PANZURA

FOR TECHNICAL DECISION MAKERS

Migrate from NetApp Global File Cache to Panzura CloudFS

Technical Migration Brief for Storage Architects and Infrastructure Engineers

The discontinuation of NetApp Global File Cache (GFC)—recently rebranded as BlueXP Edge Cache—has created what can be seen as a forced transition for storage architects. This requires a critical search for a successor platform that preserves collaborative integrity. While NetApp directs customers toward internal features like Volume Caching, and alternatives like Nasuni, CTERA, or Egnyte struggle to fill the gap, only Panzura CloudFS delivers the real-time global locking and AI-powered Threat Control required for high-concurrency technical workflows.

The Successor Search:

- **GFC rebranded & service discontinued:** Rebranding occurred in Nov. 2022; NetApp removed "Edge Caching" from BlueXP on Aug. 7, 2024
- **The successor search:** Organizations must now choose between ONTAP-native features or a dedicated Global File System
- **FlexCache/Volume Caching:** NetApp's current lead offering potentially creates a single point of failure by suspending global writes if any cache site goes offline.

Zombie Infrastructure: The NetApp GFC Situation

The NetApp GFC technical situation is problematic from our perspective. NetApp Global File Cache (formerly Talon) provided distributed file caching with centralized storage and global file locking. The architecture consolidated branch file servers into a "golden copy" in Azure NetApp Files (ANF) or Cloud Volumes ONTAP (CVO), with Windows-based edge caches at branch locations.

However, legacy GFC deployments are now potentially operating as "zombie" infrastructure:

- No vendor support SLA for critical failures
- No security patches or Windows compatibility updates
- No cloud orchestration updates when Azure/AWS APIs change
- Risk of complete service failure with no recovery path
- Unpatched systems are primary targets for ransomware and malware

FlexCache Technical Analysis

NetApp directs GFC customers to FlexCache or FlexGroup (Volume Caching), technologies designed primarily for read acceleration or capacity scaling rather than high-performance collaboration.

> **From NetApp documentation:** "With global file locking enabled, modifications to the origin volume are suspended until all FlexCache volumes are online. You should only enable global file locking when you have control over the reliability of the connections between cache and origin.

- **The FlexCache Locking Problem:** FlexCache requires all cache sites to maintain connectivity to the origin to permit write operations.
- **Failure Cascade:** If a single branch office loses internet, a global write lock engages, freezing save access for the entire organization.
- **Lacks Granular Locking:** FlexCache lacks the byte-range locking necessary for simultaneous editing in critical applications.

Competitive Landscape: Why CloudFS Wins

Panzura CloudFS is the only platform that provides real-time consistency without the "synchronization lag" found in alternative hub-and-spoke models.

Solution	Why Panzura CloudFS is Superior
Nasuni	Nasuni utilizes a hub-and-spoke architecture that possibly suffers from latency in heavy collaborative environments. CloudFS's peer-to-peer mesh provides faster metadata sync and superior byte-range locking.
CTERA	CTERA is often positioned as a gateway-heavy solution. CloudFS's software-defined approach provides greater flexibility without proprietary hardware lock-in.
Egnyte	Egnyte is primarily an EFSS solution. It potentially lacks the block-level global locking and high-performance local caching required for large-scale CAD/Revit workloads.

Panzura CloudFS Architecture

CloudFS represents an architectural shift from cache-on-top-of-storage to a true global file system. It creates a single unified namespace where every location accesses the same authoritative dataset stored in cost-efficient object storage.

Unlike hub-and-spoke architectures, CloudFS deploys local nodes (often called "filers") that connect to both the cloud object store and to every other node in the network. This full mesh topology enables:

- **No Single Point of Failure:** Individual site failures don't cascade.

- **Peer-to-Peer Data Exchange:** Sites retrieve data directly from each other, not just the origin.
- **Real-Time Metadata Sync:** Every 60 seconds, all nodes within the CloudFS attempt to synchronize changes across the network. Note: This interval is configurable upon consultation with a Panzura expert; metadata snapshots are not stored indefinitely and are subject to retention policies.
- **AI-Powered Threat Control:** Proactively detects ransomware and exfiltration using ML-based behavioral fingerprinting. It interdicts threats by disabling compromised users only, preventing the "blast radius" from freezing site operations.
- **FIPS 140-3 Certified Security:** CloudFS provides military-grade encryption for data at rest and in transit, ensuring compliance for highly regulated engineering and government environments.

Applications Supporting Byte-Range Locking:

- Adobe Illustrator (.ai)
- ArcGIS (.lock)
- Autodesk AutoCAD (dwg .dst .dwl .dwl2 .dbl .bak)
- Autodesk Revit (.dat .rvt .rws .slog)
- Bentley Microstation (.dgn)
- Geopak (.jou .dgn .sem .upf .gpk)
- Microsoft Excel (.xlsx .xlsm)
- Newforma (.pdf .ntz)
- Plant 3D (.dbl .dst .dwg .dwk .dwl .dwt .dwl2 .edf .pspx .pspc-journal .dcf-journal .xml .dcf .plck .pspc)
- Solidworks (.sldprt .sldasm)
- Tekla (.db .db1 .db2 .xml)

Management and Resilience

Technical decision makers can choose from multiple resilience tiers to match site criticality:

- **Instant Node:** Leverages existing virtual resources for <5-minute failover without dedicated hardware.
- **Local & Global HA:** Ensures zero-downtime failover by maintaining hot metadata standbys. Global HA nodes hold only the metadata "card catalog" until called into service, allowing for massive resilience with a low resource footprint.

