

The Magic of Fluency[®]

Aladdin

A Look at Fluency's Innovative Database



At RSA Conference 2018, Fluency demonstrated its latest database version, Aladdin. Aladdin builds on Fluency's native Amazon Web Services (AWS) database by adding smart storage management and on-demand search processing. Even from a warm start, Fluency outperforms Elastic. And for long-term retention compliance requirements, Fluency also provides loading from a cold start. Once searching is fully loaded, massive searches are performed in just a few seconds. While other log management solutions rely on older database technologies and manual cold loads, Fluency continues to consistently outperform and out deliver.

A few advantages of Fluency include:

- 90% reduction of total cost of ownership
- Cost-effective use of storage and processors
- 50x faster indexing over Elastic
- Multi-year integrated Storage
- Multi-month detailed searching in seconds
- Searching is isolated from indexing
- Full PCI data retention compliance

Fluency was first recognized for its database innovations in 2015 by being named an RSA Innovation Sandbox Finalist, when we were known as SecurityDo. This year, Cyber Defense Magazine recognized Fluency's leap in technology at RSA Conference 2018 by presenting us with its coveted Cool Company Award that's presented to annually to groundbreaking, innovative companies.

Fluency out performs: Cold, Warm or Hot

Cost and speed are impacted by the type of storage used. Most solutions have an always-on hot-storage to compensate for slow database technology. The customer pays twice for such a solution. First, the search window is significantly lowered to reduce costs and increase performance. Second, the customer is paying for higher-end disk storage and search processing power, even when there is no user performing searches.

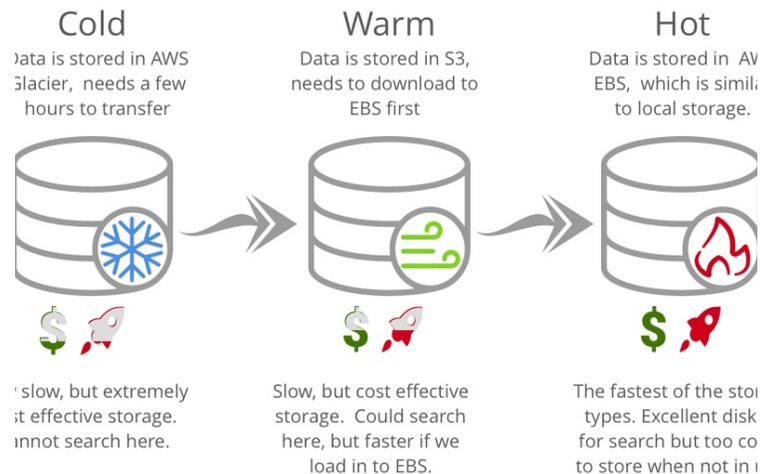


FIGURE 1: TYPES OF AWS STORAGE

Fluency is the only log management system that reduces cost and increases speed by moving data to the appropriate storage and assigning processing when needed. In Fluency, warm data is kept on S3 storage and moved into faster EBS storage once a user begins to search. When a user requests a search outside of the 90-day window, the system reaches back into AWS Glacier storage and moves cold data into EBS storage. When data is placed in EBS storage, Fluency assigns a higher amount of processing power.

The result of EBS searching with Fluency's LavaDB is so fast that searching a four-week block of warm data (2 minutes) is faster than a fifteen (15) node Elastic cluster searching the same period with always-on hot storage (about 5 minutes). Once loaded, Fluency searches will take two to eleven seconds, while the Elastic cluster takes minutes.

Cold Searching with Fluency

While storing data in cold storage is an optimal means of keeping log retention costs down, access to cold data when needed can become an issue. Due to how traditional relational and big data databases format and index data, it's not as simple as moving data from one

storage device to the other. Data is often kept raw and then reloaded, which can take hours or days. This is further complicated as data needs to be re-indexed and often cannot be properly correlated.

10.1.2.190:20180330;	Launch	Download	Decompression	Search
10.1.118.168:20180329;	Launch	Download	Decompression	Search
10.1.102.179:20180329;	Launch	Download	Decompression	Search
10.1.240.107:20180331;	Launch	Download	Decompression	Search

FIGURE 2: COLD STORAGE SEARCH

Fluency addresses cold storage by its streaming design. Data remains in a database format, performs a simple move and then decompressed onto the faster drive. The four-phase process of instantiating a process (launching), moving the data (downloading), decompressing and searching can be seen in the timeline tracker of the Fluency interface (Figure 2: Cold Storage Search). At the first use of this data, the database is moved into EBS storage, making future searches on this dataset extremely fast.

