ISSN: 2096-3246
Volume 58, Issue 01, 2026

CLOUD REPATRIATION: A SNAPMIRROR-DRIVEN MIGRATION FRAMEWORK FROM AZURE CLOUD VOLUMES ONTAP TO ON-PREMISES NETAPP STORAGE

Arjun Bhargav Devan
Independent Researcher, USA

Abstract

As enterprise storage deployments at cloud scale reveal persistent cost overruns, a growing cohort of organizations is pursuing cloud repatriation: the structured migration of workloads from the public cloud back to on-premises infrastructure. Cloud Volumes ONTAP (CVO), deployed on Microsoft Azure, delivers the full ONTAP feature set over Azure Premium SSD managed disks under a capacity-based licensing model. While operationally familiar, its all-in annual expenditure for high-capacity, high-availability workloads significantly exceeds the equivalent on-premises total cost of ownership (TCO) over a multi-year horizon. This article presents a technically rigorous examination of the repatriation pathway using SnapMirror as the native migration engine. The analysis covers the economic case for repatriation with a quantitative cost comparison; the SnapMirror Extended Data Protection (XDP) replication architecture and its block-level transfer mechanics; a phase-structured migration procedure from pre-assessment through cutover; the reverse-resync strategy that converts the vacated CVO instance into a warm disaster recovery standby; and post-migration governance. The article draws on foundational distributed storage design principles, scalable system architecture theory, and recent empirical studies on cloud data placement optimization, migration downtime reduction, and distributed file system performance to contextualize the SnapMirror methodology within the broader enterprise storage literature. The article demonstrates that SnapMirror-based repatriation eliminates third-party tooling dependency, preserves the full snapshot history on the destination, and constrains user-visible downtime to a window of fifteen to thirty minutes.

Keywords: Cloud Volumes ONTAP, SnapMirror, Cloud Repatriation, On-Premises Storage Migration, NetApp ONTAP, Azure ExpressRoute, Total Cost of Ownership, Disaster Recovery, Extended Data Protection, Data Sovereignty

1. Introduction

Enterprise storage strategy in the mid-2020s is increasingly shaped by the accumulated financial consequences of cloud adoption at scale. Cloud Volumes ONTAP on Microsoft Azure offered organizations the complete ONTAP data management platform, including SnapMirror, deduplication, compression, and FlexClone, deployed as a software-defined storage layer on Azure Premium SSD managed disks under a capacity-based per-gibibyte (GiB) licensing model [1]. The operational appeal was genuine: teams familiar with on-premises ONTAP could replicate their workflows without retraining, and the absence of upfront capital expenditure removed a traditional procurement barrier. What early business cases rarely modeled accurately, however, were the compounding costs of sustained, high-capacity operation across multiple simultaneous billing dimensions, including software licensing, compute, managed disk provisioning, and data egress [18]. The financial divergence between projected and actual CVO costs has become one of the primary drivers of cloud repatriation, a trend now documented across industry surveys and CIO research: a 2024 survey of enterprise technology leaders found that 83% of CIOs planned to shift at least some workloads from the public cloud to private or on-premises environments [17].

The architectural foundations of distributed storage, established by landmark works including the

design of redundant disk arrays [2], the Ceph distributed file system [3], cluster-based scalable network services [4], the Google File System [5], and the Hadoop Distributed File System [6], all share a common principle: storage system design must align data access patterns with the physical and economic properties of the underlying storage medium. CVO on Azure inherits the familiar ONTAP interface but overlays it on Azure-managed disks, whose cost model is optimized for elasticity and short-term workloads rather than the sustained, high-throughput patterns typical of enterprise primary storage. When workload characteristics are well understood and data volumes are large and stable, the economics of on-premises ownership consistently outperform the pay-as-you-go cloud model, a finding supported by recent empirical work on cost optimization in hybrid multi-cloud environments [19].

What makes the CVO-to-on-premises migration structurally distinct from heterogeneous cloud migrations is the tool it can use: SnapMirror. Because both source (CVO on Azure) and destination (on-premises ONTAP array) run the same ONTAP operating system, SnapMirror functions as a native, zero-additional-cost migration engine. It replicates data by transferring snapshot deltas at the block level, preserves the full snapshot history on the destination, and enables cutover through a break-and-mount operation that any ONTAP-certified administrator can execute [8], [9]. This stands in sharp contrast to heterogeneous migrations, which require third-party tooling, format conversion, and extended parallel-run periods. This article examines the repatriation pathway in depth, covering the economic rationale in Section 2, the SnapMirror replication architecture in Section 3, the migration procedure in Section 4, and the reverse-resync strategy and post-migration governance in Section 5.

2. Cost Economics of Cloud Storage at Scale

2.1 Azure CVO Total Cost of Ownership

A CVO deployment on Azure accumulates charges from four distinct billing dimensions simultaneously, and it is the compounding interaction of those dimensions that drives the TCO to levels that frequently exceed pre-migration projections. The first dimension is the ONTAP software license. Under the capacity-based pricing model, the Essentials tier carries a per-GiB-per-month charge, while the Professional tier, which includes unlimited BlueXP Backup and Recovery to Azure Blob, carries a substantially higher per-GiB fee [1]. Both tiers are billed against provisioned capacity rather than logical utilization, meaning that ONTAP deduplication and compression savings reduce the logical data footprint but do not reduce the provisioned GiB against which the license meter runs [1].

