HPE Alletra Storage

Technical white paper



Superior economics and performance for Splunk SmartStore with HPE Alletra Storage MP X10000

Two-site resilience and reduced local hot and cache storage for Splunk SmartStore deployments

Contents

Executive summary 3
Introduction 3
Splunk SmartStore fundamentals 4
Solution architecture 5
Tested configuration 9
Deployment overview 11
Validation approach 14
Operating a two-site deployment 15
Conclusion 18
References 19
Appendix A: Validation scope and required caveats 20
Appendix B: Recommendations summary 20
Appendix C: Splunk configuration reference 21
Appendix D: X10000, endpoint, and network reference 22
Appendix E: Test method detail 23
Appendix F: Glossary 24

Executive summary

Splunk deployments keep growing. Ingest rises every year, and security and compliance requirements keep data for years. Splunk SmartStore keeps recent data in each indexer's local cache and the rest in S3 object storage, so storage can grow without adding indexers. That benefit depends on a fast object store; otherwise, the cache must be oversized to compensate.

HPE Alletra Storage MP X10000 has completed Splunk SmartStore validation for the submitted single-site and two-site on-premises multisite configuration. The validation ran a multisite Splunk indexer cluster across two sites, with an X10000 cluster at each site and every Splunk virtual machine on HPE Morpheus VM Essentials Software.

— **Migration.** Existing indexes moved to SmartStore with search results unchanged.

— **Search.** Searches returned the same results from cache and fully cold, with no failures; search time rose steadily as less data was cached.

— **Scale.** Doubling ingest to the 2 TB/day test profile completed with no throttling observed.

— **Two sites.** When an X10000 cluster or a whole site failed, ingest and searches continued, deletes made during the outage reconciled, and both clusters ended with matching object counts and stored bytes.

These results describe the tested configuration; they are not benchmarks or sizing guarantees. Replication between sites is asynchronous, so after an outage, cold searches on the recovering site can briefly return stale or incomplete results.

Introduction

Splunk indexes machine data, such as logs, metrics and events, for security operations and IT observability, often at very large ingest volumes. Classic Splunk deployments keep all indexed data on the indexers' own disks, so storage grows only by adding indexers, and the cluster's replication and search factors multiply every copy.

SmartStore separates the two. Each indexer keeps a local cache, and an S3-compatible object store holds the data. For resilience across data centers, the indexer cluster can span two sites, with an object store at each site replicating to the other; Splunk supports this on premises for two sites.

HPE validated this design with Splunk in 2026, using two X10000 clusters behind a single S3 endpoint. This paper describes the architecture, the tested configuration, deployment, how the validation was run, and operation.

Purpose and audience

Target audience: Presales consultants, solution architects, data center managers, enterprise architects, and deployment and implementation engineers.

Document purpose: This paper presents the Splunk-validated architecture for SmartStore on HPE Alletra Storage MP X10000 at one or two sites, with deployment steps and guidance for operating across two sites.

Validation scope

Splunk's validation decision is summarized in the following approved statement and scope note.

Validation statement. HPE Alletra Storage MP X10000 has completed Splunk SmartStore validation for the submitted single-site and two-site on-premises multisite configuration. The validated configuration demonstrated SmartStore deployment, migration, remote data lifecycle, cache eviction and retrieval, search completion, multisite failover, object replication, delete-marker reconciliation, and recovery across the submitted scenarios.

Scope. Validation applies only to the submitted product version, Splunk version, configuration, topology, and test scope. It is not a general S3 API certification, performance benchmark, production sizing recommendation, RPO guarantee, or approval of every HPE Alletra Storage MP X10000 configuration. Multisite replication is asynchronous, and cold searches on a recovering store can temporarily return stale or incomplete results until replication and delete-marker reconciliation complete.

Appendix A lists the caveats that accompany them.

Splunk SmartStore fundamentals

How SmartStore manages data

Indexers write incoming data to hot buckets on local storage. When a hot bucket rolls to warm, the indexer uploads it to the object store, which then holds the master copy; the local copy stays in the cache until the cache manager evicts it. A search reads hot buckets and cached warm buckets locally and first downloads any other warm buckets it needs into the cache. Splunk does not search the object store directly. Figure 1 shows both paths in a two-site deployment.

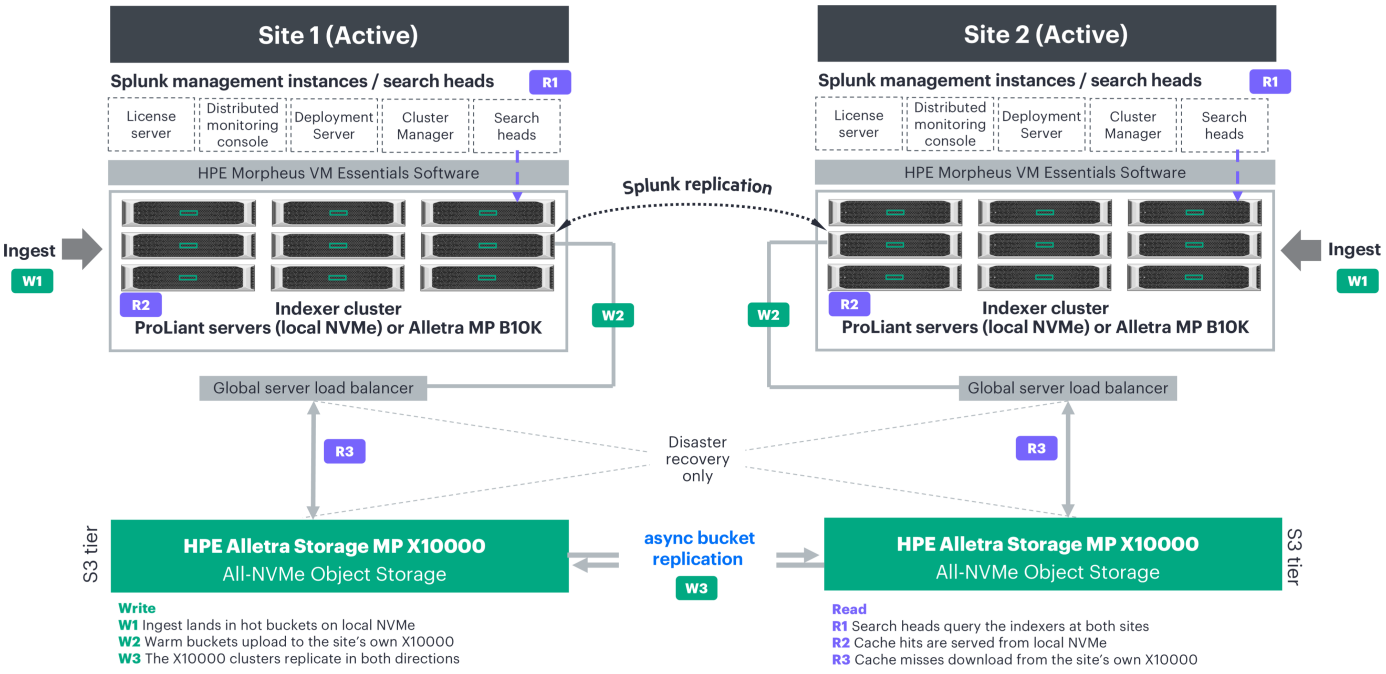


Figure 1. Write and read paths in a two-site SmartStore deployment

SmartStore also changes where copies live. The indexer cluster applies its replication and search factors to hot buckets only; the object store is responsible for the high availability and disaster recovery of warm buckets (SmartStore architecture overview). Across two sites, that protection comes from replication between the X10000 clusters, not from Splunk.

Splunk buckets and S3 buckets

The word bucket means two things here. A Splunk bucket is a directory of index files: hot while it is written, warm once it has rolled and uploaded, and frozen at the retention limit; SmartStore has no cold tier. An S3 bucket is the container on the X10000. All indexes share one S3 bucket, each under its own key prefix, and each Splunk bucket becomes several objects under that prefix.

Cache eviction removes only an indexer's local copy; the X10000 copy remains. Retention is different. When an index exceeds `maxGlobalDataSizeMB` or `maxGlobalRawDataSizeMB`, Splunk freezes its oldest buckets and deletes them from the X10000 and every cache. Unless an archive is configured, that data is gone. Whichever retention limit is reached first applies, so set size limits high enough not to cut time-based retention short. `maxTotalDataSizeMB` and `maxWarmDBCount` have no effect on SmartStore indexes (Configure data retention for SmartStore indexes).

The local cache

Hot buckets and the cache share each indexer's local storage. The cache holds recently used warm buckets; the cache manager evicts the least recently used ones when space runs low, and a later search that needs them downloads them again. Eviction never loses data, because the object store holds the master copy.

