Technology Brief

Extreme Search® by Lewis Rhodes Labs

Blazingly Fast Self-Searching Storage for the Public Sector

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.

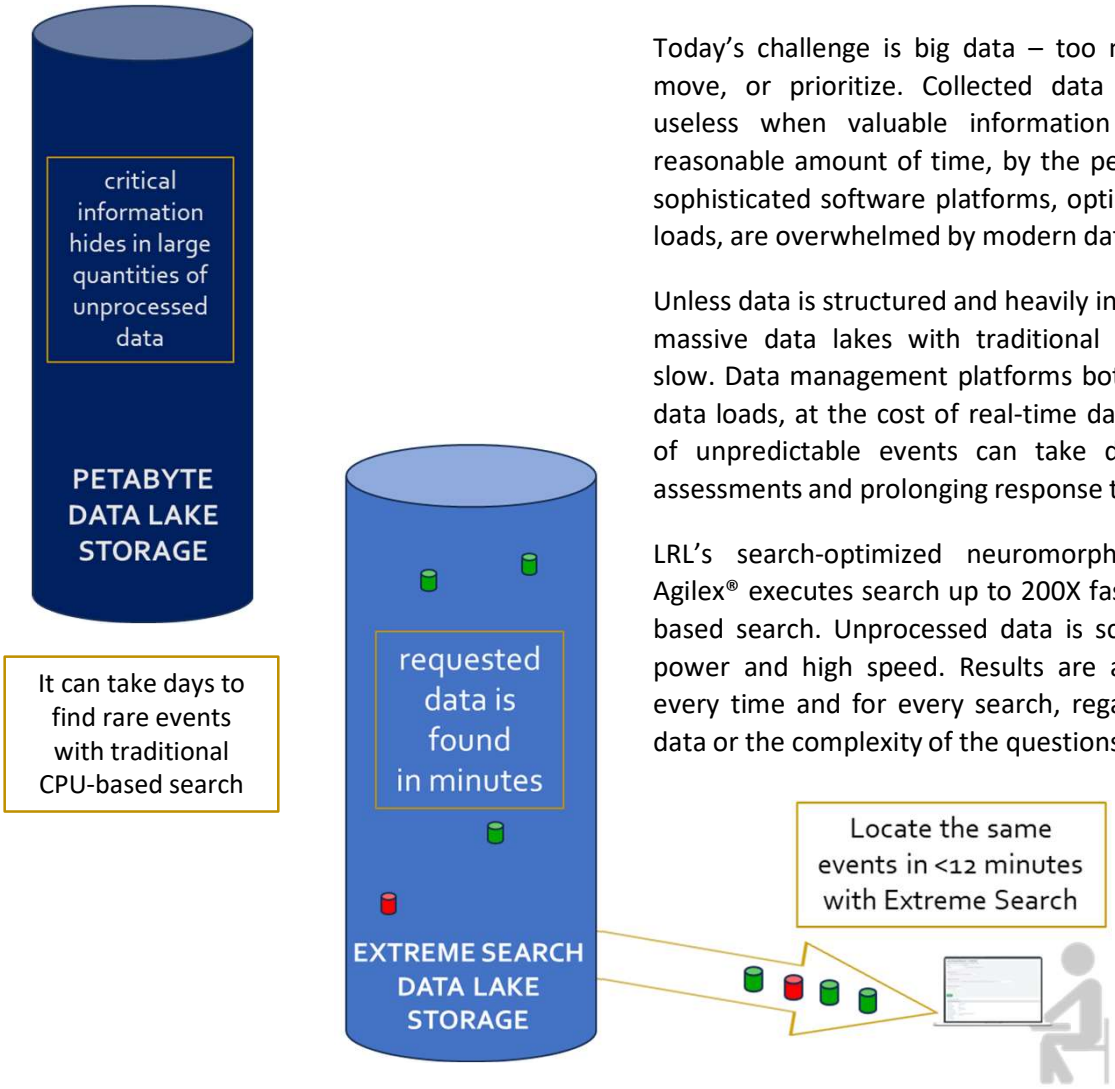

Figure 1. Find Rare Events Fast in Unprocessed Data

Introduction

Today's challenge is big data – too much to organize, index, move, or prioritize. Collected data can become effectively useless when valuable information can't be found, in a reasonable amount of time, by the people who need it. Highly sophisticated software platforms, optimized for moderate data loads, are overwhelmed by modern data volumes.

Unless data is structured and heavily indexed, searching through massive data lakes with traditional methods is prohibitively slow. Data management platforms bottleneck due to excessive data loads, at the cost of real-time data visibility. Identification of unpredictable events can take days or weeks, delaying assessments and prolonging response time.