The second dimension is Azure compute. CVO requires DS-series virtual machine (VM) instances to host the ONTAP controller process. A high-availability (HA) pair, which most production deployments require for resilience, runs two such instances continuously. The third and most substantial dimension is Azure Premium SSD managed disk provisioning. CVO aggregates are constructed on P-series managed disks, which are billed per GiB per month regardless of the logical data occupancy on the volume. The fourth dimension is data egress: Azure charges for every gigabyte transferred out of the Azure environment, including replication traffic to on-premises disaster recovery (DR) targets and application data access from systems outside the region. These four dimensions interact to produce an all-in annual cost for a 100 terabyte (TB) HA CVO deployment that is substantially higher than the equivalent on-premises TCO, a gap that widens as data volumes grow [18].

Research on cost optimization in hybrid multi-cloud environments confirms that the pay-as-you-go

model imposes a structural cost penalty on organizations running stable, always-on workloads with predictable demand profiles. Zibitsker and Lupersolsky found that for such workloads, a well-designed hybrid architecture combining on-premises compute and storage with selective public cloud usage achieves superior cost-per-transaction ratios compared to a purely public cloud deployment, particularly once the break-even point for on-premises capital expenditure is reached [19]. Similarly, Goswami documented that for organizations with high-capacity storage workloads, hidden costs including egress fees, cross-zone traffic, and layered managed services erode the apparent savings of the cloud subscription model, making repatriation financially tangible rather than merely theoretical [18].

2.2 On-Premises NetApp Array Cost Model

An on-premises NetApp FAS or AFF array operates on a capital expenditure (CapEx) model that is structurally different from the CVO subscription. The hardware is acquired once and amortized over a five- to seven-year lifecycle. There is no per-GiB software licensing charge beyond the initial ONTAP bundle, no compute VM rental, and no egress cost for data accessed by any connected system. ONTAP licensing, which includes SnapMirror, SnapVault, FlexClone, FabricPool, and cloud tiering capabilities, is bundled with the hardware purchase [8]. Annual ongoing costs after acquisition consist of the hardware support contract, power and cooling, and staff allocation for administration, which for a system sharing the same ONTAP operating system as CVO requires no incremental retraining overhead [9].

The capital intensity of the on-premises model creates a financial profile distinction from the cloud subscription: the organization bears the hardware acquisition cost in the first year, whereas CVO charges are distributed uniformly across months. For organizations with predictable, sustained, high-capacity storage workloads, this distinction resolves decisively in favor of on-premises ownership when analyzed over a multi-year horizon [19]. Llorente documented that the cost economics of cloud versus on-premises shift at a workload tipping point: cloud is advantageous for elastic, unpredictable, or short-tenure workloads, but once usage patterns stabilize and data volumes are large and consistent, private infrastructure delivers lower unit costs [17]. Microsoft itself acknowledges this dynamic in its Azure hybrid guidance, noting that organizations frequently achieve optimal cost profiles through hybrid architectures that retain stable, high-capacity workloads on-premises while leveraging the cloud for elastic or burst requirements [16]. The break-even calculation measures the point at which avoided CVO subscription costs recover the sum of the on-premises hardware acquisition and migration project costs. At the conservative mid-range of these estimates, avoided CVO costs recover the total acquisition and migration investment within a small number of months. Once the break-even point is surpassed, each additional year of on-premises operation results in significant avoided costs that accumulate over the hardware lifecycle [19].

2.3 Non-Financial Drivers of Repatriation

Cost is the primary driver, but several non-financial factors reinforce and in some cases independently justify the repatriation decision. Latency is a persistent structural limitation of CVO: Azure Premium SSD performance tiers govern the input/output operations per second (IOPS) and latency floor of any CVO volume, while inter-service network hops within the Azure region add variable delay that is absent from direct-attached on-premises storage. Distributed storage research has well established the significance of latency and throughput consistency: Wang et al. demonstrated that load-aware directory migration in distributed file systems improved IOPS by 67.4% and bandwidth by 69.2% compared to state-of-the-art methods, with a completion time reduction of 41.6%, underscoring that storage access pattern alignment with physical infrastructure

ISSN: 2096-3246
Volume 58, Issue 01, 2026

has quantifiable, large-magnitude performance consequences [15]. For latency-sensitive workloads, including databases and high-frequency analytics, the consistency offered by a dedicated on-premises all-flash array represents a functional improvement alongside the cost reduction.

Data sovereignty presents a second non-financial driver of growing regulatory weight. The EU Data Act (Regulation (EU) 2023/2854), which entered into force in January 2024 and became fully applicable in September 2025, establishes requirements for cloud providers to implement technical and organizational measures preventing unlawful non-EU governmental access to data stored in EU cloud regions [7]. For regulated organizations, on-premises provides clear data residency control. The cloud doubles down on data sovereignty complexities where multiple jurisdictions are at play, placing conflicting national laws in opposition to cloud provider policies, creating compliance risks better addressed by on-premises deployments [14]. The combination of regulatory pressure from the EU Data Act [7] and the documented financial case for stable workloads [18] creates a compounding motivation for repatriation that transcends pure cost analysis.

