

OpenSim S3 Asset Migration

Public Guide: What to Expect During the Live Migration

This document explains the technical reason for temporary asset slowdowns during the transition from the historical OSGrid asset storage to S3 object storage, what the grid prioritizes during the migration, and what users can expect while the system progressively becomes faster.

Key message: During the live migration, the absolute priority is to keep the grid serving the in-world experience. Background migration and administrative bulk operations are deliberately lower priority.

1. What Is Changing?

The grid is progressively moving its asset storage to S3 object storage. The historical asset service remains available as a fallback while the migration is underway. This is a live migration: the grid does not stop operating while the asset population is moved.

Assets are not simply copied as a single multi-terabyte filesystem operation. The system progressively discovers assets that are actually needed, checks whether they are already present in S3, and migrates missing assets from the historical service when required.

2. The Two Asset Paths

Normal path: an asset is requested by a region or viewer. The connector checks S3 first. If the asset is already there, it can be served directly from the new storage.

Migration path: if S3 does not yet contain the asset, the historical OSGrid service is used as a temporary fallback. The asset is retrieved and can then be stored in S3 so that future requests no longer need the old source.

Conceptually:

```
Viewer / Region -> Grid -> S3 -> asset
                                   \-> S3 miss -> historical OSGrid service -> asset -> S3 migration
```

3. In-World Service Has Absolute Priority

Throughout the live migration, the most important task is serving the active grid: region activity, viewer requests, asset retrieval, and normal in-world operation. Background migration is intentionally throttled so it does not consume all available CPU, database, network, disk, or S3 capacity.

This means the migration is intentionally **not** an unlimited-speed bulk transfer. The system would rather take longer to migrate several terabytes than create a situation where normal in-world activity is competing with a massive background copy.

Administrative operations such as large IAR exports or imports are therefore not treated as the primary workload during this phase. They can wait; the live grid cannot.

4. Why the Beginning of the Migration Can Be Slow

At the beginning, a large percentage of the assets requested by users may still exist only on the historical storage system. A request for such an asset can therefore require an additional network and storage step before the asset is available through S3.

A typical first-time request can look like this:

```
S3 lookup -> MISS -> historical asset server -> download -> S3 upload -> serve / continue processing
```

That path is necessarily slower than a direct S3 hit. The effect is expected to be strongest at the beginning of migration.

5. Performance Should Improve as Migration Progresses

Migration is cumulative. Once an asset has been transferred to S3, later requests can use the new storage directly. As more frequently used assets are migrated, the fraction of requests that need the historical fallback decreases.

The practical progression is therefore expected to be:

Early phase: more fallback accesses -> more variable load times
Middle phase: more S3 hits -> fewer fallback accesses -> improving consistency
Later phase: S3 becomes the normal path -> historical fallback becomes exceptional

The migration is therefore expected to start slowly and gain effective performance as the S3 population grows.

6. Why You May See "Slow handling" Messages

OpenSim can log a message such as Slow handling of GET /assets/... took 6148ms when an HTTP request takes longer than the configured diagnostic threshold.

During the migration, this can happen because a requested asset is not yet in S3 and must be fetched from the historical asset service. A slow-handling message is therefore a diagnostic timing indicator. It is not, by itself, proof of data loss, S3 failure, or a broken asset.

In particular, a temporary slowdown can occur even when the S3 system itself is healthy, because the request may be waiting on the legacy source.

7. How S3 and the Historical Service Are Used

Situation	Action
Asset already in S3	Serve from S3; no historical lookup is required.
Asset missing in S3, available historically	Read from historical service, then migrate to S3.
Asset missing in both	Treat as unavailable / missing; do not invent data.
Temporary S3 error	Use normal retry/error handling and keep fallback behaviour available.

8. Background Inventory Migration

The system also has a low-priority background process that reads asset references from the OpenSim inventory database. Its purpose is to move assets that are likely to be useful before a user requests them.

This process is deliberately limited by batch size, delay, and concurrency. It is not supposed to behave like an unrestricted bulk transfer consuming the whole server.

