Glass

Splunk Certified App

Powered by
Spectra Logic BlackPearl and Geyser Data Tape as a Service

Create and manage a frozen tier for Splunk, reducing storage costs by 9x and improving search performance. Available on premises with Spectra Logic BlackPearl or as a service via Geyser Data.

Overview

Glass is a Splunk Certified App that automates the creation of a frozen tier for Splunk by periodically moving frozen Splunk buckets to S3 or S3 Glacier compatible storage. When an investigation or compliance review requires older data, a simple operator-driven restore brings it back. Glass supports Spectra Logic BlackPearl and Geyser Data Tape as a Service.

The need for security log archiving

IBM's Cost of a Data Breach Report 2024 found that breaches involving stolen credentials took an average of 292 days to identify and contain. Investigations need months or years of log history, which is prohibitively expensive to keep on Splunk cold tiers. Glass lets users search frozen data by time and retrieve it without reindexing.

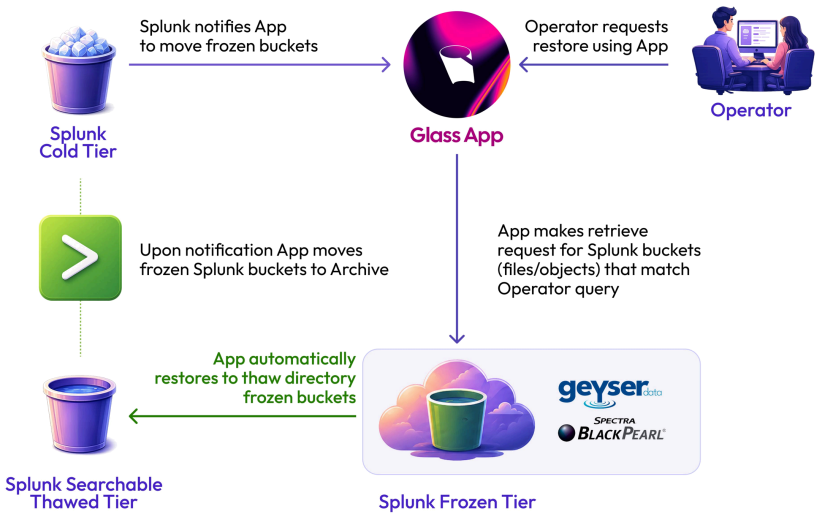

Glass is ideal for

Threat hunting

Search logs months or years old to uncover long-running or previously unnoticed malicious activity.

Regulatory compliance

Meet HIPAA, PCI-DSS, CJIS and other mandates without over-provisioning Splunk infrastructure.

Performance

A smaller active index directly speeds up real-time Splunk searches and dashboards.

Cost control

LTO tape delivers the lowest $/TB for multi-year retention.

Why Glass?

Archive and restore

Automates frozen-tier archiving from Splunk to the archive tier via the S3 API. Operator-driven restores need no manual bucket management.

No reindexing

Frozen logs are restored directly to the Splunk thawed directory, so a full search is available in minutes, not hours.

Flexible deployment

Supports both on-premises and cloud deployments of Splunk.

Validated configs

Pre-validated 1, 3 and 5 TB/day configurations for rapid deployment and predictable total cost.

Speed up Splunk

Moving data to the frozen tier reduces the size of Splunk indexes and accelerates performance.

Splunk-native experience

Managed entirely within Splunk and available on the Splunk Marketplace. No additional UI or operator training.

Splunk Certified App

Glass for Splunk is the only Splunk Certified App that cuts storage costs and speeds up search performance across every Splunk deployment.

Get started

Contact your **Spectra Logic** representative, or find Glass on the **Splunk Marketplace** or at **searchbak.com**.

Prism

Retain. Investigate. Defend.

Two gaps have opened in your cyber security armor, between what your tools see in real time and what you need to answer months later. Both end the same way: **the record you need no longer exists. That is what Prism closes.**

Mind the agent gap

An AI agent inherits its launcher's permissions, so nothing it does looks anomalous, and agents that exceed those permissions often leave no audit trail. Endpoint vendors analyze the runtime graph, but it carries no data classification and lasts only 7 days. You can see what an agent did this week, not which of 400 files it touched held sensitive information. Next month, you see nothing at all.