3. SnapMirror Architecture and Replication Mechanics

3.1 XDP Relationship Model and Block-Level Transfer

SnapMirror is an ONTAP-native replication engine that operates at the storage volume level by replicating the delta between successive ONTAP Snapshot copies from a source volume to a destination volume [9]. Because both CVO and on-premises ONTAP instances run the same underlying operating system and speak the same replication protocol over TCP ports 11104 and 11105, no format conversion, agent installation, or third-party tooling is required at either endpoint. This is a structural advantage that distinguishes same-ONTAP migration from heterogeneous approaches. Ansley documented that SnapMirror replication is policy-driven and schedule-based, with the policy framework governing which Snapshot copies are transferred and retained at the destination and the schedule governing the frequency of incremental delta transfers [8].

Modern ONTAP deployments use the Extended Data Protection (XDP) relationship type, which supersedes the legacy Data Protection (DP) type. XDP supports mirror, vault, and combined mirror-vault protection within a single policy framework [20]. For the CVO-to-on-premises migration, an XDP relationship using the MirrorAllSnapshots policy is the recommended configuration: it replicates every Snapshot copy present on the source volume to the destination, ensuring that the on-premises array arrives with a complete point-in-time recovery history on the first day of production operation [9]. This is a significant operational advantage over file-copy utilities, which transfer only the current data state and provide no historical recovery points on the destination until new snapshots accumulate after cutover.

Block-level delta transfer is the property that most distinguishes SnapMirror from file-copy approaches in a large-scale migration context. SnapMirror tracks changes at the 4 kilobyte (KB) block level using the change log maintained by the ONTAP write-anywhere file layout (WAFL) file system and transfers only those changed blocks between successive Snapshot copies [9]. This means that incremental transfer volume is determined by the data change rate, not by the number of files. This characteristic is relevant to the broader distributed storage literature: the scalability challenges of file-level enumeration in large distributed systems were documented as far back as the design of the Google File System, where consistent, scalable data access required moving away from per-file operation semantics toward larger data units [5]. SnapMirror resolves the same enumeration scalability problem at the block level, making its throughput characteristics independent of namespace complexity. Figure 1 illustrates the XDP replication architecture and the role of the

ISSN: 2096-3246
Volume 58, Issue 01, 2026

common snapshot anchor.

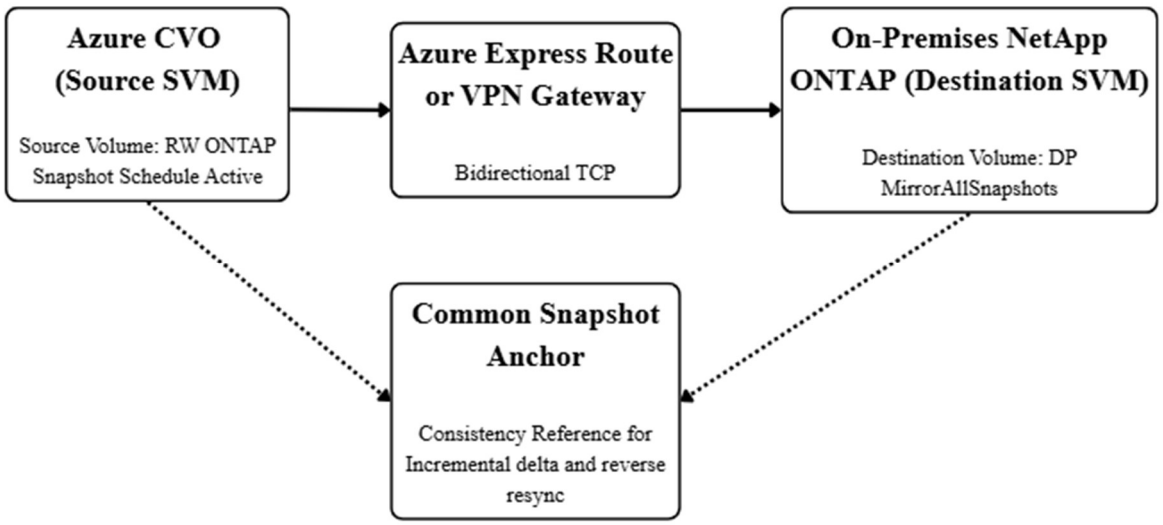

Figure 1: SnapMirror XDP Replication Architecture: Forward Migration and Reverse Resync [Author's Synthesis from 8, 9]

3.2 Network Connectivity and Intercluster Configuration

SnapMirror replication from Azure CVO to an on-premises ONTAP cluster requires network connectivity in both directions. The two transport options are Azure ExpressRoute and Azure VPN Gateway. The recommended way to migrate the data is to use ExpressRoute, a private and dedicated circuit from the on-premises data center to the Azure region. It provides consistently low-latency bandwidth and no internet egress charges for the replication traffic. ExpressRoute is the preferred option for migrating high-throughput workloads where consistent performance and predictable pricing are important [16]. Azure VPN Gateway is lower cost for smaller workloads or organizations that do not already have an ExpressRoute circuit but with a lower maximum throughput limit and standard internet egress charges for replication traffic between on-premises and Azure.

