elastic | DISCOVER

SUCCESS STORY

Discover Financial Services cuts storage costs and accelerates data retrieval with Elastic Observability

Region
United States

Industry
Financial Services

Solution
Elastic Observability, Generative AI

Cuts storage requirements by 50%

With Elastic Observability, Discover deployed a hot-cold-frozen architecture that reduced storage costs by 50%.

Reduces data retrieval from days to minutes

Discover reduced the time to retrieve historical data from more than 24 hours to just minutes with Elastic Observability.

Accelerates mean time to repair

With Elastic Observability, Discover can minimize false positives and prioritize alerts reducing meantime to repair.

 |

elastic

Leading online bank deploys Elastic Observability to build a centralized logging platform that reduces storage costs and accelerates long-term data retrieval

Discover Financial Services (Discover) owns and operates Discover Bank, which offers millions of people in the U.S. online services such as checking and savings accounts, personal loans, and credit cards.

The bank must constantly refresh products and services in this highly competitive marketplace to retain and attract new customers. These efforts lead to exponential growth of application log data, putting pressure on existing IT infrastructure. Sunny Singh, Sr. Manager, Emerging Technologies, Discover Financial Services, says, “We are constantly challenged to balance platform scalability with cost optimization.”

To overcome these issues, the bank turned to a centralized logging platform based on Elastic Observability running on Elastic Cloud Enterprise. The logging pipeline begins with Filebeat data collection agents that capture and forward data to the bank’s Kafka platform where it is categorized into Kafka topics. Logstash then pulls data from Kafka, performs additional processing, and pushes it to one of two endpoints for log data storage and retrieval, each with its own retention period.

The primary endpoint, built on Elasticsearch, is composed of 10 separate clusters to ensure users only have access to logs relevant to their business role. Kibana serves as the visualization and analysis tool, enabling teams to create dashboards and view logs.

The second endpoint is an Amazon S3 bucket. This is used for longer-term log storage, typically 13 months, although some data is seven years old.

elastic

> Elastic immediately understood our request to significantly increase data storage capacity while maintaining efficient data retrieval times.

