SOLUTION SHEET

Scality RING & Splunk Enterprise

Index larger data sets and retain them longer — at a fraction of the cost — while reducing complexity and increasing flexibility.

Digital transformation is dependent on data — lots of it — and, most notably, machine-generated data. Splunk excels at enabling organizations to extract value from such data for business insights. Splunk collects, monitors and analyzes everything from log files, clickstreams, and transactions to network activity and call records — data that enterprises use to make critical decisions to protect and grow their business.

The problem is incoming data sets are getting larger, and retention times are growing longer (to enable deep searches and meet compliance requirements). The demand for storage is outpacing demand for compute, creating an imbalance in the traditional, classic Splunk architecture. Splunk SmartStore mitigates this issue and reins in storage costs.

SmartStore splits compute and storage resources, enabling a hot cache of data to be available to Splunk's indexers for analysis while maintaining warm and cool data on object storage, such as Scality RING. Data movement between Scality and Splunk is transparently managed so the right data is available for analysis by Splunk and simultaneously protected with extreme data durability by Scality.

SmartStore allows compute and storage resources to be scaled independently, which increases agility while lowering costs. How? The more expensive compute tier and the less expensive storage tier can be sized separately based on the appropriate balance of data analysis and data storage. Additional Splunk compute power can be added without impacting storage, and additional Scality storage can be added without impacting compute.

Another way SmartStore reduces storage costs is by offloading older, less frequently used data (warm buckets) to Scality RING, without impacting search performance. Scality RING holds the master copy of warm buckets while the local storage associated with Splunk indexers is used for the hot cache.

With Splunk SmartStore, the majority of data is stored on cost-effective, highly durable Scality RING, which can be deployed on a single site, replicated across two sites, and even synchronously written to three sites for maximum data durability. The Splunk indexers are stateless and can be started and stopped based on business requirements. When started, the indexers fetch the required data from Scality RING.

Key Benefits

Lower Cost

Scality RING and Splunk SmartStore allow for decoupling of storage and compute resources, so each can scale independently. Data storage costs are further optimized by offloading warm buckets from premium indexers' local storage into cost-effective object storage. The result? 70%+ reduction in Splunk Indexers, which can save millions of dollars in infrastructure costs.

Improved Searches

Scality RING's limitless scalability and Splunk SmartStore's native, transparent data tiering allow for larger data sets to be on-hand and for a longer period of time, paving the way for deeper and more accurate searches.

Better Operational Flexibility and Agility

Scality RING and Splunk SmartStore together render indexers stateless for warm and cold data. This boosts operational flexibility and agility, and streamlines the deployment and management model with higher elasticity.

Reduce Risks

With Scality RING, there is no need to back up warm and cold data any longer, which makes it easier to meet RPO/RTO targets. Furthermore, the extreme resiliency and data durability Scality RING offers (through erasure coding and geo-distribution) protects against disk, server, and even entire site failure, ensuring that Splunk always has access to the data it needs.

SCALITY

COST SAVINGS WITH SCALITY AND SPLUNK SMARTSTORE

Reduce Splunk Indexers by 70%+

splunk> **Classic vs SmartStore**

1TB daily ingest	**365** vs **30** TB Splunk storage	**31** vs **8** Indexer Nodes	**$465K** vs **$80K** Total Cost
2TB daily ingest	**730** vs **60** TB Splunk storage	**62** vs **16** Indexer Nodes	**$930K** vs **$160K** Total Cost
3TB daily ingest	**1.095** vs **90** TB Splunk storage	**93** vs **24** Indexer Nodes	**$1.395K** vs **$240K** Total Cost
4TB daily ingest	**1.460** vs **120** TB Splunk storage	**124** vs **32** Indexer Nodes	**$1.860K** vs **$320K** Total Cost