Both transport options require that TCP ports 11104 and 11105 be open bidirectionally between the on-premises firewall and the Azure Network Security Group protecting the CVO virtual network [8]. These ports carry SnapMirror control and data traffic respectively and must remain open for the entire duration of the migration run-up period. Ansley documented that intercluster logical interfaces (LIFs) must be configured on both clusters and must be reachable from the opposite cluster before any SnapMirror relationship can be established [8]. These intercluster LIFs are dedicated network interfaces that carry only cluster peering and replication traffic. Their IP addresses must be in subnets that are routable across the ExpressRoute or VPN connection. The data placement and scheduling work of Zhang et al. on edge-cloud collaborative systems is instructive here: their analysis showed that the cost and feasibility of data movement between storage tiers depend critically on the bandwidth and latency characteristics of the connecting network and that data placement optimization must account for transfer constraints as a primary design parameter [11].

3.3 Cluster and SVM Peering Prerequisites

Cluster peering is the foundational trust relationship in any ONTAP-to-ONTAP replication topology. It authenticates the two clusters to each other and authorizes data replication between their Storage Virtual Machines (SVMs). The peering process requires confirming intercluster LIF reachability,

ISSN: 2096-3246
Volume 58, Issue 01, 2026

exchanging a cluster peer passphrase for mutual authentication, and verifying that the peer relationship status is Available [8]. A failed reachability check at this stage, most commonly caused by a missing firewall rule or an incorrect route in the on-premises network, prevents SnapMirror initialization regardless of all other configurations. SVM peering is then established between the specific source SVM on CVO and the destination SVM on the on-premises array, specifying snapmirror as the allowed application [20].

The peering configuration is a one-time setup that supports any number of subsequent SnapMirror volume relationships between the paired clusters. This architectural pattern, where a single authenticated trust relationship enables multiple data movement operations without per-volume re-authentication, reflects a principle documented in scalable distributed system design: the separation of control-plane authentication from data-plane transfer reduces overhead and enables horizontal scaling of the replication topology [4]. The Hadoop Distributed File System employed a similar architectural separation between its NameNode control plane and DataNode data transfer paths to achieve scalable throughput without proportional control-plane overhead [6]. Table 3 summarizes the complete migration phase structure, including peering as Phase 1, with key activities and success criteria.

Phase	Title	Key Activities	Key ONTAP Commands	Success Criteria
1	Pre-Migration Assessment	Inventory CVO volumes, export rules, SMB shares, and AD integrations; create destination SVM; provision aggregates; open TCP 11104/11105; and verify LIF reachability.	cluster peer create vserver peer create	Cluster and SVM peer status = Available
2	SnapMirror Initialisation	Create DP destination volumes; create XDP relationships with the MirrorAllSnapshots policy; trigger baseline transfer; monitor progress.	snapmirror create -type XDP snapmirror initialize	Baseline completes; lag-time reported; source Snapshots present on destination
3	Incremental Run-Up (1–4 weeks)	Monitor hourly incremental updates, resolve unhealthy-reason errors, configure NFS/SMB on the destination, and test client mounts.	snapmirror show -fields lag-time snapmirror update	Lag returns to 00:00:00 after each cycle; there are zero errors; test mounts succeed
4	Final Update and Cutover	Notify stakeholders, quiesce writes, run the final update, verify zero lag, break the relationship, mount the volume, update DNS/DFS, and validate client access.	snapmirror quiesce snapmirror update snapmirror break volume mount	Lag time = 00:00:00 before break; destination RW; all clients validated within 15–30 min window
5	Reverse Resync	From on-premises (now the source), resync CVO (now the destination) using a common snapshot anchor; no full re-baseline is required.	snapmirror resync -quick-resync true	CVO volume = DP; reverse lag within schedule interval
6	Observation and Decommission	30-day production stability, DR test break-and-mount, application owner sign-off, delete CVO Working Environment in BlueXP, and verify Azure billing stops.	BlueXP: Delete Working Environment	All three decommission criteria met; Azure compute, disk, and licensing billing terminated

Table 1: Migration Phase Overview [8, 9, 10, 20]

4. Migration Procedure: Phases and Execution

4.1 Pre-Migration Assessment and Environment Preparation

ISSN: 2096-3246
Volume 58, Issue 01, 2026

A rigorous pre-migration assessment determines the difference between a predictable cutover and an extended, problematic one. The assessment begins with a complete inventory of the source CVO environment: every volume name, logical size, Snapshot policy, junction path, access protocol (Network File System (NFS) or Server Message Block (SMB)), export policy rule set, SMB share definitions, Active Directory (AD) integration, and application-specific storage service settings. This inventory is the blueprint from which the destination SVM on the on-premises array is configured. Every access control, export rule, and share permission on the CVO source must be reproduced on the on-premises destination SVM before cutover, because the cutover sequence itself does not allow time for configuration discovery [20].

On the on-premises array, the destination SVM is created with the same name and language settings as the source SVM to preserve NFS client mount path continuity and SMB share access path consistency. Aggregate provisioning must account for the full raw volume capacity at the volume level, because SnapMirror replicates at the volume level before storage efficiency savings are visible in the destination aggregate space report [9]. Firewall rules permitting TCP 11104 and 11105 between the on-premises intercluster LIF addresses and the CVO intercluster LIF addresses are opened, and reachability is verified with a port-level connectivity test before any peering commands are issued [8]. The preparation phase also covers DNS and namespace configuration staged for activation at cutover: DNS A-record changes and Distributed File System (DFS) namespace target updates are prepared but not applied, enabling them to be activated within seconds of the SnapMirror break command during the maintenance window.