Sunny Singh
Sr. Manager, Emerging Technologies, Discover Financial Services

Centralized Logging Pipeline

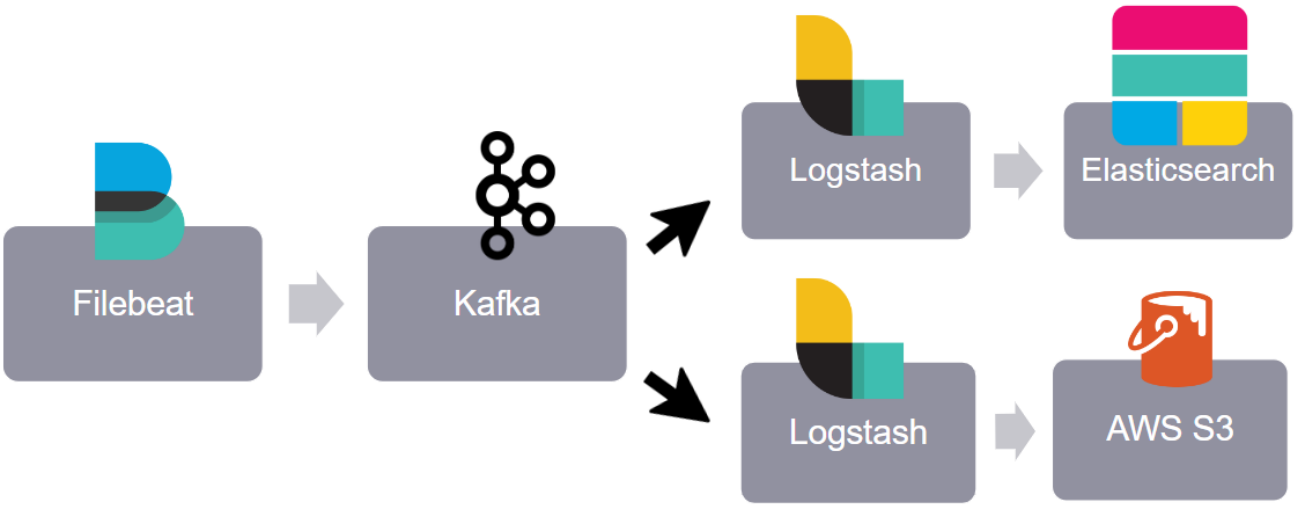

Managing log data volumes

The bank also wanted to improve log data management, as volumes can fluctuate from five million events to 50 million events within a three-hour window. Matt Keelan, Expert Application Engineer, Discover Financial Services, says, “Factors such as time of day, day of week, and application activity lead to spikes in data volumes. Unpredictable performance testing within development environments can also lead to a surge.”

Keelan and his team implemented several solutions based on Elastic Observability. One approach involved autoscaling Logstash based on the consumer lag of Kafka partitions. “By pushing consumer lag to CloudWatch metrics, we can trigger automatic scaling of Logstash instances in AWS. This ensures optimal resource utilization by dynamically adjusting Logstash capacity based on real-time data flow,” says Keelan.

Another solution tackles excessive logging from misbehaving applications. “We implemented rate limiting within Filebeat, which restricts the number of log messages received per second per file. This prevents a single application from overwhelming the system with unnecessary log data.”

Keelan also chose to migrate to data streams, enabling automatic index creation that eliminates the often cumbersome and time-consuming task of manually maintaining lists of indices, rollovers, and aliases. This move allowed Discover to institute Index Lifecycle Management (ILM) policies that further streamline and improve index management.

Another challenge was alerting for log events. “We needed to provide two outputs for these alerts: email and the Moogsoft alerting system that leads to our paging network,” says Keelan.

He achieved this goal by deploying Elastic Watcher, an advanced alerting and notification feature. “We’ve integrated with Watcher to create alerts, send emails, and integrate with Moogsoft so that users can be paged if there’s a problem with their application,” says Keelan.

Cooling down costs with a hot-cold-frozen architecture

Discover has also transitioned from an all-hot node architecture to a hot-cold-frozen setup for its Elasticsearch clusters. The hot nodes run on AWS Graviton servers for fast processing of large volumes of data. Cold and frozen nodes run on i3en attached storage to return query data.

Previously, the bank could only store up to 30 days of live data. Historical data retrieval involved manual re-indexing, often taking up to 24 hours or more to fulfill user requests. The new hot-cold-frozen architecture empowers users to access data directly from all three tiers. It also eliminates the need for re-indexing and significantly reduces retrieval times.

Frozen tier searches now only take a few minutes. “As well as improving performance, we have reduced storage requirements by 50%,” says Keelan. “In addition, the team can spend more time on developing new features instead of re-indexing data.”

Elastic Observability also plays a key role in helping to deliver uninterrupted services to customers. “With more efficient logging, analysts can quickly detect issues, eliminate false positives, and engage an engineer to resolve them,” says Keelan. “Since deploying the unified full stack from Elastic Observability, our mean time to repair has been reduced.”

Elastic Observability: A foundation for the future

One of the biggest benefits that emerged from the project is the trust and confidence that now connects Discover and Elastic stakeholders. Singh says, “The relationship that we built with Elastic and its engineering team is incredibly valuable and we have learned enormously from their expertise. It really gives our engineers confidence in the technology.”

It also gives Singh and the Discover team confidence for the future. “Elastic has helped us every step of the way, helping us scale our platform and take it to the next level,” says Singh. “They also give us valuable insights into new products and features, and how we can apply them to our own deployment.”

Singh also sees Elastic Observability as a key component of the organization’s AI strategy as it develops increasingly personalized services for its customers. “Discover is actively exploring the integration of generative AI models within its platform. By leveraging Elasticsearch’s scalability and integration capabilities, we can continue to deliver a cutting-edge online experience for our customers.”

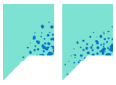

> Elastic definitely makes a positive difference to what we are doing. There is a constant release of new features that enable us to accelerate the development of new products and services for customers.

Matt Keelan
Expert Application Engineer, Discover Financial Services

