CASE STUDY

Predictive IT Operations at a Leading Global IT Services Provider

A leading global IT services provider wanted to reduce operations costs by predicting customer-impacting incidents from events in their data.

SVDS applied agile data science methods to detect patterns of anomalous events, and build a foundational framework for continued investigations.

Background and Business Problem

SVDS was engaged by a leading IT group that provides technology, backend operations, and IT services. The client has customers interacting with their platform and backend services, producing tens of billions of electronic transactions per day. These transactions can include: critical customer facing actions such as search, reservations, and internal processes such as inventory management. The challenge to reduce operational costs while maintaining a stringent service level is getting orders of magnitude more complex as the company evolves. A critical issue can take days to resolve by an operator, potentially resulting in millions of dollars in loss of revenue for the client's customers.

The client hypothesized that the log files, created through customer transactions processed on the backend, could be used to decipher operational insights and challenges, predicting service outages leading to customer incidents. Based on discussions with the client, we determined that the most immediate and valuable area of the system to target was the internal enterprise service bus (ESB). If anomalous system behavior could be detected before an incident record was raised, a lot of value could be captured: better ESB incident categorization, prioritization, and ultimately automated resolution could reduce staffing demands, and improve customer experience dramatically.

The client's ESB routes all customer transactions to their backend services. It is heavily instrumented, and every routed transaction creates a log entry containing the routing path of the transaction, involved backend components, and response times, among other things. Based on the magnitude of complex workflow components and the volume of transaction data on a daily basis, there were many instances where an anomaly could occur and not be immediately obvious to an individual or mitigated before impacting customers.

The Challenge

Client needed to understand patterns of customer transactions within backend services

Malfunctions in system components were resulting in customer-impacting service interruptions

It was critical to automate understanding of when issues in the system occurred, because growth and complexity were driving increased operating effort and cost

Although the client had performed some studies to investigate the anomalous conditions in the log files, attempts to analyze the logs were not done in real time, analysis produced a high false positive rate, and no one had attempted to relate the logs to critical customer incidents programmatically. Additionally, the size and complexity of the log files made analysis of the data over time challenging.

The objective of the engagement was to apply our data science methodologies to investigate and identify which models yielded predictive power against the ESB logs, with the ultimate intent of serving market facing SLAs at lower cost and with lower reputational risk due to service outages.

Solution

SVDS executed an agile data science investigation to understand how customer facing incidents or disruptions manifested in the client's log files, and ultimately to detect these in the system before a customer raises an incident, as doing so would help reduce operational costs. Our team performed text processing on the semi-structured incident reports dataset to identify high-priority incidents. Various derivations of the raw log data were iteratively developed and evaluated for their ability to detect incidents. Both supervised and unsupervised machine learning techniques were applied to surface the patterns of normal and abnormal system behavior present in the data.

The client did not have an existing data science exploration and development environment to support the analysis. SVDS quickly instantiated a lightweight data science environment and ingested the log data in order to get tooling in place and expedite analysis. Components of this environment include:

- custom software for ingesting and parsing raw transaction logs
- a storage schema that enabled fast, flexible query processing
- a pipeline for extracting and engineering relevant features from the parsed log files
- a sampling strategy to prioritize the most customer-relevant transactions

Our approach served as a foundational proof of concept for extracting meaningful signal from the mass of noisy, semi-structured operational data captured from the ESB. We provided a systematic approach to data science to prove the underlying assumption that the message logs were indeed most promising at exhibiting anomalies that lead to incidents. It was a first-of-its-kind analysis of features over a longer

Our Approach

We performed text processing on the semi-structured incident reports data set to identify high-priority incidents

We applied supervised learning approaches to classify events; unsupervised approaches were applied to denoise events

We created a lightweight framework for ingesting and parsing raw transaction logs, and a storage schema that enabled fast, flexible query processing in Spark time window. Previously, internal studies were limited to small subsets of data and much shorter time windows.

SVDS provided a highly extensible, lightweight data science feature engineering and model evaluation framework that handles high volume and complexity of data, as well as transformation processes that unified and standardized inputs and outputs. This framework was a reproducible approach to the problem, and was easily transferable to other use cases and data sources. The insights uncovered by this project allowed the client to better understand the process of detecting customer-impacting incidents, and constituted a first step toward an operational solution.