The importance of thorough pre-migration preparation is well supported in the migration literature. Manda documented that in a large-scale Oracle database migration to cloud infrastructure, downtime was limited to one hour through the use of careful pre-migration preparation combined with Oracle GoldenGate for change data capture, and that query performance improved by 30 percent post-migration alongside a 40 percent reduction in administrative overhead [13]. While the toolchain differs for NetApp ONTAP, the preparation discipline is analogous: success in the cutover window depends almost entirely on the completeness of the preparation work performed in the weeks preceding it.

4.2 Baseline Initialisation and Incremental Synchronisation

For each volume being migrated, a destination volume of type DP is created on the on-premises array, and an XDP SnapMirror relationship is established between the source CVO volume and the destination on-premises volume. The snapmirror initialize command triggers the baseline transfer, which copies the complete contents of the source volume as an ONTAP Snapshot [9]. The baseline is the most network-intensive operation in the migration and must be scheduled well before the planned maintenance window. Hibbard documented that the SnapMirror replication workflow follows the same sequence for all XDP relationship types: create a destination volume, create a job schedule, specify a policy, and create and initialize the relationship [20].

```
# On an on-premises ONTAP cluster (destination):
volume create -vserver onprem-svm -volume vol_data01 -aggregate aggr1 -size 10T -type DP
snapmirror create -source-path cvo-svm:vol_data01 \
  -destination-path onprem-svm:vol_data01 \
  -type XDP -policy MirrorAllSnapshots -schedule hourly
snapmirror initialize -destination-path onprem-svm:vol_data01
```

ISSN: 2096-3246
Volume 58, Issue 01, 2026

Once the baseline completes, the SnapMirror schedule drives hourly incremental updates. The primary monitoring metric during the run-up period is the lag time: the age of the oldest unconfirmed Snapshot at the destination, reported by the snapmirror show command [9]. A healthy incremental schedule produces lag times that return to zero after each update cycle completes. If lag grows across successive cycles, the incremental transfer duration is exceeding the schedule interval, and the cause must be investigated and resolved before the cutover window is entered. During this period, all destination-side configuration is completed: Before the cutover day, validate NFS export policies, SMB shares, AD domain joins, and test mounts from representative client systems to the on-premises data LIFs to confirm that protocol access works. Figure 2 maps the cutover decision flow, including the go/no-go criteria at each phase.

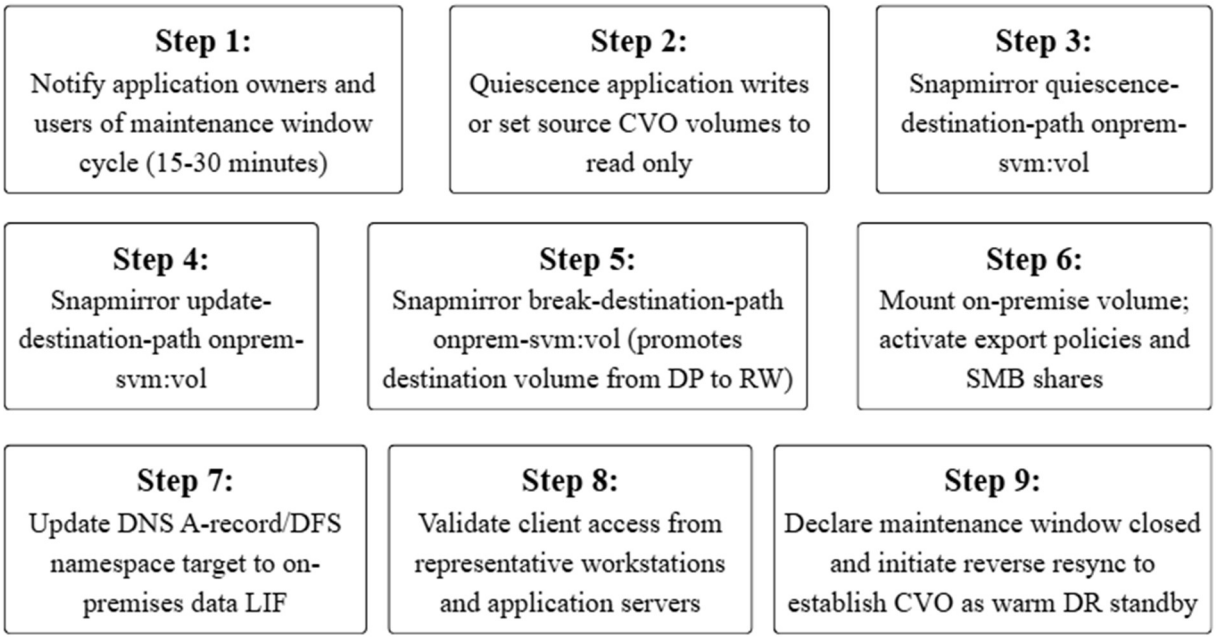

Figure 2: Cutover Execution Timeline and Decision Flow [9, 13, 20]