LRL's search-optimized neuromorphic processing on Intel Agilex® executes search up to 200X faster than traditional CPU-based search. Unprocessed data is scanned in storage at low power and high speed. Results are available in <12 minutes every time and for every search, regardless of the amount of data or the complexity of the questions.

Central visibility on geographically distributed data.

Data is accumulating in remote locations without sufficient technology to fully utilize data value. Maintenance of essential security requires frequent and rapid identification of information that is sparsely distributed across these multiple locations. Escalating threat information further bottlenecks data visibility. Extreme Search storage provides a central operations analyst with visibility and access to PBs of data regardless of how many assets or where they are located.

Unprecedented access to remote data.

Remote assets can collect too much data to move, too much data to ingest, and too much data to search. Essential data may remain uncharacterized until an event becomes mission critical, or an asset is under attack.

Extreme Search® storage provides visibility across remote assets without indexing or moving data. Queries may be issued remotely and completed locally. Alerts can be generated without compromising data security. Visibility on events in geographically dispersed data collections is available on demand in <12 minutes. Extreme Search delivers a previously unavailable level of situational awareness for remote data.

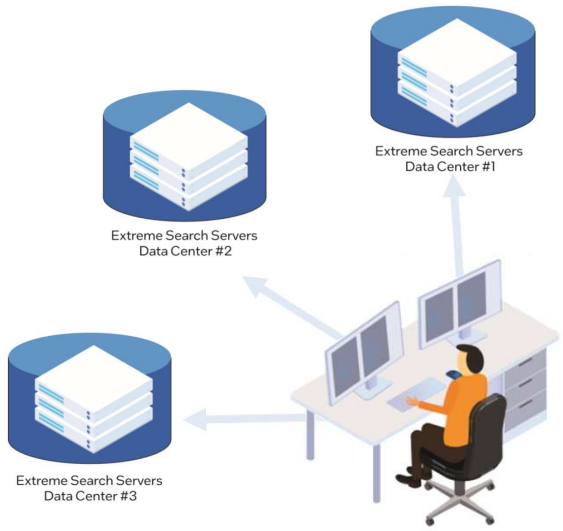

Figure 2. Unprecedented Visibility on Distributed Data

Combination search for evolving threats.

Cyber attack footprints change over time, requiring new capabilities to stay ahead of threats. The capacity to rapidly search for partial sets of expressions can provide real-time awareness of an evolving situation. Extreme Search combination functionality supports greater visibility on shifting indications of compromise.

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

In addition to standard regular expression search capabilities, Extreme Search allows users to search for files which match a minimum combination of a larger set of expressions. A user can search for all files which has text that matches at least *n* of the *m* expressions that they input, or all files which match some logical combination of the expressions that they input, such as “expr0 AND (expr1 OR NOT expr2)”.

Cyber attack footprints change over time, requiring new capabilities to stay ahead of threats. The capacity to rapidly search for partial sets of expressions can provide real time awareness of an evolving situation. Extreme Search combination functionality supports greater visibility on shifting indications of compromise.

Real Time Analytics for Public Sector

Slow data analysis delays mission response time. Specialized analytics on large data collections are backlogged by the long latency of data preparation. Lewis Rhodes Labs' Extreme Search was designed specifically to solve this problem.

Extreme Search servers generate the list of files containing matches for the patterns requested, where the file list can then be automatically fed directly into existing analytics toolsets and incident management platforms, or used to direct a specialized search with a small subset of files. And Extreme Search scales easily to accommodate data set growth from TBs to PBs without impacting search time.

Extreme Search rapidly identifies the data patterns of interest and feeds the next level of analytics – whether that be refined search, correction, outlier analysis, or AI/ML. And distributed search can include both local and remote data sets – all in real time.

Accelerate SIEM Tools

Slow search limits security and surveillance capacity while increasing risk of successful bad actor activity and cyber compromise. Data security and surveillance need fast search of raw data, because improved data visibility means decreased risk from cyberattacks.

Mission Benefits

- Unprecedented visibility on remote data
- Transform cyber security response
- Optimize utility of distributed data
- Improve SWAP for Big Data usage
- Scalable technology with efficient TCO

Globally distributed assets generate more and more data each day, leading to performance challenges – especially when trying to index all incoming data. Extreme Search finds indications of compromise in raw data in minutes instead of days, without complex ingest or indexing.

Specialized security information and event management (SIEM) solutions provide essential and robust response platforms, but raw data search remains necessary for unindexed events, and regular expression search can be painfully slow. Lewis Rhodes Labs' Extreme Search solution finds evidence of rare events rapidly and on demand, allowing more optimized threat assessment, detection, and remediation.

