Hewlett Packard Enterprise | splunk>

Solution brief

Check if the document is available in the language of your choice.

SCALE UP YOUR DATA WITH HPE EZMERAL RUNTIME ENTERPRISE AND SPLUNK ENTERPRISE

Unlock the power of analytics-driven security

Solution overview
HPE Ezmeral and Splunk partnered to leverage the power of the HPE Ezmeral Runtime Enterprise and open-source Kubernetes to bring agility and scale to the Splunk Enterprise.

*** "...organizations believe that 68% of their data is dark, meaning that only 32% of their data is operationalized and able to deliver value."**

– "What Is Your Data Really Worth?" Splunk 2020

The exponential data growth from increased infrastructure, application, compliance requirements, and network traffic have outpaced the ability for many organizations to effectively collect and leverage data to detect, alert, and prevent security threats. As data rates exceed an organization's ability to index, store, and analyze the data, it is estimated that more than half* of the security-related digital exhaust isn't even being utilized in security analytics because the data simply can't be loaded and processed in time. This means there are huge blind spots in IT security and operations.

The problem is that legacy infrastructure not optimized for Splunk has resulted in the underutilization of oversized systems. This excessive infrastructure bloat means that data centers are nearing maximum capacity as IT support teams struggle to scale out ingestion, processing, storage, and analysis of the data. The increased infrastructure coupled with the growing backlog of data and security insights is forcing organizations to find ways to optimize their delivery and consumption of Splunk analytics.

HPE Ezmeral and Splunk partnered to leverage the power of the HPE Ezmeral Runtime Enterprise and open-source Kubernetes to bring agility and scale to the Splunk Enterprise. This has two immediate benefits: the solution deploys new indexer and search heads in a matter of minutes and independently scales them up within a host to fully saturate the infrastructure and scale out across the entire available information estate.

The Splunk Operator for Kubernetes (SOK) within the Splunk Enterprise allows administrators to deploy and operate enterprise deployments of Splunk in a Kubernetes environment. SOK is a container and set of instructions that has been designed to deploy and manage the resources of even the most complex instances of Splunk and integrates it with the underlying Kubernetes tools and processes.

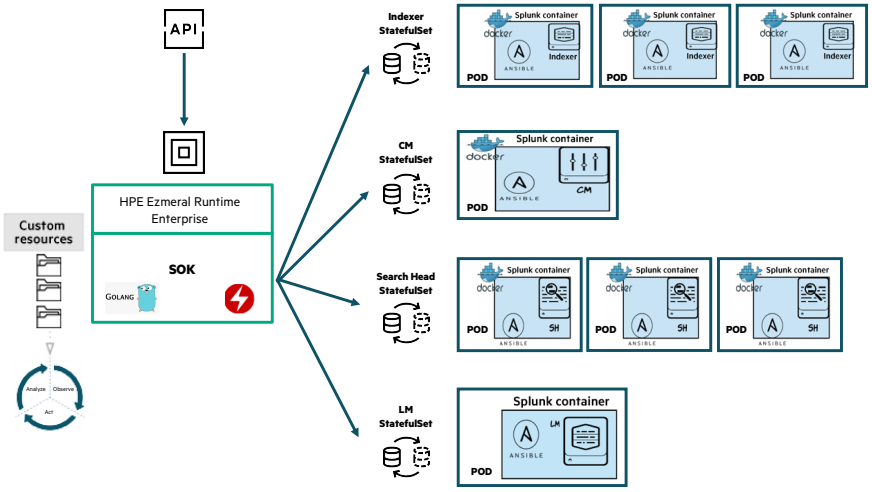

FIGURE 1. Graphic representation of SOK and HPE Ezmeral Runtime Enterprise solution architecture

KEY CAPABILITIES OF SPLUNK ENTERPRISE

Splunk Enterprise collects data from any source, including metrics, logs, clickstreams, sensors, stream network traffic, web servers, custom applications, hypervisors, containers, social media, and cloud services. It enables users to search, monitor, and analyze that data to discover powerful insights across multiple use cases such as security, IT operations, application delivery, industrial data, and IoT.

Additionally, with the power of machine learning baked in, it helps to make faster, more informed decisions across the organization.

- **Collect and index data:** Collect data from virtually any source and location. Convert logs to metrics and freely analyze and correlate data without the limitations of conventional database structures. Import data from relational databases and data warehouses for a complete business view.
- **Search, analyze, and visualize:** The Splunk search processing language supports all of search needs—from the simplest to the most sophisticated. Point-and-click analysis brings insights to business users. Rich visualizations then make results understandable and actionable for all audiences.
- **Monitor, alert, and report:** Set thresholds to monitor for incidents or proactively signal potential issues. Use alerts to launch applications or custom actions. Interact with your data with custom dashboards that can be shared or embedded into other applications as PDFs.
- **Apps and premium solutions:** Apps deliver a targeted user experience for typical use cases and data sources. Splunk Premium Solutions combine real-time analytics and rich features to manage your security posture, IT operations, and more.

KEY CAPABILITIES OF HPE EZMERAL RUNTIME ENTERPRISE

HPE Ezmeral Runtime Enterprise is designed for both cloud-native and distributed non-cloud-native applications, enabling true hybrid cloud operations across any location: on-premises, public clouds, and edge.

- Bring the speed and efficiency of containers to both cloud-native microservices apps and non-cloud-native monolithic apps
- Deliver new code releases faster with one-click container deployment
- Build once and run anyplace (on-premises, public clouds, and edge), providing hybrid cloud portability

The HPE Ezmeral Runtime Enterprise has been jointly validated** with SOK bringing additional benefits to open-source CNCF Kubernetes but without modifying the open-source code, bringing the value add around and to Kubernetes.

- **Ease of access:** Create, control, and manage solution environments through a single pane of glass
- **Rapid cluster creation:** Spin up ready-to-run unmodified clusters in minutes via self-service UI or APIs
- **Multitenancy:** Run completely isolated application clusters and distributions on shared infrastructure
- **Elastic:** Add or remove nodes on demand, with elasticity to quickly expand or contract cluster size based on workload
- **Data access:** Pre-integrated persistent container storage and data fabric; access data sets without data duplication via DataTap and IOBoost technology
- **Security:** Complete end-to-end security with Active Directory (AD)/ Lightweight Directory Access Protocol (LDAP), Kerberos, and Transparent Data Encryption (TDE) integration

WHY SPLUNK ON HPE EZMERAL RUNTIME ENTERPRISE?

The tight collaboration between HPE and Splunk brings dark data to light by making it simple to collect, analyze, and act upon the untapped value of the Big Data generated by your technology infrastructure, security systems, and business applications—giving you the insights to drive operational performance and business results. The enhanced delivery of Splunk creates a single datastore of all machine data that leverages open-source Kubernetes and S3 and is available as a fully managed as-a-service (aaS) solution from HPE.

- Shrinks the infrastructure footprint and significantly lowers TCO with a loosely coupled architecture that independently scales search heads, indexers, and storage
- Adds new use cases deploying new indexers and search heads in minutes
- Leverages Splunk SmartStore to efficiently balance hot cache and S3 object storage for exabyte scale cold data

LEARN MORE AT

hpe.com/us/en/software/marketplace/splunk.html

** splunk.github.io/splunk-operator/

Make the right purchase decision.
Contact our presales specialists.

Chat Email Call

Get updates

Hewlett Packard Enterprise

a50004569ENW, Rev. 1