive information, legal exposure and substantial recovery costs. In some environments, especially those supporting defense, public safety, healthcare, or industrial operations, RAID failures may directly impact operational readiness and continuity of operations.

This white paper outlines recommended best practices for RAID management designed to reduce risk, improve uptime and increase the probability of successful recovery in the event of failure.

---

1. RAID Is Not a Backup

One of the most common misconceptions in enterprise storage environments is the belief that RAID itself constitutes a backup solution. RAID provides redundancy against certain hardware failures, depending on the RAID level implemented, but it does not protect against:

- Accidental file deletion
- Ransomware attacks
- Malware corruption
- Fire or flood damage
- Power events
- Controller failure
- Firmware corruption
- Insider threats
- Simultaneous multi-drive failures

Organizations should maintain independent and verified backup systems in addition to RAID storage. Best practices include:

- Offline backups
- Immutable backup storage
- Cloud replication
- Offsite backup retention
- Routine restoration testing

A backup system that has never been tested should not be assumed to be reliable.

---

2. Perform Regular RAID Integrity Testing

Routine RAID monitoring and integrity testing is one of the most important preventative measures available.

At a minimum, RAID systems should undergo weekly integrity verification. Mission-critical systems may require daily monitoring and alert review.

Recommended monitoring practices include:

- SMART diagnostics review
- Read/write error analysis

- Controller health monitoring
- Thermal monitoring
- Rebuild status verification
- Event log review
- Automated failure alerting

Modern RAID controllers and enterprise storage platforms typically include predictive failure capabilities that can identify degrading drives before complete failure occurs. Administrators should configure automatic email or SMS alerting whenever possible.

Early detection is often the difference between a simple drive replacement and a catastrophic array failure.

---

3. Replace Aging Drives Proactively

Hard drives and SSDs have finite operational lifespans. Waiting for drives to fail before replacement substantially increases organizational risk.

As drives age, the probability of latent sector errors, rebuild failures, and simultaneous drive degradation increases significantly. This risk becomes particularly dangerous in large RAID arrays where rebuild stress can overload aging drives.

Best practices include:

- Scheduled enterprise drive replacement cycles
- Maintaining identical spare drives onsite
- Using enterprise-grade storage devices
- Tracking drive age and operational hours
- Replacing drives showing early warning indicators

Many organizations implement preventative drive refresh programs every three years depending on environmental conditions, operational workloads, and manufacturer guidance. Spinning media MTF is three years, perhaps less for HAMR and shingle drives. Also, for NVMe SSDs, SMART data should be used for examining reported “bad blocks.”

For NVMe SSDs, the Self-Monitoring, Analysis and Reporting Technology (SMART) reporting system is standardized via the NVMe SMART / Health Information Log. Unlike older SATA drives that use numbered IDs (like Attribute 5 or 197), NVMe drives use specific named fields defined in the NVMe specification.

The following parameters are the primary indicators for "bad blocks" and media health on an NVMe drive, specifically the percentage used as described below:

- Core "Bad Block" Parameters
    - Media and Data Integrity Errors: This field tracks the total number of occurrences where the controller detected a data integrity error. This includes unrecovered ECC errors, CRC checksum failures, or internal device-level failures that resulted in a "bad" or unreadable block.
    - Available Spare: This represents the remaining capacity of the "spare" area on the SSD, expressed as a percentage. SSDs ship with extra NAND blocks to replace those that wear out or become "bad." As the controller identifies bad blocks and swaps them for spares, this percentage drops.
    - Available Spare Threshold: This is a user-defined or factory-set threshold. If the Available Spare falls below this level, the drive triggers a "Critical Warning" to signal that its ability to handle further bad blocks is nearly exhausted.
    - Number of Error Information Log Entries: This is a counter of the total number of entries in the Error Information Log. While not a direct count of bad blocks, a high number often indicates frequent hardware or transmission issues that lead to block retirement.

- Related Health Indicators
    - Percentage Used: While technically an endurance indicator (measuring how much of the drive's rated write life has been consumed), as this number nears or exceeds 100%, the probability of encountering "later bad blocks" (blocks that fail due to wear) increases significantly.
    - Critical Warning: This is a bitmask where specific bits are flipped if the drive's health is compromised. Bit 0 is flipped if the Available Spare has fallen below the threshold, indicating the drive is running out of ways to manage new bad blocks.

- Key Logic Comparison (SATA vs. NVMe)

  In traditional SATA SSDs, you might look for Reallocated Sector Count (ID 05) or Current Pending Sector (ID 197). In NVMe, these are consolidated:

    - Confirmed bad blocks that have been swapped are reflected in the reduction of Available Spare.
    - Failed read/write operations caused by those blocks are counted in Media and Data Integrity Errors.

---

4. Respond Immediately to RAID Warning Signs

RAID systems rarely fail without warning. Common indicators of pending failure include:

- Degraded RAID status
- Rebuild failures
- Increasing bad sectors
- Slow performance
- Controller alerts
- Unexpected drive dropouts
- Unusual drive noises
- Repeated filesystem inconsistencies

Organizations should avoid ignoring these warning signs or continuing heavy production workloads on unstable arrays.

Improper reboot attempts, forced rebuilds, firmware experimentation or unauthorized repair attempts can significantly worsen the condition of a failing RAID system and reduce recovery success rates.

---

5. Engage Qualified RAID Recovery Experts Early

When a RAID system begins exhibiting signs of instability or failure, early intervention by experienced professionals is strongly recommended.

Specialized storage recovery firms such as CPR Tools possess advanced expertise in:

- RAID diagnostics
- Metadata reconstruction
- Multi-drive recovery
- Controller failure analysis
- Firmware incompatibility resolution
- Secure forensic recovery
- Enterprise storage restoration

In many cases, early expert evaluation substantially increases the probability of successful recovery while minimizing additional damage to the array.

Organizations supporting sensitive or classified environments should also ensure that recovery providers maintain secure handling procedures and chain-of-custody protocols where applicable.

6. Maintain Proper Environmental Controls

Environmental conditions directly affect RAID reliability and drive longevity.

Best practices include:

- Proper server room cooling
- Clean airflow management
- Humidity control
- UPS battery backup systems
- Surge suppression
- Redundant power supplies
- Vibration reduction
- Dust mitigation

Excessive heat remains one of the leading causes of premature storage device failure.

Power irregularities, including brownouts and surges, may also damage RAID controllers and corrupt active write operations.

7. Develop Written RAID Management Procedures

Organizations should maintain formal RAID management and incident response procedures that define:

- Monitoring schedules
- Alert escalation paths
- Replacement policies
- Recovery vendor contacts
- Backup verification procedures
- Documentation standards
- Emergency response protocols

Written procedures improve continuity of operations, reduce response time during emergencies, and ensure consistent handling of critical storage infrastructure.

Conclusion

RAID systems remain an essential component of enterprise and government storage infrastructure, but they require disciplined management and proactive maintenance to remain reliable.

Routine integrity testing, preventative drive replacement, environmental controls, verified backups and rapid expert intervention collectively form the foundation of effective RAID management best practices.

Organizations that invest in preventative RAID management significantly reduce operational risk, minimize downtime, and better protect their mission-critical data assets.

For additional information, contact the authors at info@cdrominc.com or support@cprtools.com