Application Brief
Extreme Search®

Blazingly Fast Self-Searching Storage for Cyber Hunt Kits

Lewis Rhodes Labs' next generation scalable search technology on Intel Agilex® searches raw data in storage in under 12 minutes without indexing, accelerating all Big Data tasks.

Extreme Search® is a branded product produced by Lewis Rhodes Labs, Inc. (LRL) that uses proprietary LRL NPU*search*™ technology to execute in-storage searches at extremely high speed while using less than 1% of the power required by traditional methods. This second-generation product is a 2U appliance based on a Dell 750xa that performs regular expression-based search of 64TB of SSD storage at sustained rates exceeding 90 GB/s. Search bandwidth is independent of the data content or expression set so searches consistently complete in less than 12 minutes.

Current instantiations of NPU*search* are optimized to scan file systems for sparse content. This consideration greatly simplifies integration into complex applications and mission spaces. Prefiltering files in storage requires no modifications or new ATOs in critical platforms such as Elastic. The practical impact is that an application which previously may have ingested 3,000,000 files need only ingest the 120 or so that are relevant. The system level impact on power, space and response time can be game-changing.

LRL's NPU*search* technology is soft IP including FPGA images of the NPU processor with an extensive software stack including drivers, compilers, and other elements. The NPU*search* processors appear as devices on a PCIe bus. The software API presents as two python commands that may be used to integrate NPUsearch into traditional user interfaces such as those of Elastic and Splunk.

Mission Drivers

- Rapid threat assessment and response
- Improved SWaP and sustainability
- Control and prioritize awareness data
- System scalability for data and analytics
- Mission advantage and faster response

Application Overview

Rapid evolution of sensor technology and the resulting telemetry streams have vastly outpaced the network, storage, and computation required to make them actionable. Size, weight and power (SWaP) limitations impact data triage at every stage of the analytics pipeline. Cyber missions are struggling with this challenge. Vast data flows lead to reduced situational awareness in a theater where efficiency and response time are critical. Extreme Search directly addresses SWaP restrictions, converting large data sets into actionable events from the sensor edge to the computational core. Extreme Search capability deployed in a Cyber Hunt Kit provides a differential mission advantage by reimagining Big Data utility at the edge.

Deployment

Cyber Hunt Kits (CHKs) are mobile Security Operations Centers (SOC) optimized for rapid field deployment. An essential cyber mission capability is PCAP capture and analysis. A modern CHK ingests a petabyte of PCAP per week and runs a variety of security applications such as Elastic Stack, Splunk, Suricata, Stenographer, and Zeek. These kits must be ruggedized appropriately for field deployment, be easily transportable, and function within stringent SWaP requirements. Traditional technology requires CHKs to rely heavily on creative indexing methods to provide analysts with data visibility; however, these techniques are failing to keep step as data volumes increase.

Consider a comparison workload where ingesting 30TB of uncompressed PCAP into a state-of the art server using an Elastic Stack architecture (**Fig. 1**) takes 7 days to complete. Extrapolating this benchmark, a CHK would require an unrealistic 30-35 equivalent servers for a full duplex 10GB Ethernet cyber tap. There are current workarounds for this problem including sampling, ingest lag, and reduced indexing, but all significantly degrade mission capability.

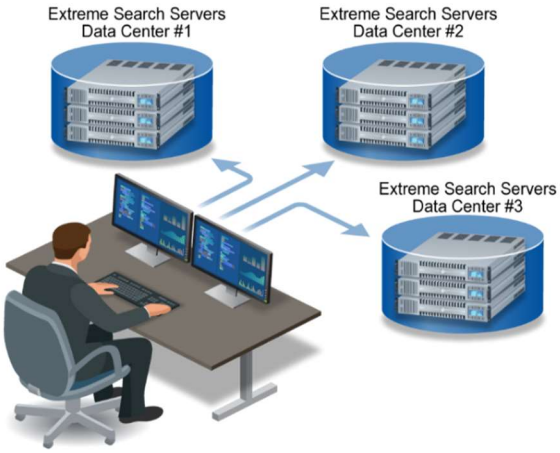

Figure 1. Find rare events fast in unprocessed data for unprecedented situational awareness on distributed data.

Since solutions such as Elastic Stack and Splunk are only efficient at retrieving indexed data; full and timely visibility requires all potentially useful data be indexed upon ingest. In practice, however, most of the information indexed will never need to be accessed. Extreme Search upends this burden by rapidly retrieving unindexed data within massive storage arrays and indexing the additional files only *as needed*.

When Extreme Search was incorporated into existing Elastic Stack architecture (**Fig. 2**), the previous 30 TB, seven-day ingest benchmark completed in under an hour while consuming a fraction of the power. This lower SWaP data processing pipeline allows raw PCAP to be transferred directly in storage.

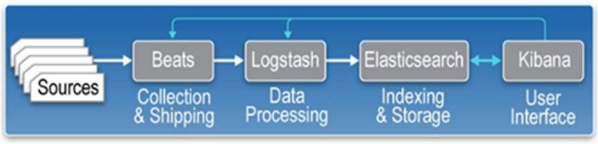

Figure 2. Elastic Stack data processing pipeline is computationally intensive but is required for timely data visibility using traditional methods.

The Extreme Search Elastic plug-in (**FIG. 3**) rapidly identifies files of emerging interest as needed, optimizing computational resources, reducing ingest SWaP, and enabling specific, high resolution ingest.

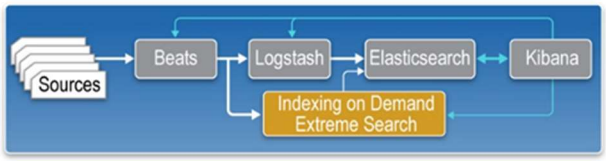

Figure 3. Schematic of the traditional architecture plus Extreme Search that reduces ingest, manages SWaP, and enables efficient data visibility by rapidly ingesting data as needed for evolving indexing requirements.

Summary

CHKs incorporate a large variety of complementary tools to collect, analyze, store, and manage PCAP. These tools in turn generate extensive logs, metadata, and alerts to help inform analysts. Existing toolkits struggle to process the increasingly large volume of diverse data containing relatively sparse actionable elements needed to generate a timely response within limited SWaP capacity. The addition of Extreme Search to a CHK provides a low SWaP data sync that extracts and inserts enriched files on demand—as needed—into the powerful, higher SWaP, SIEM and SOAR environments for timely visibility and efficient use of actionable data. The availability of rapid search and predictable time-to-results with the Extreme Search plugin increases the capacity to automate mission critical tasks.