Mind the age gap

A SIEM is built for real-time defense. To stay fast and economical, the hot index can't keep growing, so popular platforms purge indexed data on a 90- or 180-day cycle. The intrusions that matter outlive that window: BRICKSTORM averaged 393 days before discovery, and SUNBURST persisted for more than fourteen months. Google's BRICKSTORM guidance asks for at least 400 days of retained syslog.

53%
of organizations have had AI agents exceed their intended permissions.
Cloud Security Alliance, April 2026

33%
keep no audit trail for agent activity, so nothing survives to reconstruct.
Kiteworks 2026, 459 organizations

393 days
average BRICKSTORM dwell time before discovery, longer than most SIEMs keep logs.
Google Threat Intelligence, 2025

30 months
federal log-retention requirement, against the 180 days your index holds.
OMB M-21-31; NARA GRS 3.2

Three Prism solutions on your Lakehouse

AI Agent Defense

Every session recorded, classified and correlated

- Omnigent, the open meta-harness from Databricks, covers Claude Code, Codex, Cursor, OpenCode, Pi and your own agents.
- Joins agent traces with EDR and SIEM telemetry and your sensitivity labels.
- Models sessions over time to surface drift, sending deviations to your SIEM.

Long-Timeline Threat Analytics

Every new detection, replayed across years of history

- Replays new and updated ATT&CK checks across your history, so every detection becomes a regression test.
- Checks your history for indicators when a new attack is published.
- Pinpoints when a breach began, not when it was noticed; findings go to your SIEM or ticketing.

Log Transformation & Retention

One normalized copy, kept for the mandated years

- Connects to the major SIEMs and EDRs, normalizing every event to OCSF in your Lakehouse.
- Meets the 30-month federal mandate and multi-year identity retention requirements.
- Deeper baselines for threat hunting, so fewer deviations turn out to be noise.

How the pieces fit together

Prism transforms events from connected SIEMs into OCSF and writes them in open Apache Parquet format to any S3-compatible storage. The Lakehouse, storage tier and query engine are your choice, and Prism isn't required to read the tables once written.

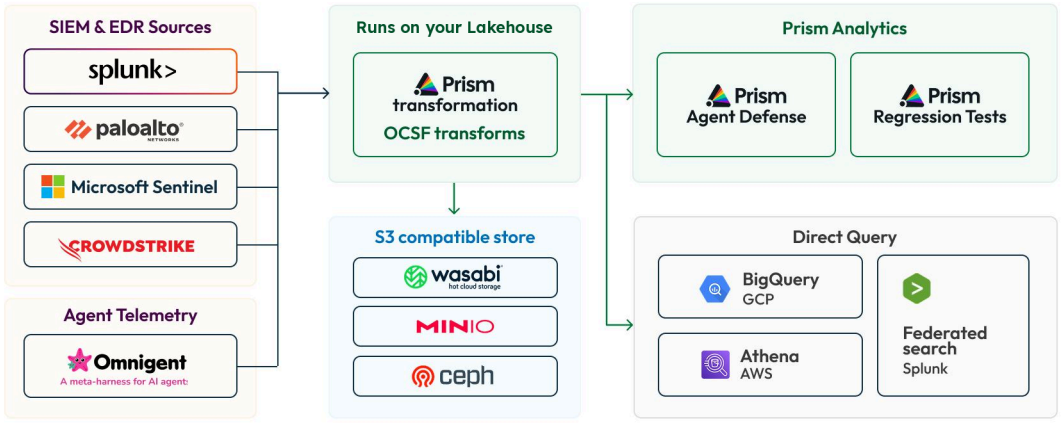

Storage tiers
Wasabi, MinIO, IBM Ceph, NetApp StorageGRID, Quantum ActiveScale, Geyser Data

Lakehouse targets
Databricks (Built On Databricks partner) and AWS at launch; Google BigQuery and Microsoft Fabric read the same open tables

Sensitivity labels
Databricks Unity Catalog and Microsoft Purview