4.3 Cutover Execution and Validation

The cutover sequence is executed within a pre-approved maintenance window. The target window for a well-prepared single-volume cutover is fifteen to thirty minutes [8]. The sequence begins with stakeholder notification, followed by application write quiescence. With writes quiesced, the SnapMirror quiesce command pauses the SnapMirror schedule on the destination relationship, and the SnapMirror update triggers a final transfer that captures all changes since the last scheduled update. The critical validation at this point is the lag-time field in the snapmirror show output: it must reach 00:00:00 before the break command is issued [9]. A non-zero lag at break indicates that the on-premises destination is not fully up to date compared to the source, leaving a gap.

```
snapmirror quiesce -destination-path onprem-svm:vol_data01
snapmirror update -destination-path onprem-svm:vol_data01
# Wait for transfer to complete, then verify zero lag:
snapmirror show -fields lag-time   # Target: 00:00:00
snapmirror break -destination-path onprem-svm:vol_data01
```

ISSN: 2096-3246
Volume 58, Issue 01, 2026

Once lag is resolved, the snapmirror break command promotes the destination volume from DP read-only mode to read-write mode. The volume is mounted with its junction path, export policies and SMB shares are activated, and the DNS record or DFS namespace target is updated to point clients to the on-premises data LIF [20]. The final step before cutover is testing access from representative workstations and application servers to mount, read, and write to the on-premises volume. This phased go/no-go gated approach to cutover is also a reflection of the lifecycle discipline defined in the enterprise migration study, where structured evaluation, proof of concept, and staged execution phases are primary drivers of migration success and downtime mitigation [13].

The secure migration approach documented by Altahat et al. for containerized workloads is instructive as a parallel: their two-stage scheduling and migration solution achieved a reduction in running host count of 4% to over 18% through careful pre-migration optimization, with power consumption savings of 3.5 to 16.25 megawatts [12]. While the workload type differs, the principle transfers: the quality of the pre-migration optimization and scheduling work determines the magnitude of the benefits realized at and after cutover.

5. Reverse Resync, Risk Mitigation, and Post-Migration Governance

5.1 Reverse Resync and the CVO Warm Standby Model

The most useful feature of SnapMirror in this context is the ability to reverse the direction of replication after the cutover without having to baseline from scratch. Once the on-premises volume is serving the production traffic, the old CVO volume is kept in a dormant state with a snapshot of its data from the time the relationship was broken. A reverse-resync operation reestablishes the SnapMirror relationship in the reverse direction such that the on-premises volume is the source and the CVO volume is the destination. A reverse resync operation uses the common Snapshot anchor retained during the original forward replication operation [10]. As both volumes share this anchor as a reference point, only the delta that has occurred on the on-premises primary since the break is replicated, dramatically reducing the bandwidth and time that would otherwise be needed to re-baseline [10].

```
# On the on-premises cluster (now source), after cutover:
snapmirror resync -source-path onprem-svm:vol_data01 \
  -destination-path cvo-svm:vol_data01 -quick-resync true
```

Vorwerk documented that the resync operation uses the most recent common Snapshot to compute the minimum delta required to bring the destination into consistency with the source, and that the -quick-resync true parameter activates this optimized path for XDP relationships [10]. Once the reverse resync completes, the on-premises array is the SnapMirror source (primary), and the CVO instance is the destination (warm standby). This converts what would otherwise be a stranded, decommissioned cloud resource into a productive risk-management asset during the post-migration observation period.

The design principle of retaining a consistent secondary copy during a transition period reflects a pattern well established in distributed storage systems research. The foundational Ceph architecture maintained data availability through replication across placement groups, with the replication state machine managing consistency across primary and replica objects throughout any topology change [3]. The RAID model similarly established that maintaining redundant copies during transitions is

ISSN: 2096-3246
Volume 58, Issue 01, 2026

the defining property that separates recoverable from unrecoverable failure modes [2]. SnapMirror reverse-resync implements this same principle applies at the volume replication level: the organization never reduces its DR posture below dual-site consistency during the entire observation period.

5.2 Rollback Procedure and Risk Boundaries

The reverse-resync relationship provides a clean, tested rollback path if issues emerge after cutover. The scenarios that most commonly trigger rollback consideration are an unanticipated application dependency on CVO-specific Azure connectivity, a performance regression on the on-premises array due to misconfigured storage efficiency policies or incorrect Quality of Service (QoS) settings, or an access control mapping error preventing a specific user population from reaching their data. In all these cases, the rollback procedure mirrors the original cutover sequence: quiesce writes on the on-premises primary, trigger a final reverse update, break the reverse relationship to promote the CVO volume back to read-write, and update DNS or DFS targets to redirect clients back to the CVO path [10].

Because the common Snapshot anchor eliminates the need for a full re-baseline, this rollback can be completed within the same fifteen to thirty minute window as the original cutover, provided that the on-premises primary has been running for less than the Snapshot retention policy allows before the common anchor is expired [10]. The CVO instance should not be decommissioned until the on-premises system has been stable in production for the agreed observation period, all application owners have confirmed full functionality on the on-premises path, and the reverse-resync relationship has been validated through a test break-and-mount exercise. When decommissioning proceeds, the CVO working environment is deleted through the BlueXP console, which terminates the Azure VM instances, disassociates the managed disks, and stops the ONTAP licensing meter [1].

The data placement strategy of Zhang et al. provides a relevant parallel for this staged approach. Their DE-PSO algorithm for edge-cloud collaborative storage placement was designed under deadline constraints, explicitly modeling the cost and risk of data residing in a non-optimal location during the transition from one placement scheme to another [11]. The observation period following SnapMirror cutover serves the same function: it is the period during which the organization validates that the new placement is fully functional before permanently releasing the previous one, managing transition risk through temporal overlap rather than a hard cutover with no fallback.