The background path is lower priority than live in-world traffic:

LIVE requests -> highest priority
Inventory pre-migration -> background / lower priority
Administrative bulk tasks -> should be scheduled with care during the migration period

9. Why We Ask Users Not to Repeatedly Load or Save Large IAR Files

During live migration, large IAR imports and exports are not the priority workload. They can create substantial additional asset activity on top of normal region and viewer traffic. Repeatedly saving and reloading the same large IAR archives can therefore amplify the amount of asset work the grid has to perform while the migration is already running.

Practical request to users: please avoid repeatedly exporting and re-importing IAR archives simply to test the storage system during the live migration. Unless there is a real operational need, one IAR backup/export approximately every three months is more than enough for normal backup practice.

This is not a restriction on legitimate maintenance. If an administrator genuinely needs an IAR for recovery, transfer, or another specific operational purpose, the operation can of course be performed. The request is simply to avoid continuous, repetitive load/save cycles while the infrastructure is performing a multi-terabyte live migration.

10. Why Repeated IAR Operations Can Affect Users

An IAR operation can reference a large number of assets. Depending on the operation, this may create additional asset reads, asset writes, database activity, compression/decompression work, network traffic, and S3 operations. When combined with live migration, this can increase resource contention.

For that reason, temporary slowness during this migration is not a good reason to repeatedly start another import or export. Doing so can increase the workload that is already being handled.

11. What Users Should Expect During the First Months

- Some asset loads may initially be slower than the eventual steady-state performance.
- Some region or viewer operations may occasionally wait for an asset to be recovered from the historical service.
- Server consoles may display diagnostic "Slow handling" messages for long-running asset requests.
- Performance should become more consistent as frequently used assets accumulate in S3.
- The migration may take a long time because the historical asset population is large and the transfer is intentionally throttled.
- The priority remains maintaining a usable in-world grid while this process runs.

12. What Administrators Monitor

Administrators monitor S3 and fallback statistics during the transition. Useful indicators include S3 latency, S3 errors, timeouts, throttling, server-side errors, fallback requests, fallback latency, and migration progress.

A healthy migration period can therefore show a large number of S3 misses or fallback requests while still showing zero actual errors. The important distinction is between a normal cache/migration miss and a real transport or service failure.

13. Long-Term End State

The target state is for S3 to become the normal asset store while the historical server becomes a secondary safety source. Once the required assets are present in S3, normal asset delivery requires fewer trips to the historical system and should be much less dependent on its disk and network performance.

The historical server can then remain available primarily as a read-only or emergency fallback source, depending on the final operational policy.

14. In One Diagram

LIVE USER TRAFFIC (highest priority)

Viewer / Region -> Grid -> S3 -> Asset

\-> MISS -> Historical OSGrid -> Asset -> S3

BACKGROUND INVENTORY MIGRATION (lower priority)

Inventory DB -> assetID -> S3 check -> Historical OSGrid if missing -> S3

ADMINISTRATIVE BULK OPERATIONS

Use when needed, but avoid repeated IAR load/save cycles during the live migration.

15. Final Message to Users

The migration is intentionally designed to protect the in-world experience first. The grid is not attempting to move several terabytes as fast as possible at the expense of live service. Instead, the system is moving assets progressively while keeping the historical service available as a fallback.

The beginning of the process may therefore feel slower. The expectation is that the effective performance will improve as more assets are available in S3 and fewer requests need the historical source.

Please allow the migration to do its work. Avoid continuously exporting and importing large IAR archives just to test the system. For normal backup purposes, an IAR export about once every three months is a reasonable operational guideline unless there is a specific reason to perform one more often.

The goal is simple: keep the grid usable today, migrate safely in the background, and progressively reach a faster and more sustainable S3-based asset service.

Technical note: exact timings depend on asset size, historical service performance, database activity, network conditions, S3 latency, and the concurrent workload of the grid. The migration is therefore intentionally measured and adjusted rather than promised as a fixed-duration transfer.