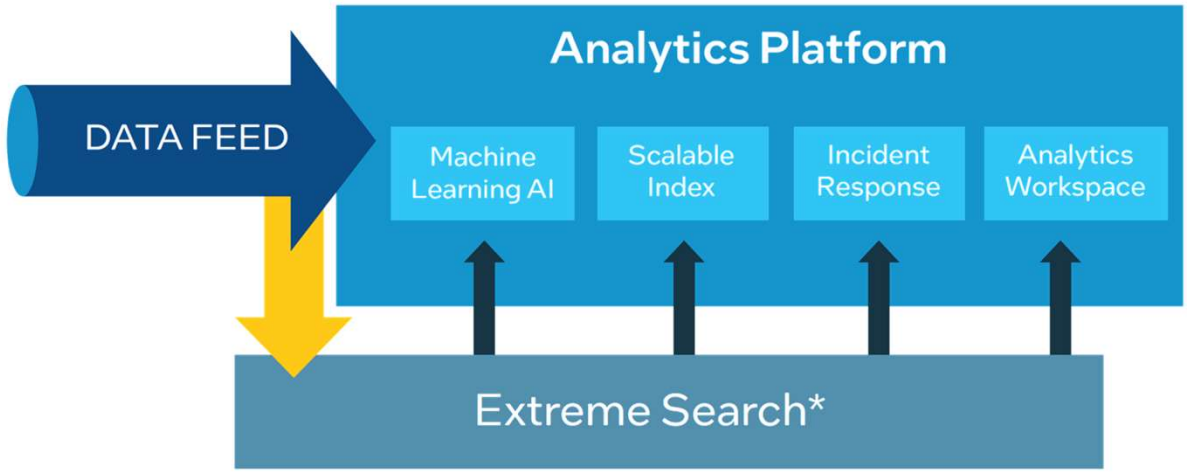

Figure 3. Rapid Search Accelerates Existing Analytics Platforms

Mission response depends on harnessing data potential. High performance, low power neuromorphic search for Big Data is a next generation approach that unleashes a new capacity for mission advantage.

Key Advantages of Extreme Search®

- Lewis Rhodes Labs search in storage technology, available in Extreme Search appliances, and in board level NPUsearch™ powered solutions through our partners, provides an entirely new capability to search massive amounts of unprocessed data extremely fast.
- High speed, low power neuromorphic processing generates very little heat, so full bandwidth search is possible within storage, without moving data.
- Full visibility is gained on huge data sets without the burden of ETL or indexing. The integrity of raw data collections can be fully maintained. No data is hidden by the compromises inherent in indexing. No data needs to be discarded to manage ingest rates.
- Extreme Search is easily accessed locally or remotely. Searches can be completed while data remains on location in storage.
- Regular expression-based search completes on all storage data in <12 minutes, regardless of the location, the diversity or type of data structures, or the complexity of the query. Search is accessed through a pair of python command lines, and specialized search features are available including Boolean logic and *n* of *m* analysis.
- Customer comparisons demonstrate Extreme Search storage servers use less than 1% of the power, and less than 1% of the rack space, as their existing systems for similar tasks.
- Search capability is fully scalable. Each 4TB SSD is linked to the neuromorphic search processor needed to scan it. Additional storage contains additional search capacity, and searches complete in parallel, so time to results doesn't increase. More storage can be added as needed within the distributed file system, and search still completes in <12 minutes.

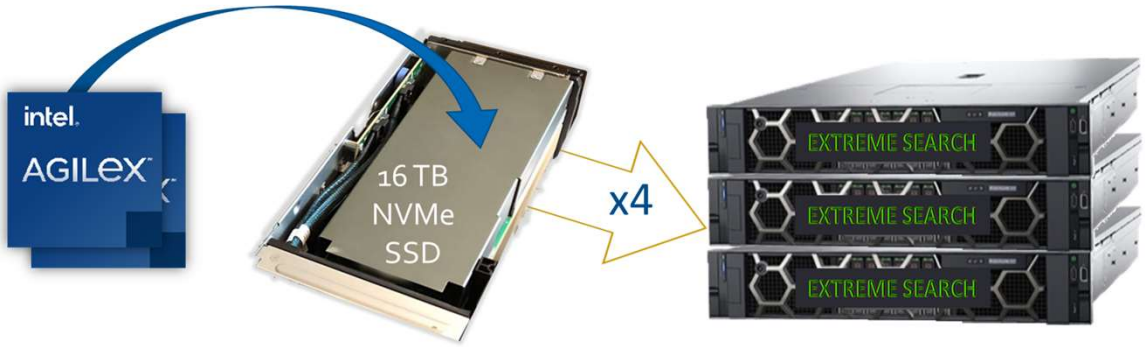

Figure 4. ExtremeSearch® storage is available in the latest servers with Intel® Xeon® processors.