

Product Bulletin 111

6-00960-174, Rev. A, September 2026

Bulletin Reference	SNXT-2592 (underlying issue: SNXT-2580)
Title	Flusher Daemons May Fail to Start on SNFS Clients Using Multiple Buffer Caches
Affected Product	StorNext File System (SNFS) clients
Affected Versions	StorNext 7.2.5 and earlier
Fixed In	StorNext 7.2.6
Severity	We can have delayed/ineffective cache flushing when bursts of writes cause buffer cache pressure, and in extreme cases (it has been observed by a single customer) to cause write stall.

Summary

On StorNext File System (SNFS) clients, file systems normally share a single buffer cache and its associated flusher daemon. Certain client-side configurations cause the client to allocate multiple, distinct buffer caches — each of which requires its own dedicated flusher daemon. On systems configured with enough distinct caches, the client could run out of available daemon queue slots. In StorNext 7.2.5 and earlier, when this limit was reached, the affected file system(s) would still mount successfully, but without a functioning flusher daemon and without any warning to the administrator.

Conditions That Trigger This Issue

A distinct buffer cache — and therefore an additional flusher daemon — is allocated per file system mount when either of the following is true:

- **Mixed cachebufsize values** are used across file systems mounted on the same client (for example, one file system mounted with **cachebufsize=512k** and another with **cachebufsize=1m**).
- **The privatebufcache** mount option is used, which explicitly allocates a private cache (and flusher daemon) to that mount.

Flusher daemons are scheduled from a shared daemon queue that is also used for I/O completion and other periodic client tasks. That queue has a fixed capacity, so a client mounting many file systems with distinct caches (heavy use of **privatebufcache** in particular) can exhaust the available flusher daemon slots.

Customer Impact

On affected versions (7.2.5 and earlier), once the flusher daemon queue was exhausted, additional file systems requiring a distinct cache would still mount normally, but their buffer cache would never be flushed by a dedicated daemon. This condition produced no error, warning, or log event visible to the administrator at mount time — the file system appeared healthy and usable while operating without proper write-back cache flushing.

This was identified in production at a customer site making heavy use of **privatebufcache** across multiple file systems on the same client.

Resolution in StorNext 7.2.6

StorNext 7.2.6 addresses this issue in two parts:

1. **Fail-safe mount behavior:** If a flusher daemon cannot be started for a file system's buffer cache, the mount now fails explicitly instead of silently succeeding without a working flusher. This ensures administrators are immediately notified of the condition rather than unknowingly running a file system without proper cache flushing.
2. **Increased flusher daemon capacity:** The number of flusher daemon slots available on the client has been increased, raising the effective limit on file systems that can use **privatebufcache** (or otherwise require a distinct cache) on a single client to 62.

Together, these changes eliminate the silent-degradation scenario: clients can support significantly more distinct-cache file systems than before, and if a configuration ever does exceed the new limit, the affected mount will fail loudly rather than run in a degraded, unflushed state.

Recommended Action

- **Upgrade to StorNext 7.2.6** at the earliest opportunity, particularly on clients that use **privatebufcache** on multiple file systems or mix **cachebufsize** values across file systems.
- **Until upgraded**, review client mount configurations and keep the number of file systems using **privatebufcache** (or non-default **cachebufsize** values) low, to reduce the risk of exceeding the flusher daemon queue on 7.2.5 and earlier.
- **After upgrading**, if any file system fails to mount, check client logs for a "Can't enqueue CvBCFlusher" message — this indicates the number of distinct-cache file systems on that client still exceeds the (now much higher) supported limit of 62, and the mount configuration should be adjusted.

References

Internal tracking: SNXT-2580 (fix), SNXT-2592 (this bulletin).

If needed, contact [Quantum Technical Support](#) to obtain additional assistance to guide you through the required software update.