Time Series Anomaly Detection in Devo

Monitor your operations and act at scale

The Current State of Time Series Anomaly Detection

Time Series Anomaly Detection (TSAD) is the process of detecting abnormal behavior - anomalies - in time series data. Time series data is data that captures the value of a metric at a point of time - for example, number of errors in a TV stream, CPU utilization of a server, temperature recorded by a data center sensor, and more.

Traditional approaches to TSAD typically employ manually-set thresholds or simple mathematical calculations such as moving averages and standard deviations. Such approaches can work for small scale environments with dozens of series to monitor but fail as the complexity of the environment increases: imagine 10,000 routers with different application workloads varying by hour of day, day of week and geographic region. Many different “normals” will need to be defined and maintained.

Furthermore, traditional techniques create large numbers of false positives, fail to account for well understood and expected periodic behavior, and fail to learn new trends and behaviors as they establish themselves. This leads to wasted time and effort triaging systems simply because thresholds do not accurately capture the “normal” state of the system.

Enter Time Series Anomaly Detection in Devo

Devo solves the machine-scale TSAD problem by applying Machine Learning (ML) techniques to learn unique profiles of what is "normal" per metric being measured. Devo can be configured to monitor thousands of metrics concurrently. A metric can be as granular as a single device performance metric (e.g. CPU utilization every minute), or as broad as a global KPI across an entire infrastructure stack (e.g. number of devices online every 30 seconds). Each metric has a unique ML model and alerts are generated only when a specific metric begins behaving anomalously. Models are trained using an online algorithm - as new data arrives models are updated to learn the new normal. This allows Devo to capture dynamically evolving environments without the need for operator intervention.

Devo is uniquely positioned to offer TSAD capabilities due to its architecture: all historical and real time data is available in a single pane of glass and can be accessed at high speed. This powers Devo's ability to monitor thousands of time series metrics concurrently, learning for each what is normal expected behavior and detecting anomalies as they occur in real time.

When something anomalous does occur, custom actions can be triggered. The Devo platform provides a wide and extensible set of actions for a given alert, including notifications, triggering commands and generation of support tickets in external systems. Actions can be tiered by anomaly severity - slightly unusual metric values are assigned a lower severity than extreme spikes or drops.

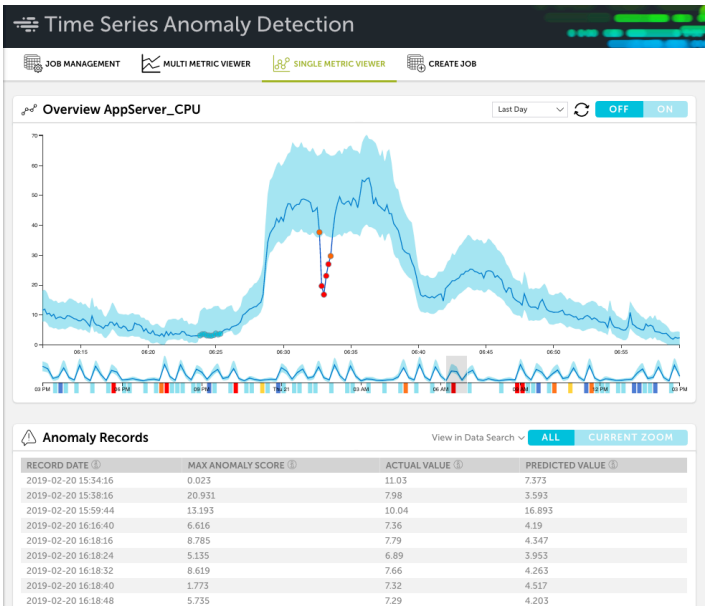

Figure 1: Time-Series Anomaly Detection Single-Metric

Operators are able to tune Devo's TSAD models by providing feedback on anomalies. Using human-in-the-loop techniques, anomalies can be marked as known anomalies so that they are not incorporated into real-time model updates. Operators can also trigger fast learning of a new environment if there is a known fundamental change in the underlying - e.g. a new system is brought online that doubles the value of a metric (e.g. network bandwidth utilization) in a day. This can be marked as the new normal and Devo will update thresholds at an accelerated pace.

Devo Time Series Anomaly Detection enables operators to identify anomalies and take action, whilst limiting false positives and reducing the human effort typically associated with configuring and maintaining anomaly monitoring thresholds.

Multi Metric Time Series Anomaly Detection

Modern application and infrastructure stacks consist of many individual components that work together to deliver a service. Monitoring individual components separately, while possible, does not provide comprehensive, holistic visibility into the service. To address this, Devo enables users to logically group components that comprise a service's stack, treating each component as a single metric with its own ML model and anomalous thresholds. Devo then aggregates all individual metric components, with configurable weightings for each, to derive an overall health score. With this capability, Devo empowers users to monitor complex interactions between components in a technology stack; operations teams can identify root cause failures, learn how issues propagate to affect other components in the stack, and receive alerts before anomalies threaten the operational integrity of the entire stack.

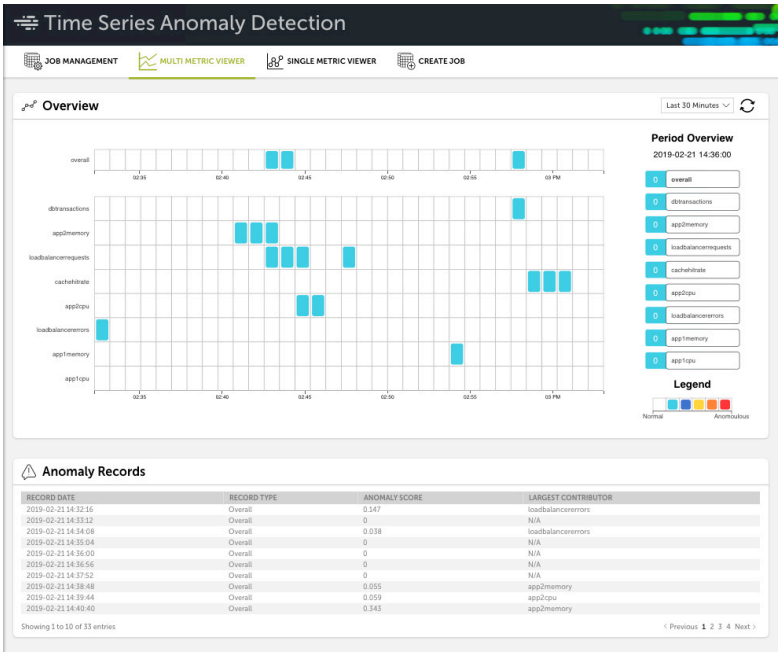

Figure 2: Time-Series Anomaly Detection Multi-Metric

Get started today

It's time to monitor everything and gain predictive insights at scale. No heavy-lifting required or advanced modeling skills are required. Devo for Time Series Anomaly Detection is included as a standard module within the Devo Data Operations platform.

ABOUT DEVO

Devo Technology is the data engine behind today's digitally-driven enterprises, helping organizations maximize the economic and operational value of their machine data. The Devo Data Operations Platform delivers real-time analytics on streaming and historical data to turn machine data into actions that help enterprises achieve sustained performance and growth. By collecting, enhancing and analyzing machine data, Devo provides business-driving insights for IT, security, and business teams at the world's largest organizations. For more information visit www.devo.com