Figure 2 compares SmartStore on a capacity-oriented object store with SmartStore on the X10000. When cache misses are slow, the cache is sized for the worst case and local SSD capacity sets the indexer count. When they are fast, the cache can be sized to the working set, and compute sets the indexer count.

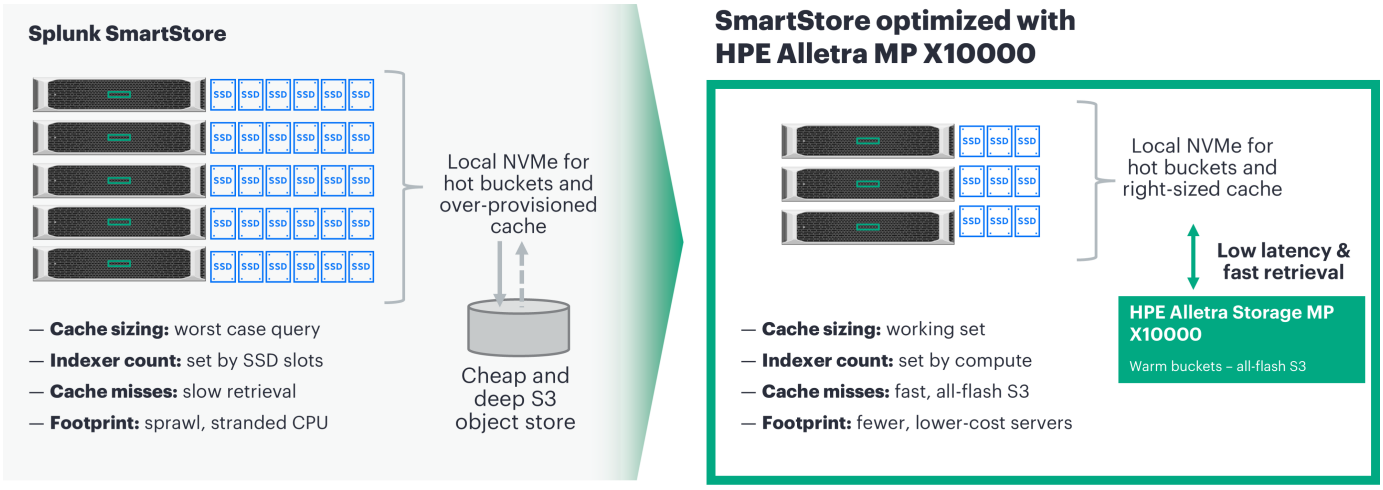


Figure 2. SmartStore with a capacity-oriented object store compared with HPE Alletra Storage MP X10000

This applies where local SSD capacity, not compute, sets the indexer count. If indexers are already CPU-bound at about 300 GB/day, a smaller cache saves SSDs but not servers.

Why customers deploy SmartStore

Splunk lists these advantages of SmartStore (About SmartStore):

— Storage and compute scale independently, so capacity grows without adding indexers.

— Long retention sits on object storage, with one master copy instead of replicated local copies.

— Recovery from indexer failure is fast, because warm data needs only metadata fixup.

— Warm data survives even if as many indexers as the replication factor fail.

Solution architecture

The solution runs a multisite Splunk indexer cluster across two sites, with an HPE Alletra Storage MP X10000 cluster at each site as the SmartStore remote store.

Solution components

— **Splunk forwarders** send data to the indexers, load-balanced across both sites.

— **Splunk indexers** index and search data. Each holds hot buckets and a local cache and uploads warm buckets to the object store over S3.

— **Splunk search heads** run searches across the indexers at both sites.

— **Splunk cluster managers** coordinate the indexer cluster and distribute the SmartStore configuration to the indexers. There is one at each site.

— **Splunk monitoring console** monitors the deployment, including SmartStore activity.

— **HPE Alletra Storage MP X10000** is the S3 object store for SmartStore, one cluster at each site.

— **Loadbalancer.org GSLB** presents one S3 endpoint name to every indexer and routes each indexer to its own site's X10000 cluster.

— **HPE GreenLake** manages X10000 buckets, S3 users and replication policies.

— **HPE Morpheus VM Essentials Software** hosts the Splunk virtual machines.

HPE Morpheus VM Essentials Software

Every Splunk virtual machine in the validation (indexers, search heads, cluster managers, monitoring console and forwarders) ran on HPE Morpheus VM Essentials Software across three hosts. Figure 3 shows them in the Instances view, where the status, health and resource use of each one can be checked from a single console.