5.3 Post-Migration Governance and Performance Optimisation

A successful cutover is not the end of the migration project. The first seventy-two hours after cutover require active monitoring of the on-premises array performance characteristics: IOPS, throughput, latency at the storage layer, and cache hit rates for AFF systems. Workloads that exhibited variable latency under Azure CVO will typically show improved and more consistent response times on a dedicated on-premises all-flash array, as the Premium SSD performance tiers and inter-service network hops introduce a minimum and maximum latency [8]. The IOPS improvement and latency consistency documented by Wang et al. in their load-aware distributed study of file system migration, where IOPS improved by 67.4% and completion time was reduced by 41.6% after alignment of data placement with actual access patterns, provide a reference frame for the magnitude of performance improvement achievable when storage architecture is aligned with workload characteristics [15].

Reconfiguring the data protection policy is a mandatory post-migration task. CVO backup policies managed through BlueXP Backup and Recovery, which vault Snapshot copies to Azure Blob Storage, are no longer operational once the CVO instance is decommissioned [1]. On-premises data protection must be re-established through SnapVault to a secondary on-premises ONTAP system,

SnapMirror to a cloud destination for off-site DR, or integration with existing NDMP backup infrastructure [8]. The selected option should be activated and validated with a test restore during the observation period, not deferred until after CVO is decommissioned. Storage efficiency validation confirms that deduplication, compression, and compaction are active on all migrated volumes. FabricPool tiering provides an additional mechanism for reducing the effective capacity footprint of cold data on on-premises ONTAP systems with a cloud tier configured to Azure Blob or compatible object storage [8].

The financial validation of the repatriation decision should be formalized through structured post-migration reviews at thirty and ninety days. The thirty-day review compares the Azure invoice for the month following CVO decommission against the pre-migration baseline and quantifies the avoided expenditure. The ninety-day review extends the review period across the first quarter and examines whether on-premises performance, data protection, and storage efficiency outcomes align with the pre-migration TCO model [19]. Goswami documented that organizations that conducted structured post-repatriation cost reviews consistently found that avoided egress fees, managed disk charges, and software licensing costs materialized as projected, with the primary source of deviation being underestimated on-premises support contract costs [18]. Table 2 provides the post-migration governance checklist with quantitative targets.

Governance Area	**Metric**	**Target / Threshold**	**Validation Method**
Performance (0–72 h)	Storage latency	Meet or exceed CVO baseline, consistent with dedicated all-flash characteristics	ONTAP System Manager; Active IQ latency dashboard
Performance (0–72 h)	IOPS and throughput	Workload IOPS aligned with pre-migration model; bandwidth improvement with aligned storage placement	volume statistics show active IQ performance graphs
Data Protection	Backup policy active	Policy running on all migrated volumes before CVO decommission; RPO meets SLA	The snapshot policy shows SnapVault relationship status
Data Protection	Test restore	Successful restore within RTO; no checksum mismatch	Manual restore + MD5/SHA-256 checksum validation
Storage Efficiency	Dedup + compression ratio	Efficiency policies active; savings ratio matches pre-migration TCO model	volume efficiency show
DR Reverse Resync	CVO warm standby lag	Lag within hourly schedule interval; common Snapshot anchor preserved; no re-baseline triggered	snapmirror show -fields lag-time on on-premises cluster
Cost — 30-day review	Azure invoice reduction	Avoided CVO costs materialise: egress, managed disk, compute, and licensing charges eliminated	Azure Cost Management portal: month-over-month comparison
Cost — 90-day review	On-premises TCO vs. model	Actual costs within acceptable deviation of pre-migration estimate; primary variance is support contract	Internal finance reconciliation; formal 90-day review
Operational Stability	Support ticket	No sustained spike in storage-related incidents beyond the 72-hour transition	Helpdesk weekly ticket category report vs. pre-

	volume	window	migration baseline

Table 2: Post-Migration Governance Checklist [15, 18, 19]

Conclusion

The migration of storage workloads from Azure Cloud Volumes ONTAP to on-premises NetApp arrays represents a financially compelling and technically tractable project for enterprise IT organizations operating high-capacity, stable workloads. The financial rationale is grounded in the structural difference between the CVO capacity-based licensing model and the amortized CapEx model of on-premises hardware: for sustained, predictable workloads, the on-premises TCO is substantially lower over a multi-year horizon, a gap that compounds with data volume. The CIOs who planned to repatriate workloads in 2024 are responding to precisely this financial reality, reinforced by data sovereignty obligations from the EU Data Act and the performance consistency advantages of dedicated on-premises storage documented in distributed file system research. The technical case rests on SnapMirror functioning as a native zero-additional-cost migration engine: block-level delta transfer, full Snapshot history preservation, a cutover window of fifteen to thirty minutes, and the reverse-resync capability that converts the vacated CVO instance into a warm DR standby without re-baselining. These properties are unique to same-ONTAP migrations and are not replicable by file-copy utilities or generic migration platforms. The foundational principles of distributed storage design and recent empirical work on migration optimization collectively support the conclusion that aligning storage architecture with workload characteristics and exploiting native replication capabilities rather than generic tooling is the defining principle of a successful, low-risk repatriation.