Capability	NetApp GFC	NetApp FlexCache	Panzura CloudFS
Locking Architecture	Central lock server (latency)	Origin-dependent (brittle)	Peer-to-peer full mesh (resilient)
Locking Granularity	File-level only	File-level only	Byte-range (co-authoring)

WAN Failure Behavior	Possibly degrades over time	Potential global failure	Continues operating
Backend Storage	Block (ANF/CVO)	Block (ONTAP)	Object (S3/Blob)
Duplication	Wire-level only	Volume-level	Global block-level (all sites)
Data Loss Defense	Possibly Passive/Abandoned	Snapshots (reactive)	AI-powered Threat Control (proactive)
RPO	Snapshot intervals (hours)	Snapshot intervals (hours)	60-seconds (standard)
RTO	Hours (restore required)	Hours (restore required)	Sub-5-minutes (failover)
Cloud Flexibility	NetApp clouds only	ONTAP only	Any cloud or on premises
Protocol Support	SMB only	SMB, NFS	SMB, NFS, S3 (native)

Distributed Byte-Range Locking

CloudFS uniquely supports byte-range locking for applications that support it, including Autodesk Revit, AutoCAD, Civil 3D, Tekla Structures, Bentley MicroStation, and Microsoft Excel.

- Data and metadata physically decoupled in architecture
- Every node holds complete metadata copy for entire file system
- When file created, originating node designated as file owner
- Origin tracks current lock holder (Data Owner)
- Data Owner state transported via metadata snapshots
- Multiple users can edit different byte ranges simultaneously

Example Use Case: User A edits the 3rd floor of a Revit model while User B modifies the roof structure. Both work in the same file concurrently without conflict.

Immutable Object Storage Backend

CloudFS stores data in standard object storage (AWS S3, Azure Blob, Google Cloud Storage, on-premises S3-compatible) using a Write-Once-Read-Many (WORM) architecture:

- **Data Never Overwritten:** New blocks appended; metadata pointers updated
- **Ransomware & Malware Immunity:** Encrypted files can’t overwrite base data blocks
- **Recovery = Pointer Reset:** Revert metadata to pre-attack snapshot in minutes
- **60-Second RPO:** Maximum 60-second data loss vs. hours with traditional backup

The Final Recommendation

Traditional file systems break under the weight of global teams. CloudFS changes that dynamic with its distributed architecture. Research from Great Place to Work analyzing 1.3 million employees found that among Fortune 100 Best Companies—97 of which support remote or hybrid work—productivity is 42% higher than typical workplaces. The key enabler is infrastructure that supports seamless collaboration regardless of location.

- **Single Global Namespace:** Every site—London, Munich, Los Angeles—accesses the same file system in real time. No more version conflicts, no more "which file is current?" confusion.
- **Local Performance Everywhere:** CloudFS nodes cache active data for sub-millisecond access at each location. Team members can save minutes and even hours daily in time that can be redirected to productive work. CloudFS ensures they have the same file access performance no matter where they are located.
- **Intelligent File Locking:** Byte-range locking enables true concurrent collaboration without version conflicts. Multiple team members can work on the same file simultaneously without overwriting each other's changes.
- **Efficient Synchronization:** Only changed data blocks move between locations; everyone works on the same authoritative version. This dramatically reduces WAN bandwidth requirements while keeping all sites current.
- **Remote Workforce Ready:** Field engineers, media teams, and remote staff access files securely without VPN latency. With 30% of distributed workers citing collaboration difficulties in distributed environments, CloudFS eliminates the underlying infrastructure barriers.

Measured Outcomes:

- 40% faster multi-site project completion
- 75% fewer file access support tickets
- 50–80% reduction in redundant data copies

Compliance & Governance, Automated

The deprecation of NetApp GFC forces a platform decision. Moving to FlexCache sacrifices the collaborative locking integrity that distributed AEC, manufacturing, and engineering workflows require. The "all-sites-online" constraint creates an unacceptable brittleness for organizations with global operations.

CloudFS provides the architectural evolution needed to eliminate single points of failure while transforming the storage economic model.

Get an immediate Proof of Concept for Panzura CloudFS focusing on:

- Validating byte-range locking for complex applications and workflows
- Testing Threat Control interdiction to verify blast radius minimization
- Confirming Resilience during simulated WAN failures

This analysis is based on publicly available information, vendor documentation, industry research, and independent technical evaluations. Organizations should conduct their own assessments based on specific requirements and environments. All product and company names are trademarks or registered® trademarks of their respective holders. Use of those names does not imply any affiliation with or endorsement by their owners. The opinions expressed above are solely those of Panzura LLC as of January 7, 2026, and Panzura LLC makes no commitment to update these opinions after such date.