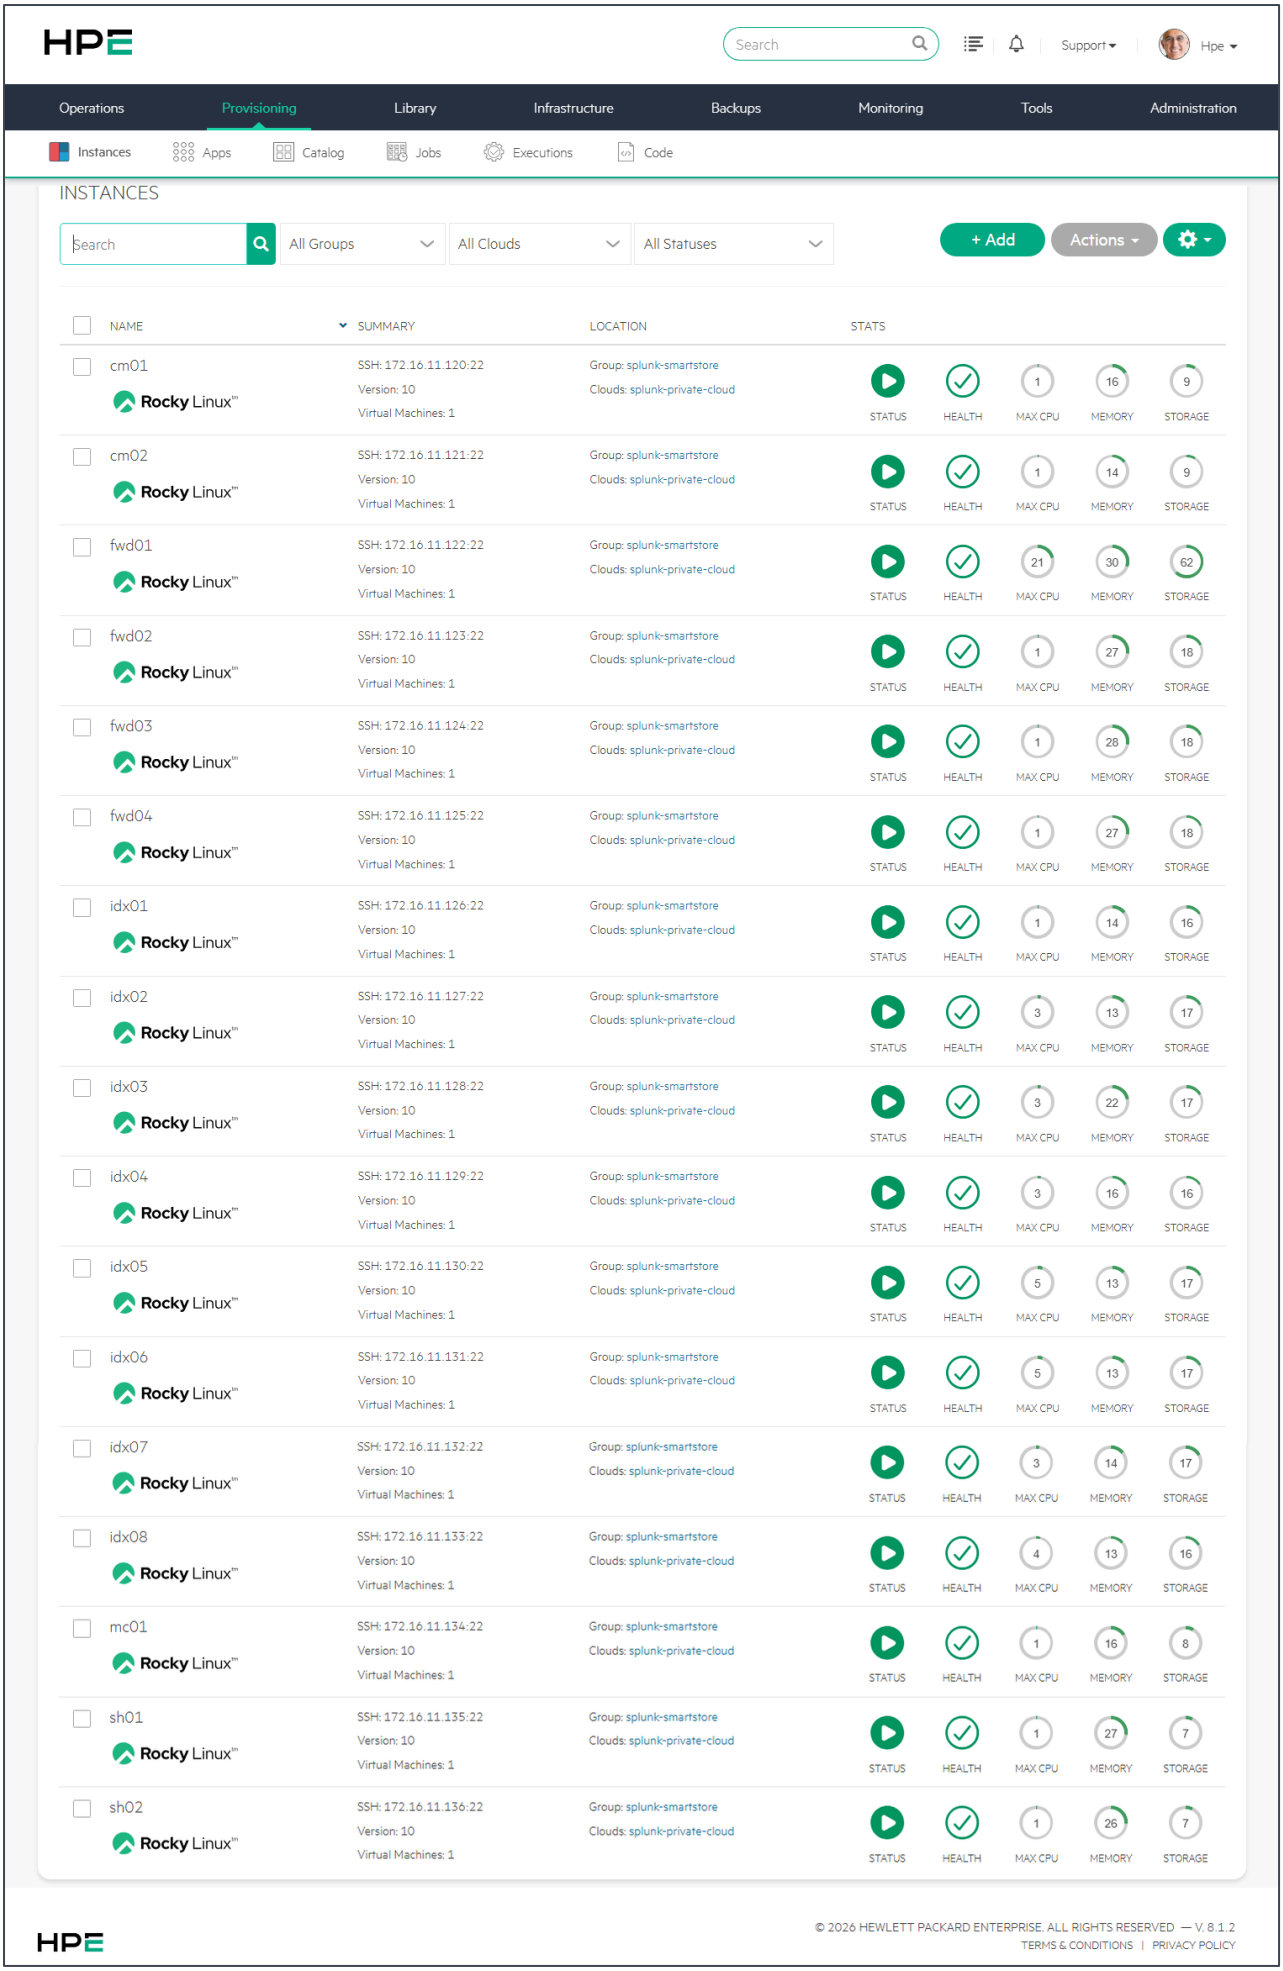


Figure 3. Splunk virtual machines in the HPE Morpheus VM Essentials Software Instances view

Reference architecture

Figure 4 shows the validated architecture. Each site hosts indexers, a cluster manager and an X10000 cluster. Search heads run with site affinity disabled (`site=site0`), so they search the indexers at both sites. The cluster's site replication and search factors keep at least one searchable copy of every bucket at each site.

All indexes share one remote volume: one endpoint and one S3 bucket, with each index stored under its own prefix. A single-site deployment is the same design without the second site, the replication and the GSLB.

Splunk limits on-premises SmartStore multisite to two sites, each in its own data center (Deploy multisite indexer clusters with SmartStore). Designs with more sites apply only to public cloud deployments.

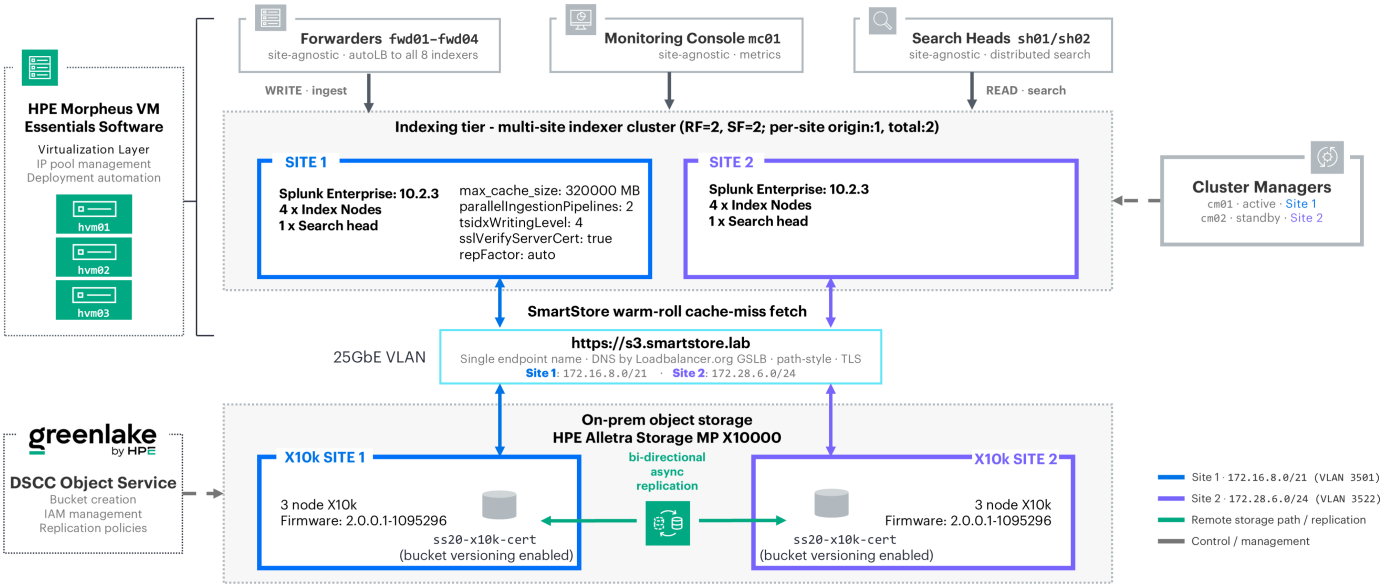


Figure 4. Two-site Splunk SmartStore reference architecture with HPE Alletra Storage MP X10000

A single S3 endpoint for both sites

Splunk requires every indexer to use the same remote store endpoint. Here, that endpoint is `https://s3.smartstore.lab`, and the GSLB answers DNS queries for it with:

— **Site-local routing.** Each indexer receives a data IP on its own site's X10000 cluster.

— **Health-checked failover.** When a cluster stops answering on TCP 443, the GSLB answers with the surviving cluster. No indexer configuration changes.

— **A short record lifetime,** so indexers pick up a failover quickly.

Failover works only if the other cluster already holds the data. Before relying on it, confirm that both clusters report the same object count, especially right after enabling replication on a bucket that already holds data.

The GSLB only answers DNS. S3 traffic flows directly between the indexers and the X10000 data IPs. The validation used a Loadbalancer.org GSLB virtual appliance; any GSLB with these capabilities can be substituted. Because the record lifetime is short, indexers lose the endpoint name within seconds if no GSLB answers, so deploy the GSLB redundantly, with an instance at each site.

Replication between X10000 clusters

Each X10000 cluster holds a bucket with the same name (`ss20-x10k-cert`), and HPE GreenLake replication policies are configured. This meets Splunk's object store requirements for on-premises multisite deployments:

— Bi-directional replication, with one policy in each direction, for the bucket.