70%+ reduction

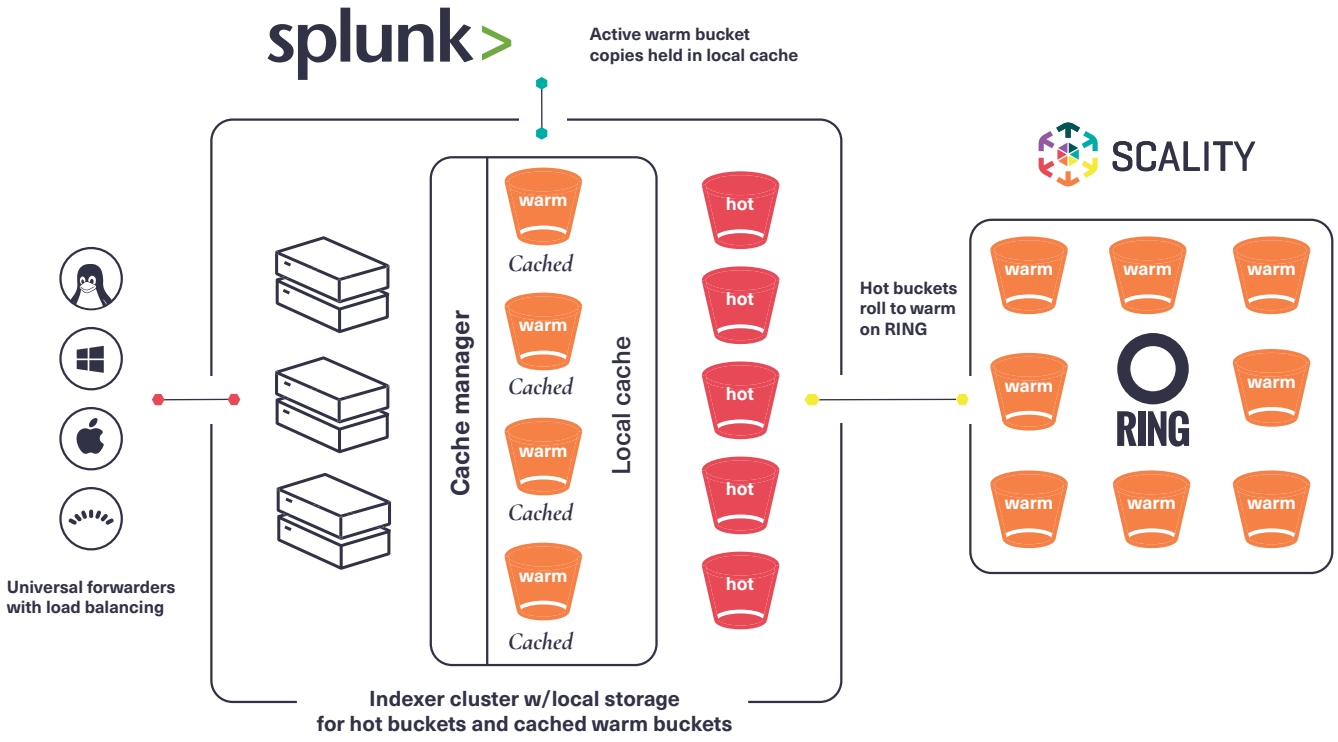

Splunk SmartStore's disaggregated architecture is a perfect match for the limitless scalability, high throughput performance, and extreme resiliency of Scality RING. Native File (NFS) support means Scality RING is a great fit for non-SmartStore deployments as well.

ABOUT SPLUNK Splunk is the world's first Data-to-Everything Platform designed to remove the barriers between data and action, so that everyone thrives in the Data Age. We're empowering IT, DevOps and security teams to transform their organizations with data from any source and on any timescale. Splunk's analytics-driven IT solutions enable customers to accelerate their IT modernization journey, providing the technical agility, speed and visibility needed to drive actionable business outcomes.

ABOUT SCALITY Scality® storage propels companies to unify data management no matter where data lives — from edge to core to cloud. Our market-leading file and object storage software protects data on-premises and in hybrid and multi-cloud environments. With RING and ARTESCA, Scality's approach to managing data across the enterprise accelerates business insight for sound decision-making and maximum return on investment. To compete in a data-driven economy, IT leaders and application developers trust Scality to build sustainable, adaptable solutions. Scality is recognized as a leader by Gartner and IDC.

Follow us on Twitter and LinkedIn. Visit www.scality.com, or subscribe to our blog.
San Francisco. Paris. Washington, D.C. Tokyo. London.

VTSS-202010

SCALITY