References

[1] Manini Mandal, "Cloud Volumes ONTAP licensing overview," NetApp Documentation, 2026. [Online]. Available: https://docs.netapp.com/us-en/bluexp-cloud-volumes-ontap/concept-licensing.html

[2] David A. Patterson et al., "A case for redundant arrays of inexpensive disks (RAID)," in Proc. ACM SIGMOD International Conference on Management of Data, 1988, pp. 109-116. [Online]. Available: https://dl.acm.org/doi/pdf/10.1145/50202.50214

[3] Sage Weil et al., "Ceph: A scalable, high-performance distributed file system," in Proc. 7th USENIX Conference on Operating Systems Design and Implementation (OSDI), 2006, pp. 307-320. [Online]. Available: https://www.usenix.org/legacy/event/osdi06/tech/full_papers/weil/weil.pdf

[4] Armando Fox et al., "Cluster-based scalable network services," in Proc. 16th ACM Symposium on Operating Systems Principles, 1997, pp. 78-91. [Online]. Available: https://dl.acm.org/doi/pdf/10.1145/268998.266662

[5] Howard Gobioff et al., "The Google file system," in Proc. 19th ACM Symposium on Operating Systems Principles (SOSP), 2003, pp. 29-43. [Online]. Available: https://static.googleusercontent.com/media/research.google.com/en//archive/gfs-sosp2003.pdf

[6] Konstantin Shvachko et al., "The Hadoop distributed file system," in Proc. IEEE 26th Symposium on Mass Storage Systems and Technologies (MSST), 2010, pp. 1-10. [Online]. Available: https://dl.acm.org/doi/10.1109/MSST.2010.5496972

[7] European Commission, "EU Data Act: Regulation (EU) 2023/2854," Official Journal of the European Union, 2023. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32023R2854

[8] Tony Ansley, "SnapMirror configuration and best practices guide for ONTAP 9," Technical

Report TR-4015, NetApp, 2024. [Online]. Available: https://www.netapp.com/media/17229-tr-4015-snapmirror-configuration-ontap.pdf
[9] Aaron Holt et al., "SnapMirror disaster recovery and data transfer concept," ONTAP 9.15, NetApp Documentation, 2025. [Online]. Available: https://docs.netapp.com/us-en/ontap/concepts/snapmirror-disaster-recovery-data-transfer-concept.html
[10] Lenida Vorwerk, "Resynchronize an ONTAP SnapMirror replication relationship," ONTAP 9.15, NetApp Documentation, 2025. [Online]. Available: https://docs.netapp.com/us-en/ontap/data-protection/resynchronize-relationship-task.html
[11] Yifei Zhang et al., "A novel cost-aware data placement strategy for edge-cloud collaborative smart systems," in Proc. IEEE 16th International Conference on Cloud Computing (CLOUD), 2023, pp. 450-456. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/10255009
[12] Mohammad A. Altahat et al., "Optimized encryption-integrated strategy for containers scheduling and secure migration in multi-cloud data centers," IEEE Access, vol. 12, pp. 51330-51345, 2024. [Online]. Available: https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=10494738
[13] Prasad Manda, "A comprehensive guide to migrating Oracle databases to the cloud: Ensuring minimal downtime, maximizing performance, and overcoming common challenges," International Journal of Research Publications in Engineering, Technology and Management, vol. 6, no. 3, pp. 8201-8209, 2023. [Online]. Available: https://ijrpetm.com/index.php/IJRPETM/article/download/97/96
[14] Alex Mathew, "Cloud data sovereignty: Governance and risk implications of cross-border cloud storage," ISACA, 2024. [Online]. Available: https://www.isaca.org/resources/news-and-trends/industry-news/2024/cloud-data-sovereignty-governance-and-risk-implications-of-cross-border-cloud-storage
[15] Yuanzhang Wang et al., "LoADM: Load-aware directory migration policy in distributed file systems," in Proc. Design, Automation and Test in Europe Conference and Exhibition (DATE), 2024, pp. 1-6. [Online]. Available: https://ieeexplore.ieee.org/document/10546634
[16] Microsoft, "Azure hybrid options," Microsoft Learn, 2026. [Online]. Available: https://learn.microsoft.com/en-us/azure/architecture/guide/technology-choices/hybrid-considerations
[17] Ignacio M. Llorente, "Cloud repatriation on the rise: 83% of CIOs plan workload shifts in 2024," EE Times Europe, 2024. [Online]. Available: https://www.eetimes.eu/cloud-repatriation-on-the-rise-83-of-cios-plan-workload-shifts-in-2024/
[18] Rajashree Goswami, "Recalibrating the cloud: Cost, control and the real ROI of repatriation," CTO Magazine, 2025. [Online]. Available: https://ctomagazine.com/roi-of-cloud-repatriation/
[19] Boris Zibitsker and Alexander Lupersolsky, "Cost optimization and performance control in the hybrid multi-cloud environment," in Proc. 16th ACM/SPEC International Conference on Performance Engineering, 2025, pp. 147-157. [Online]. Available: https://dl.acm.org/doi/pdf/10.1145/3676151.3722007
[20] Andrew Hibbard, "SnapMirror replication workflow," ONTAP 9.15, NetApp Documentation, 2025. [Online]. Available: https://docs.netapp.com/us-en/ontap/data-protection/snapmirror-replication-workflow-concept.html