— Object versioning, enabled when the bucket is created.

— Delete-marker replication, so a deletion at one site reaches the other.

— Replication that resumes after a failure. After a cluster disconnect in the validation, the backlog drained to zero with no operator action.

SmartStore keeps one copy of each bucket in the object store, whatever the replication and search factors, and replication gives each site's X10000 cluster a full copy of the retained data. The data arrives compressed, so plan X10000 capacity with no further data reduction.

Replication is asynchronous. While a failed site recovers, cold searches served from its X10000 cluster can briefly return stale or incomplete results until replication catches up. See Operating a two-site deployment.

Cluster manager availability

Splunk requires a cluster manager at each site: one active, the other standby. If the active manager's site fails, the standby is activated and the surviving indexers are pointed to it, as described in Handle manager site failure. This is separate from the S3 endpoint failover.

Network, DNS delegation, and TLS trust

— **Network.** SmartStore adds one requirement: HTTPS (TCP 443) from every indexer to the S3 endpoint. Splunk recommends no more than 100 ms of latency between indexer cluster nodes and allows up to 300 ms between sites for on-premises multisite SmartStore; higher latency sharply slows recovery from node failures (Reference hardware).

— **DNS.** The GSLB chooses the site from the address a query comes from, so each indexer's queries for the endpoint name must reach the GSLB from an address in its own site's topology. In the validation, each indexer's first DNS server was the GSLB's address on its own site's data network.

— **TLS.** Indexers verify the endpoint certificate (`remote.s3.sslVerifyServerCert = true`) and use path-style addressing. Because either cluster can answer for the endpoint name, both clusters' certificates must list it as a subject alternative name, and Splunk must trust the signing certificate authority.

Tested configuration

Tables 1 to 4 describe the validated environment. Values are the configuration of record, not sizing or tuning recommendations.

Hardware

Table 1. Hardware

Component	Configuration
Object storage	Two HPE Alletra Storage MP X10000 clusters, one per site. Each has 3 nodes and 24 × 3.84 TB NVMe drives (84 TiB raw, 62 TiB usable).
Network	25 GbE data network per site: site 1 `172.16.8.0/21`, site 2 `172.28.6.0/24`. 24 S3 data IPs per X10000 cluster (8 per node).
Virtualization hosts	Three HPE ProLiant DL380 servers, 768 GB memory each

Splunk deployment

Table 2. Splunk deployment

Role	Hosts	Configuration
Indexers	`idx01–idx04` (site 1), `idx05–idx08` (site 2)	24 vCPU, 32 GiB memory, 1.8 TiB SSD, of which about 1.5 TB holds hot buckets and the SmartStore cache
Search heads	`sh01`, `sh02`	16 vCPU, 16 GB memory; `site=site0` (site affinity disabled)
Cluster managers	`cm01` (site 1), `cm02` (site 2)	12 vCPU, 16 GB memory; `cm01` active, `cm02` standby
Monitoring console	`mc01`	12 vCPU, 16 GB memory
Forwarders	`fwd01–fwd04`	16 vCPU, 8 GB memory; load-balanced across all eight indexers

Software and firmware versions

Table 3. Software and firmware

Component	Version
Splunk Enterprise	10.2.3
Indexer operating system	Rocky Linux 10.1
HPE Alletra Storage MP X10000 firmware	2.0.0.1-1095296
HPE Morpheus VM Essentials Software (HVM)	8.1.2
Loadbalancer.org Enterprise ADC, KVM virtual appliance	LBOS-7 release 8.13.4

Configuration of record

Table 4 summarizes the validated settings. Appendix C lists the full effective configuration.

Table 4. Configuration of record (summary)

Setting	Value
S3 endpoint	`https://s3.smartstore.lab` (TLS, port 443)
S3 bucket	`ss20-x10k-cert` on both clusters; versioning enabled at creation; bi-directional asynchronous replication, including delete markers
GSLB pool	Topology-weighted round robin (own site's cluster while it has a healthy data IP, otherwise the other site's); one address per response; TCP 443 health check every 10 s (5 s timeout, 2 retries); record TTL 5 s
Indexer cluster	`replication_factor` and `search_factor` 2; `site_replication_factor` and `site_search_factor` `origin:1,total:2`
`parallelIngestionPipelines`	2
`tsidxWritingLevel`	4
`max_cache_size`	320000 (MB)
`maxDataSize`, `repFactor`	`auto`, `auto`
`remote.s3.url_version`	`v1` (path-style)
`remote.s3.sslVerifyServerCert`	`true`

Deployment overview

Deployment runs in three stages: the X10000 clusters, then the GSLB and certificates, then Splunk. Appendices C and D have the full configuration listings.

Prerequisites

— An X10000 cluster at each site, managed in HPE GreenLake.

— A Splunk multisite indexer cluster with a cluster manager at each site.

— A GSLB, and DNS that sends the indexers' queries for the endpoint name to it.

— A certificate authority to sign the X10000 endpoint certificates.

— HTTPS (TCP 443) from every indexer to the S3 data IPs of both X10000 clusters.

Prepare the X10000

In HPE GreenLake:

— **Create the bucket** (`ss20-x10k-cert`) on each cluster with versioning enabled. Versioning cannot be turned on later through the S3 API; a bucket created without it must be replaced.

— **Create one S3 user on both clusters** with Assign User to Multiple Systems, so both clusters share the single access key and secret the Splunk remote volume holds. The secret is shown only once.

— To rotate the credential or add an X10000 cluster later, create a new user the same way (the name must be new), then update `remote.s3.access_key` and `remote.s3.secret_key` on the cluster manager and push the bundle. Reset Password works on one cluster at a time and gives each cluster a different secret, so do not use it.

Figure 5 shows where Assign User to Multiple Systems is in HPE GreenLake, and Figure 6 shows the dialog, where every X10000 cluster is selected and the one-time secret is shown.