Warm Searching

Warm searching moves data from S3 data storage to EBS. For a Fluency user, this will normally be the first search using the interface, after which the data will be in EBS, and the searches will be hot.

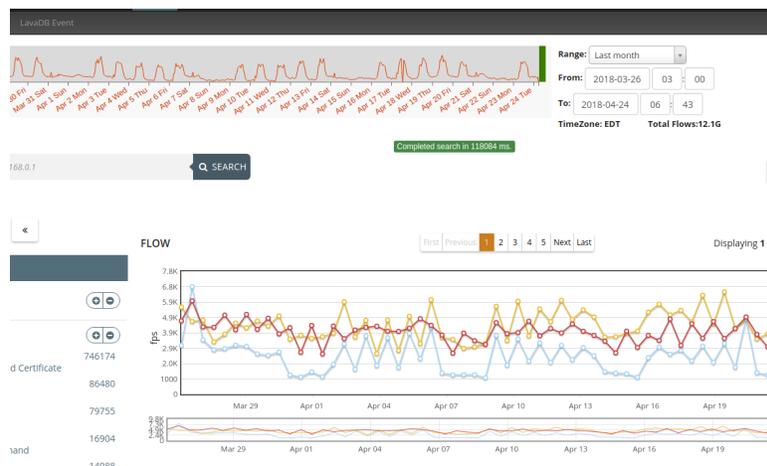


FIGURE 3: WARM SEARCHING

To the user, the warm storage search takes 118 seconds, while it covered four weeks and 12 billion events (Figure 3: Warm Searching). As stated, these speeds are faster than Elastic's hot searching. Warm searching allows that system to be efficient, keeping a 90-day window of data for analysis, while avoiding the use of more costly storage.

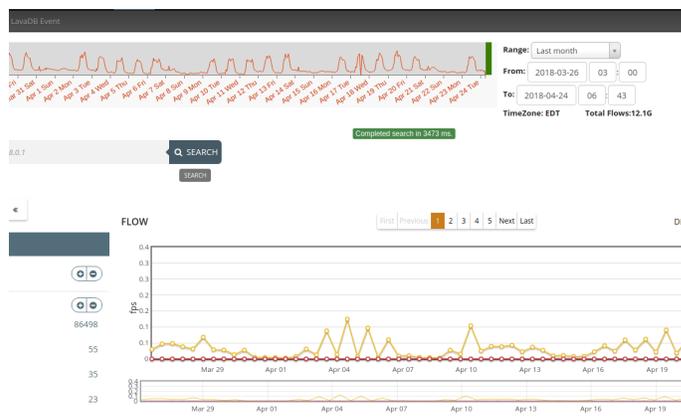
Hot Searching

The user will spend most of their time performing hot searches. When the data is hot, the user is interacting with a read-only copy of the data on the fastest data storage. As seen in a process timeline of a hot search, a series of processes operate in parallel searching the dataset. Compare that to Figure 2, and you see why Fluency's hot searching is so fast – not only is it using a fast disk, but it's doing so in parallel.

10.1.120.39:20180412;	Search			
10.1.163.223:20180417;	Search			
10.1.44.187:20180403;	Search			
10.1.249.79:20180423;	Search			
10.1.16.135:20180328;	Search			
10.1.16.85:20180406;	Search			
10.1.19.37:20180411;	Search			
10.1.238.44:20180417;	Search			

FIGURE 4: HOT SEARCH IS PURE SEARCHING

The result is a search over billions of events and multiple weeks only taking seconds. The previous example of 118 seconds is performed again, this time searching for a subset of data. The restful is 3.473 seconds. It takes longer to communicate and draw the page than it does for the search to be performed.



TYPE TO ENTER A CAPTION.

Summary

Fluency offers significant results for its customers through continual database innovation and leveraging cloud infrastructures. Fluency has always made it a point to deliver superior capability at a competitive cost. The Aladdin version allows Fluency to provide the log management industry's fastest database. It also provides the solution for long-term retention that can be searched, while keeping the cost in li