While the initial phase of the project, limited to just a few easily available data sets, didn't yield strong enough signals to support full implementation, it did surface promising insights and signals. The client was very happy that the agile approach led to this discovery very quickly and efficiently, and gave them confidence that by adding additional contextual data, significant value was very likely to be captured.

New Capabilities

Can now correlate performance degradation with specific events in the log files

Further development nearly frictionless thanks to testbed for quickly prototyping and evaluating new analysis techniques

Line-of-sight to detecting system anomalies before a customer raises an incident, and reducing operational costs

MACHINE LEARNING APPROACH

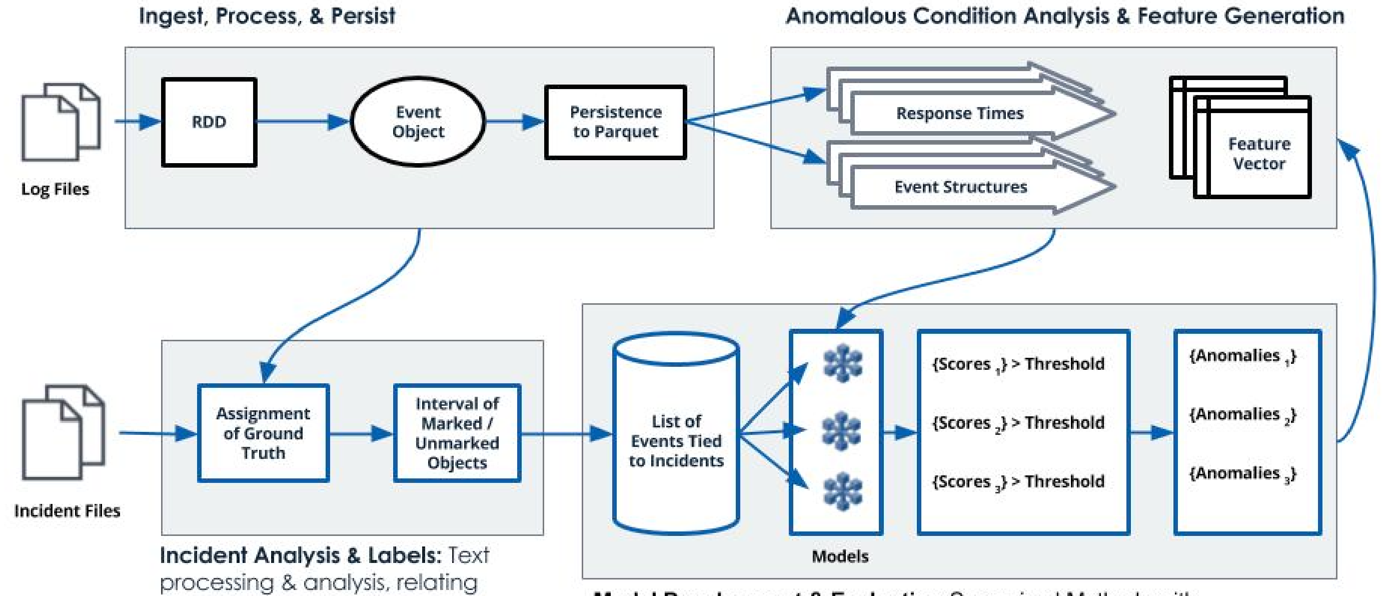
Ingest, Process, & Persist
Anomalous Condition Analysis & Feature Generation
Log Files
RDD
Event Object
Persistence to Parquet
Response Times
Event Structures
Feature Vector
Incident Files
Assignment of Ground Truth
Interval of Marked / Unmarked Objects
List of Events Tied to Incidents
Models
{Scores $_1$} > Threshold
{Scores $_2$} > Threshold
{Scores $_3$} > Threshold
{Anomalies $_1$}
{Anomalies $_2$}
{Anomalies $_3$}

Incident Analysis & Labels: Text processing & analysis, relating incident reports to event logs

Model Development & Evaluation: Supervised Methods with Unsupervised Methods to tune supervised approach