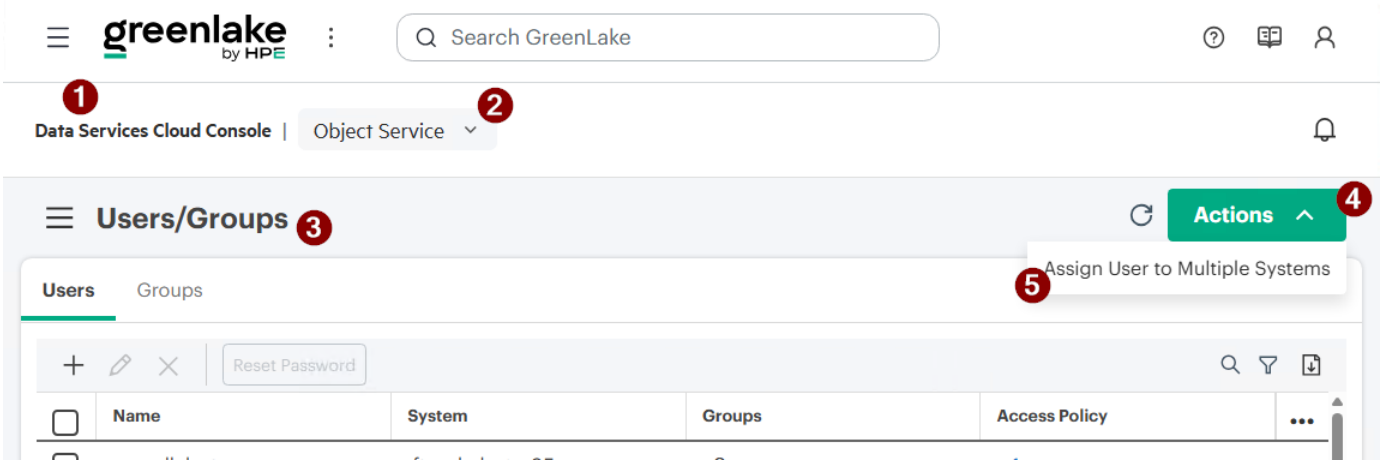


Figure 5. Assign User to Multiple Systems in the HPE GreenLake Object Service

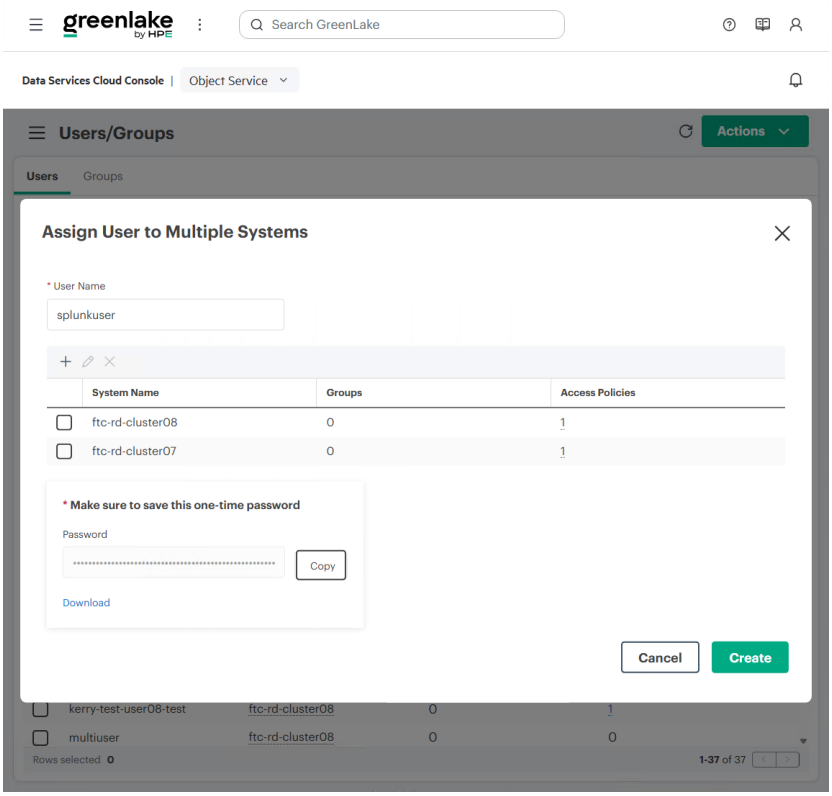


Figure 6. Creating one S3 user on multiple X10000 clusters

Configure bi-directional replication

Add a remote protection policy to each bucket, one in each direction:

— Partner: the other X10000 cluster. Target bucket: the same bucket name.

— Replicate delete markers and replicate delete operations: enabled.

— Locking: disabled.

Configure the S3 endpoint

On the GSLB, create four objects (Appendix D has the `lbcli` commands):

— **Global name** `s3.smartstore.lab`, with a short TTL.

— **Members**: every S3 data IP of both clusters, each health-checked on TCP 443.

— **Pool**: topology-weighted round robin, which answers with the indexer's own site while that site has a healthy member and with the other site otherwise; one address per response.

— **Topologies**: map each site's indexers to its local cluster's members.

Then make the GSLB answer the indexers' queries for `s3.smartstore.lab` (Appendix D, DNS delegation) and have each cluster's certificate signed with `s3.smartstore.lab` as a subject alternative name.

Configure Splunk

On the active cluster manager, define one remote volume in `manager-apps/_cluster/local/indexes.conf` and point each index at it. Credentials and the remaining settings are in Appendix C; see also Configure the S3 remote store for SmartStore.

```
[volume:remote_store]
storageType = remote
path = s3://ss20-x10k-cert
remote.s3.endpoint = https://s3.smartstore.lab
remote.s3.url_version = v1
remote.s3.supports_versioning = true
remote.s3.sslVerifyServerCert = true
remote.s3.sslRootCAPath = /opt/splunk/etc/auth/mycerts/ca.crt

[<index_name>]
remotePath = volume:remote_store/$_index_name
repFactor = auto
```

Place the CA chain at `sslRootCAPath` on every indexer, then validate and push the bundle:

```
splunk validate cluster-bundle --check-restart
splunk apply cluster-bundle
splunk show cluster-bundle-status
```

Some settings take effect only after a restart. Confirm a setting is in effect on an indexer with `splunk btool`, not only that it is in the bundle.

Migrate existing indexes to SmartStore

A new index can use SmartStore from creation (Deploy SmartStore on a new indexer cluster). An existing index is converted with Migrate existing data on an indexer cluster to SmartStore: the cluster goes into maintenance mode, all peers stop, the configuration is pushed, and all peers start and upload their data. Migration cannot be reversed.

Verify the deployment

From an indexer, confirm that:

— The endpoint presents the expected certificate and issuer:

```
openssl s_client -connect s3.smartstore.lab:443 \
  -servername s3.smartstore.lab </dev/null \
  | openssl x509 -noout -subject -issuer
```

— The bucket is reachable with the SmartStore credential, and versioning is enabled:

```
aws --endpoint-url https://s3.smartstore.lab s3 ls s3://ss20-x10k-cert/
aws --endpoint-url https://s3.smartstore.lab \
  s3api get-bucket-versioning --bucket ss20-x10k-cert
```

— Splunk can list the remote volume:

```
splunk cmd splunkd rfs -- ls --starts-with volume:remote_store
```

The Splunk monitoring console should then show remote storage as online for every indexer. If it does not, Troubleshoot SmartStore lists the dashboards, logs, and commands to check.

Validation approach

HPE ran Splunk's partner test plan for SmartStore against the configuration in Tested configuration, and Splunk reviewed the evidence for each phase.

Validation phases

The plan has nine phases: environment (0), S3 API compatibility (1), deployment (2), functional lifecycle (3), migration (4), remote-store performance (5), search performance (6), scale (7), and multisite failure and recovery (8). Figure 7 shows phases 1 to 8.

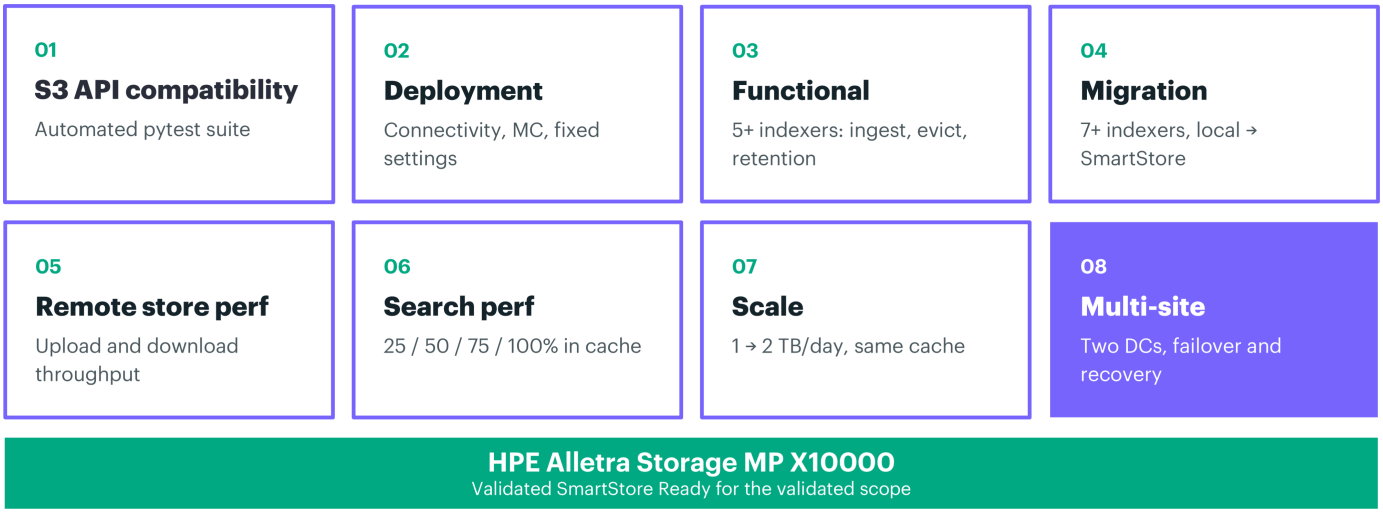


Figure 7. Splunk SmartStore validation phases

All phases ran on the multisite indexer cluster. Phases 2 to 7 used one X10000 cluster; Phase 8 added the second cluster and replication. Splunk's final decision passed every phase, seven of them with caveats; Appendix A lists the caveats.

Workload

Forwarders generated synthetic web-access events with the open-source gogen generator and sent them to all eight indexers. The baseline was at least 10 days of data at 1 TB/day or more, as the plan requires; the scale phase doubled ingest to 2 TB/day. For the migration corpus, ingest volume was confirmed three ways: the Splunk license meter, bucket sizes reported by Splunk, and the event-time distribution.

How results were verified

Correctness was checked with three fixed searches, each over a fixed time window with a known event count: two selective searches over raw events, and one answered from a report-acceleration summary. Each search ran from both search heads, first with its data in the indexer cache and again after every indexer's cache was evicted, so that the data came from the X10000. A run passed when every count matched its baseline.

Phase 8 took down a full site, and separately one X10000 cluster, while ingest and searches continued, then recovered each. A known set of events was deleted while one cluster was offline, to confirm the deletion reached both clusters after recovery. Finally the two clusters' object counts and total stored bytes were compared.

Performance and replication measurements came from the Splunk monitoring console, the Splunk job inspector and HPE GreenLake.

Operating a two-site deployment

When an X10000 cluster fails

If one X10000 cluster fails but its site stays up, the GSLB health checks mark that cluster's data IPs down and answer the endpoint name with the surviving cluster. That site's indexers then upload to and read from the surviving cluster across the intersite link. No Splunk action is needed. In the validation, the indexer cluster stayed fully healthy, and ingest and searches continued.

Replication into the failed cluster stops, and changes, including deletes, wait on the surviving cluster until it returns. Cold searches from the affected site cross the intersite link and can be slower until the cluster is restored.

To see which cluster an indexer is using, resolve the endpoint from that indexer and check the certificate subject, which names the cluster that answered:

```
dig +short s3.smartstore.lab
openssl s_client -connect s3.smartstore.lab:443 \
  -servername s3.smartstore.lab </dev/null \
  | openssl x509 -noout -subject
```

When a site fails

When a whole site fails, its indexers and X10000 cluster fail together, and possibly the active cluster manager. The surviving site keeps ingesting and searching from its own cluster. The cluster manager reports all data searchable while the site replication and search factors are not met; that is expected until the site returns.

If the active cluster manager was at the failed site, promote the standby on the surviving site (Handle manager site failure). The validation used manual switchover, as Splunk's test plan required. On the standby manager, run:

```
splunk cluster-manager-redundancy -switch-mode active
```

Peers and search heads must reach the newly active manager; Splunk describes a DNS record or load balancer for this (Implement cluster manager redundancy). If the manager restarts while a site is still down, it blocks indexing until you run `splunk set indexing-ready` (Restart indexing in multisite cluster after manager restart or site failure).

When a failed site rejoins

This applies both to a returning site and to a returning X10000 cluster. When a site returns, its peers rejoin the indexer cluster and cluster fixup runs. The returning cluster manager rejoins as a standby, so there is no need to switch back. Make sure the returning site's peers point to the active manager.

Replication then drains the backlog into the returning X10000 cluster. Replication is asynchronous, so until it finishes, cold searches on the returning site can return stale or incomplete results: data deleted during the outage can briefly reappear, and some bucket downloads can fail. In the validation both cleared on their own; deleted events reached zero on both search heads 10 minutes 21 seconds after the cluster returned. That is an observed result, not a recovery-time commitment. Searches served by the surviving cluster, and all steady-state searches, returned correct and complete results throughout.

Figure 8 shows the replication backlog on the surviving cluster during a cluster outage in the validation: it built while the other cluster was down, held steady, and drained to zero after the cluster returned.

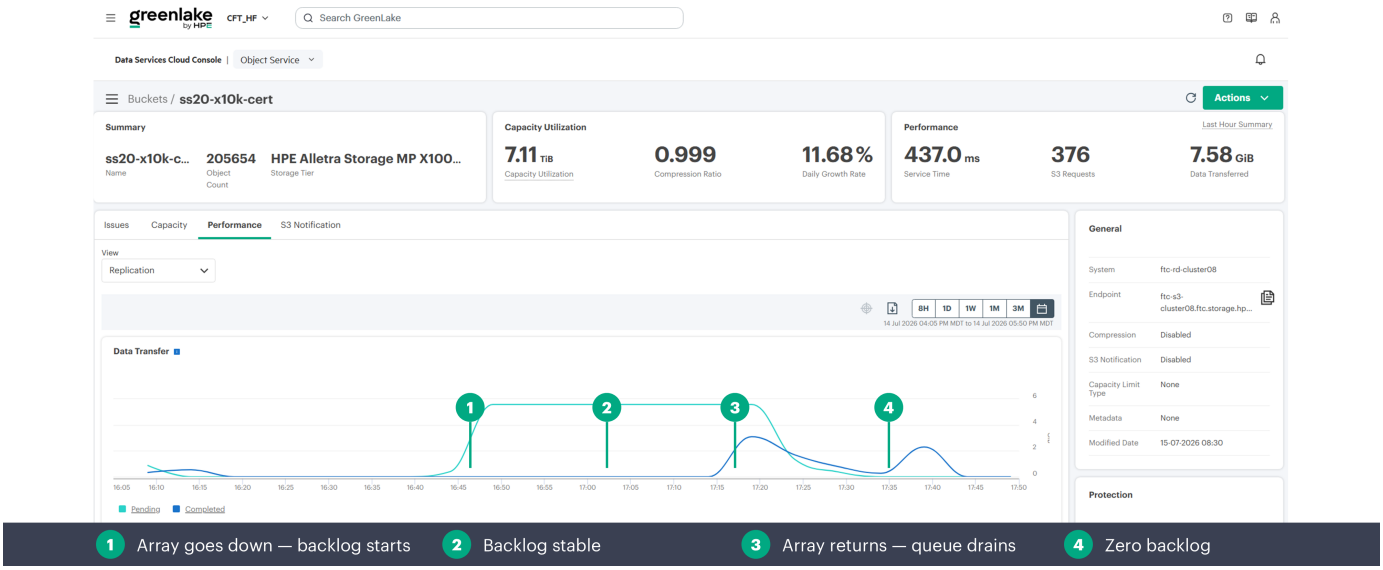


Figure 8. Replication backlog on the surviving X10000 cluster during a cluster outage and recovery, in HPE GreenLake

Treat the returning site as ready only when the checks in Confirming both sites are ready for cold searches pass.

What to monitor

Table 5 lists what to watch in normal operation and during recovery. Troubleshoot SmartStore describes the monitoring console dashboards and the related log files.

Table 5. What to monitor in a two-site deployment

Area	Where	Look for
Replication between clusters	HPE GreenLake: the bucket's Performance tab, Replication view, on each X10000 cluster	Partner status normal; pending objects and bytes near zero in both directions; after an outage, the backlog draining to zero
Remote storage activity	Monitoring console, SmartStore Activity dashboards	Remote storage online; upload and download failures
Cache behavior	Monitoring console, SmartStore Cache Performance dashboards	Eviction rate, cache misses, repeat downloads
Indexer cluster	`splunk show cluster-status --verbose` on the active manager	All peers up, all data searchable, site replication and search factors met, no fixup tasks
Cluster managers	`splunk cluster-manager-redundancy -show-status`	One manager active, the other standby
S3 endpoint	GSLB health status	Every data IP of both clusters up; each site's indexers resolving to their own cluster

Figure 9 shows the replication view in steady state: versioning enabled, the replication policy to the partner cluster, delete markers and delete operations replicated, and latency flat across operations.

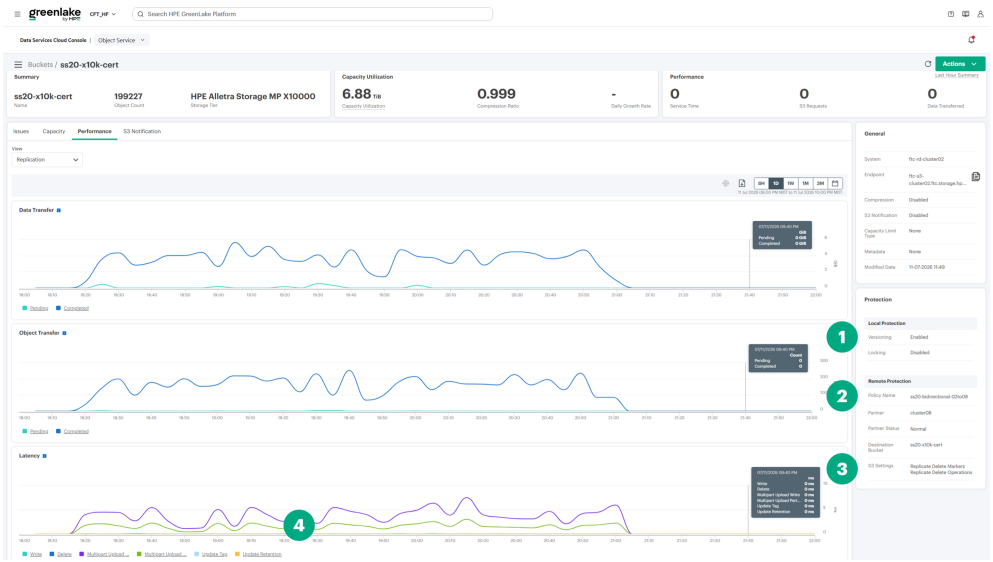


Figure 9. Steady-state replication view of the bucket in HPE GreenLake

No recovery point objective (RPO) was defined or validated. Production deployments need a customer-defined RPO, and monitoring that replication lag stays within it.

Confirming both sites are ready for cold searches

After an outage, both sites are ready for cold searches when:

— The replication backlog is zero in both directions in HPE GreenLake.

— Both X10000 clusters report the same object count and total stored bytes. Query each cluster by its own hostname, which must be on the endpoint certificate, rather than through the shared endpoint. This compares totals, not individual objects.

```
aws --endpoint-url https://cluster08.smartstore.lab \
  s3 ls s3://ss20-x10k-cert/ --recursive --summarize | tail -2
```

— The cluster status shows all peers up, all data searchable, both site factors met and no fixup tasks.

Troubleshooting quick reference

Table 6 lists common symptoms, starting from the X10000 side. For Splunk-side symptoms, see Troubleshoot SmartStore; for cache-full conditions, see SmartStore CacheManager failures.

Table 6. Troubleshooting quick reference

Symptom	First check	Likely cause
TLS verification errors	Certificate subject and issuer with `openssl s_client`	Wrong certificate authority in `remote.s3.sslRootCAPath`, or the endpoint name is not on the certificate
`403 AccessDenied`	`aws s3 ls` with the SmartStore credential	Credential not assigned to this X10000 cluster
`NoSuchBucket`	Bucket name in the remote volume path	Wrong bucket name
Versioning off	`get-bucket-versioning`	Not enabled when the bucket was created; a new bucket is needed
Replication stalled	Policy status in HPE GreenLake; write a test object and look for it on the other cluster	Policy missing, one-way, or without delete-marker replication

Symptom	First check	Likely cause
Indexers reach the wrong cluster, or reads miss after failover	`dig +short` and the certificate subject, from an indexer	GSLB site mapping, or replication not yet caught up
No failover when an X10000 cluster fails	GSLB member status and weights	Health check still passing, or the other cluster's members are missing from the pool or set to weight 0
Indexing blocked after the manager restarts while a site is down	Indexing Ready in `splunk show cluster-status`	Expected; run `splunk set indexing-ready` on the manager
"Failed to localize fileSet... Results will be incomplete"	`S3Client` and `CacheManager` in `splunkd.log` on that indexer	A bucket download failed; brief failures are expected while a site recovers
Cache partition full, or "Unable to evict enough data"	SmartStore Cache Performance dashboard; `max_cache_size`	`max_cache_size` too large for the partition, or broad searches downloading many buckets
Slow cold searches	Cache size and network bandwidth	Cache too small for the working set, or a slow link; cache misses are always slower than cache hits

Conclusion

HPE Alletra Storage MP X10000 has completed Splunk SmartStore validation for the submitted single-site and two-site on-premises multisite configuration. Across eight test phases, two X10000 clusters (one at each site) served as the SmartStore remote store for a multisite Splunk indexer cluster. The tests covered deployment, migration of existing indexes, the full data lifecycle, cold searches, and failures of a site and of a cluster.

The design is simple. Every indexer uses one S3 endpoint name, which a GSLB resolves to the site's own X10000 cluster without sitting in the data path. Each cluster holds a versioned bucket of the same name, and replication runs in both directions, delete markers included. When an X10000 cluster became unavailable, its site's indexers moved to the surviving cluster with no change on any indexer, and ingest and cold searches continued. Replication is asynchronous, so after a cluster returns, cold searches there can briefly return stale or incomplete results until replication and delete markers catch up. Operating a two-site deployment describes how to confirm both sites are ready.

Cold searches were slower than cached ones, as SmartStore's design expects, and no throttling was observed at any tested request rate. These values describe the tested environment; they are not benchmarks or sizing guarantees.

Together, HPE Alletra Storage MP X10000 and HPE Morpheus VM Essentials Software demonstrate the value of an integrated HPE platform for Splunk environments. The X10000 provides scalable S3-compatible object storage with resilient two-site data services. HPE Morpheus VM Essentials Software ran every Splunk virtual machine in this validation, and managing them all from one console made the deployment easier to build and operate. Combined in one architecture, they give enterprise-scale Splunk SmartStore deployments a consistent, end-to-end HPE foundation.

References

HPE

Scalable and resilient Splunk SmartStore multi-site storage with HPE Alletra Storage MP X10000 (solution brief)

HPE Alletra Storage MP X10000

Top 10 reasons to choose HPE Alletra Storage MP X10000

HPE Morpheus VM Essentials Software

Splunk

Splunk S3 compatibility and certification repository

SmartStore architecture overview

Deploy multisite indexer clusters with SmartStore

Configure the SmartStore cache manager

GSLB

The ultimate guide to Global Server Load Balancing (GSLB): Benefits, uses, and configurations (Loadbalancer.org)

Learn more at

HPE.com/storage

Appendix A: Validation scope and required caveats

Splunk's validation decision, dated September 29, 2026, includes the following caveats, which must remain attached to any detailed description of the validated scope:

1. Validation is limited to the submitted SmartStore configuration and does not establish full AWS S3 API parity.
2. The validated workflow requires a pre-created bucket with versioning enabled administratively. Programmatic PutBucketVersioning is not supported.
3. Normal multipart completion is validated. Repeating completion against the same consumed upload ID returns NoSuchUpload.
4. Multisite replication is asynchronous. During failback and resynchronization, cold searches on the recovering side can temporarily return stale or incomplete results, including data whose delete marker has not yet reconciled.
5. The observed delete-reconciliation time of approximately 10 minutes 21 seconds is a submitted test result, not a recovery-time commitment or service-level objective.
6. No numeric RPO was defined or validated. Production deployments require a customer-defined RPO and HPE assurance and monitoring that replication lag remains within it.
7. The reported 76.6 ms value is peak observed replication-operation latency, not end-to-end replication lag and not an RPO result.
8. Matching final object count and aggregate stored bytes were demonstrated. No per-object checksum or object-manifest comparison established byte-for-byte equality.
9. Performance and scale values describe the submitted test environment. They are not universal benchmarks, sizing guarantees, or endorsements of broader performance claims.

Appendix B: Recommendations summary

This appendix collects the recommendations made in this paper. Each is backed by a validation result or by Splunk documentation, and points to the section that explains it. These are HPE recommendations; they are not part of Splunk's validation decision, which covers only the tested configuration (Appendix A).

Table 7. Recommendations and their evidence

Recommendation	Evidence	See
Create the X10000 bucket with versioning enabled	Splunk requires object versioning for on-premises multisite (Deploy multisite indexer clusters with SmartStore). Versioning cannot be enabled later through the S3 API (Appendix A).	Prepare the X10000
Replicate in both directions, including delete markers	Splunk requires bi-directional and delete-marker replication (Deploy multisite indexer clusters with SmartStore).	Replication between X10000 clusters
Give every indexer one endpoint name, answered by a health-checked GSLB	In the validation, when one X10000 cluster failed, the GSLB marked its data IPs down and that site's indexers continued on the surviving cluster with no change on any indexer.	A single S3 endpoint for both sites
Keep TLS certificate verification on	The validation ran with `remote.s3.sslVerifyServerCert = true` on every indexer.	Network, DNS delegation, and TLS trust
Keep intersite latency within Splunk's limits	Splunk recommends no more than 100 ms between cluster nodes and allows up to 300 ms between sites for on-premises multisite SmartStore (Reference hardware).	Network, DNS delegation, and TLS trust

Recommendation	Evidence	See
Run a standby cluster manager at the other site	Splunk's on-premises multisite topology has an active manager at one site and a standby at the other (Handle manager site failure). Splunk's validation test plan calls for promoting the standby manually after a site failure.	When a site fails
Give the cache its own partition and keep it below 95% full	Splunk treats the cache as the whole partition it sits on (SmartStore system requirements). Default settings can fill a terabyte-scale cache to about 99%; Splunk support recommends staying below 95% (SmartStore CacheManager failures).	Cache configuration in server.conf
Set global size limits above the expected data volume	When an index exceeds `maxGlobalDataSizeMB` or `maxGlobalRawDataSizeMB`, Splunk freezes its oldest buckets and deletes them from remote storage (Configure data retention for SmartStore indexes).	Splunk buckets and S3 buckets
Confirm a returning site is ready before relying on its cold searches	Splunk notes that cold searches can return incomplete data during replication lag (Deploy multisite indexer clusters with SmartStore).	Confirming both sites are ready for cold searches
Define a recovery point objective (RPO) and monitor the replication backlog	Splunk requires the object store vendor to assure that replication lag stays within the customer's RPO (Deploy multisite indexer clusters with SmartStore). No RPO was defined or validated (Appendix A).	What to monitor

Appendix C: Splunk configuration reference

Effective configuration

Confirm settings on an indexer with `splunk btool <conf> list --debug` rather than only in the cluster manager's bundle; the validation confirmed the validated values this way.

Example indexes.conf

Configure Splunk shows the remote volume and the per-index stanza. The volume also holds the credential (`remote.s3.access_key` and `remote.s3.secret_key`), and each index keeps `homePath`, `coldPath` and `thawedPath`, which SmartStore indexes ignore except during migration.

The validated `indexes.conf`, as `btool` reported it on an indexer, with the credentials replaced by placeholders:

```
[volume:remote_store]
storageType = remote
path = s3://ss20-x10k-cert
remote.s3.access_key = <access key>
remote.s3.secret_key = <secret key>
remote.s3.auth_region = us-east-1
remote.s3.endpoint = https://s3.smartstore.lab
remote.s3.sslRootCAPath = /opt/splunk/etc/auth/mycerts/ca.crt
remote.s3.sslVerifyServerCert = true
remote.s3.supports_versioning = true
remote.s3.url_version = v1

[default]
tsidxWritingLevel = 4

[ss-high-volume]
homePath = /opt/splunk-local/ss-high-volume/db
coldPath = /opt/splunk-local/ss-high-volume/colddb
thawedPath = /opt/splunk-local/ss-high-volume/thaweddb
```

```
maxDataSize = auto
remotePath = volume:remote_store/$_index_name
repFactor = auto
```

Cache configuration in server.conf

Set `max_cache_size` in the `[cachemanager]` stanza of server.conf on every indexer (320000 MB in the validation). The other cache manager settings stayed at Splunk defaults, including `eviction_padding` (5,120 MB) and `minFreeSpace` (5,000 MB, in `[diskUsage]`).

Splunk treats the cache as the whole partition it sits on, so `max_cache_size` counts everything there, including the operating system, Splunk binaries and search artifacts if they share it (SmartStore system requirements). Give the cache its own partition where possible. If `max_cache_size` is set close to the partition size, a terabyte-scale cache can fill its partition to about 99%; Splunk support recommends keeping usage below 95% (SmartStore CacheManager failures).

The validated `server.conf` settings, in `system/local` on each indexer:

```
[general]
parallelIngestionPipelines = 2

[cachemanager]
max_cache_size = 320000
```

Certificate trust

Place the certificate authority chain that signed the X10000 endpoint certificates, root and intermediates in PEM format, at the `remote.s3.sslRootCAPath` path on every indexer, and keep `remote.s3.sslVerifyServerCert = true`. The path names the signing authority, not the server certificate.

Applying the cluster bundle

Edit indexes.conf once on the active cluster manager and push it with `splunk apply cluster-bundle`. Some settings take effect only after a rolling restart, so confirm each change on an indexer; migrating an existing index instead needs a coordinated full stop and start.

Appendix D: X10000, endpoint, and network reference

GSLB pool configuration

Every S3 data IP of both clusters is a pool member, health-checked on TCP 443 every 10 seconds (5-second timeout, 2 retries). The pool uses topology-weighted round robin: each indexer receives a healthy data IP at its own site, or at the other site when none at its own site is healthy. It returns one address, and the record TTL is 5 seconds; topologies map each site's indexers to its own cluster's members. The fallback setting (any) applies only if every member of both clusters is down.

The lab used a Loadbalancer.org appliance (LBOS-7 release 8.13.4). The `lbcli` commands below show the configuration; substitute your own addresses. Other GSLB products need the same four objects: a global name, members, a pool and topologies.

```
# Global name for the shared endpoint, 5-second record TTL
lbcli add-gslb-globalname --name s3.smartstore.lab \
  --hostname s3.smartstore.lab --ttl 5

# One member per S3 data IP, 24 per cluster
# (repeat for site 2 as c02-01 to c02-24)
lbcli add-gslb-member --name c08-01 --ip <site1-data-ip-01> --weight 10
```

```
...
lbcli add-gslb-member --name c08-24 --ip <site1-data-ip-24> --weight 10

# Pool: one --add-member for each of the 48 members
lbcli add-gslb-pool --name s3-pool --add-globalname s3.smartstore.lab \
  --add-member c08-01 ... --add-member c02-24 \
  --lb-method twrr --monitor tcp --monitor-port 443 \
  --monitor-interval 10 --monitor-timeout 5000 --monitor-retries 2 \
  --fallback any --max-addrs-returned 1

# Topologies: each site's cluster data subnet, plus the address
# each indexer sends its DNS queries from (see DNS delegation)
lbcli add-gslb-topology --name site1 \
  --add-ips "<site1-data-subnet>, <idx01-query-ip>/32, ..., <idx04-query-ip>/32"
lbcli add-gslb-topology --name site2 \
  --add-ips "<site2-data-subnet>, <idx05-query-ip>/32, ..., <idx08-query-ip>/32"

# Apply the configuration
lbcli reload-gslb
```

DNS delegation

The GSLB answers DNS on port 53 and matches each query's source address against the topologies. In the lab, each indexer's first DNS server was the GSLB's address on its own site's data network, so queries arrived from the indexer's own data IP; the next server, a PowerDNS recursor, answered all other names. Site resolvers can instead forward the endpoint name to the GSLB (on PowerDNS Recursor, `forward-zones+=s3.smartstore.lab=<gslb-ip>:53`). The GSLB then sees the resolver's address, unless the resolver passes the client's subnet with EDNS Client Subnet, which the Loadbalancer.org GSLB uses by default. Without it, give each site its own resolver and list it in that site's topology.

Certificate names

Each X10000 cluster's certificate lists the shared endpoint name, `s3.smartstore.lab`, as a subject alternative name, plus the cluster's own name for direct checks. The HPE GreenLake certificate request has only a common-name field, so ask the certificate authority to add the subject alternative names when it signs.

AWS CLI setup and direct-to-cluster checks

Run the AWS CLI from an indexer or another host on the data network, trusting the same certificate authority and with `addressing_style = path`. To compare the clusters, query each by its own name, as shown in Confirming both sites are ready for cold searches.

Appendix E: Test method detail

Data generation

The four forwarders ran the open-source gogen generator (version 0.13.0) with a weblog template based on gogen's weblog example. It produces Apache `access_combined` web-server events for a fictional online shop: client IP addresses and user agents come from sample lists, HTTP status codes are weighted toward 200, and response sizes range from 200 to 4,000 bytes. Events went to the index `ss-high-volume` with sourcetype `access_combined` and host `gogenblast`, spread across all eight indexers. The searches in Search definitions select on fields from this template.

Corpus volume verification

Ingest volume came from the Splunk license meter, which counted only this workload, and was cross-checked against bucket sizes from `dbinspect` and the event-time distribution.

Search definitions

All three searches ran over the index `ss-high-volume`. Search A matched `clientip=74.125.19.106 status=200 useragent=BlackBerry9000* product_id=HolyGouda`; Search B matched a disjoint set (`clientip=27.102.11.11`, status other than 200) and was the target of the delete test; Search C was a report-accelerated form of Search A.

For the cold runs that Splunk's test plan requires, the validation emptied the cache on every indexer with the command below, run twice per indexer. Normal operation does not need this; the cache manager evicts buckets as space is needed.

```
splunk _internal call /services/admin/cacheman/_evict \
  -post:mb 1000000000 -post:path $SPLUNK_DB -method POST
```

Appendix F: Glossary

Table 8. Glossary

Term	Meaning
Splunk bucket	A directory of index files for one span of indexed data. It is hot while being written, warm after it rolls and uploads, and frozen at the retention limit.
S3 bucket	The object-storage container on the X10000. All SmartStore indexes share one S3 bucket, each under its own key prefix.
SmartStore cache	Local copies of warm buckets on each indexer, managed by the cache manager and evicted when space is needed.
Remote volume	The indexes.conf stanza that names the S3 bucket, endpoint and credential. Every SmartStore index points at it.
Replication and search factors	The number of copies, and of searchable copies, the indexer cluster keeps. Under SmartStore they apply to hot buckets; for warm buckets they govern only metadata.
Site	A data center in a multisite indexer cluster. Search heads set to `site0` have site affinity disabled.
Cluster manager	The Splunk node that coordinates the indexer cluster. The validation ran one active and one standby, one at each site.
GSLB	Global server load balancer. It answers DNS for the shared S3 endpoint name with the site's own X10000 cluster and, on a health-check failure, with the other cluster.
Delete marker	The object version that a delete writes in a versioned S3 bucket. Replicating delete markers carries deletions to the other cluster.
RPO	Recovery point objective: how much recent data a deployment can afford to lose. No RPO was defined or validated.

Visit HPE.com

Chat now

© Copyright 2026 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

a50014133ENW

HEWLETT PACKARD ENTERPRISE

hpe.com

